

# The Administrator's Reference

2 Books  
in 1

# Windows Server 2016

Super-sized desktop reference  
Combines 2 training guides in 1 volume  
Expert advice at your fingertips

William R. Stanek  
Author

William R. Stanek Jr.  
Contributor



# Acknowledgments

To my readers—Thank you for being there with me through many books and many years.

To my wife—for many years, through many books, many millions of words, and many thousands of pages she's been there, providing support and encouragement and making every place we've lived a home.

To my kids—for helping me see the world in new ways, for having exceptional patience and boundless love, and for making every day an adventure.

To everyone I've worked with at Microsoft—thanks for the many years of support and for helping out in ways both large and small.

Special thanks to my son Will for his extensive contributions to this book. You've made many contributions previously, but now I can finally give you the cover credit you've earned and deserved for so long.

—William R. Stanek







# **Windows Server 2016: The Administrator's Reference**

**William R. Stanek**

Author & Series Editor

**William R. Stanek, Jr.**

Contributor







# Windows Server 2016: The Administrator's Reference

Published by Stanek & Associates  
PO Box 362, East Olympia, WA, 98540-0362  
[www.williamrstanek.com](http://www.williamrstanek.com).

Copyright © 2016 William R. Stanek. Seattle, Washington.  
All rights reserved.

No part of this book may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted by Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the publisher. Requests to the publisher for permission should be sent to the address listed previously.

Stanek & Associates is a trademark of Stanek & Associates and/or its affiliates. All other marks are the property of their respective owners. No association with any real company, organization, person or other named element is intended or should be inferred through use of company names, web site addresses or screens.

This book expresses the views and opinions of the author. The information contained in this book is provided without any express, statutory or implied warranties.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND DISCUSSION IN THIS BOOK MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING





PROFESSIONAL SERVICES AND THAT SHOULD PROFESSIONAL ASSISTANCE BE REQUIRED THE SERVICES OF A COMPETENT PROFESSIONAL SHOULD BE SOUGHT. NEITHER THE PUBLISHERS, AUTHORS, RESELLERS NOR DISTRIBUTORS SHALL BE HELD LIABLE FOR ANY DAMAGES CAUSED OR ALLEGED TO BE CAUSE EITHER DIRECTLY OR INDIRECTLY HEREFROM. THE REFERENCE OF AN ORGANIZATION OR WEBSITE AS A SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER OR THE AUTHOR ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR THE RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS BOOK MAY NOT BE AVAILABLE OR MAY HAVE CHANGED SINCE THIS WORK WAS WRITTEN.

Stanek & Associates publishes in a variety of formats, including print, electronic and by print-on-demand. Some materials included with standard print editions may not be included in electronic or print-on-demand editions or vice versa.

**Country of First Publication:** United States of America.

**Cover Design:** Creative Designs Ltd.

**Editorial Development:** Andover Publishing Solutions

**Technical Review:** L & L Technical Content Services

You can provide feedback related to this book by emailing the author at [williamstanek@aol.com](mailto:williamstanek@aol.com). Please use the name of the book as the subject line.

Version: 1.0.1.2b

**Note** I may periodically update this text and the version number shown above will let you know which version you are working with. If there's a specific feature you'd like me to write about in an update, message me on Facebook (<http://facebook.com/williamstanekauthor>). Please keep in mind readership of this book determines how much time I can dedicate to it.







# Table of Contents

## How to Use This Guide

[Print Readers](#)

[Digital Book Readers](#)

[Support Information](#)

[Conventions & Features](#)

[Share & Stay in Touch](#)

## Chapter 1. Welcome to Windows Server 2016

[Windows Server 2016 and Windows 10](#)

[Getting to Know Windows Sever 2016](#)

[Using Servers in Domains and Workplaces](#)

[Navigating Configuration Options](#)

[Using Nano Server and Containers](#)

[Nano Server Essentials](#)

[Deploying Nano Server](#)

[Container Essentials](#)

[Deploying Containers](#)

[Getting Started with Windows Server 2016](#)

## Chapter 2. Working with Windows Servers

[Preparing Networking](#)

[Setting Network Category and Sharing](#)

[Working with Networking Protocols](#)

[Navigating Management Options](#)

[Getting Started with Server Manager](#)

[Getting Started with Sconfig](#)

[Getting Started with Windows PowerShell](#)

[Managing Your Servers Remotely](#)

[Remote Management Essentials](#)

[Configuring Remote Management and Trusted Hosts](#)

[Adding Remote Servers for Management](#)

[Grouping Servers for Management](#)

[Working with Added Servers](#)

[Managing Roles and Features Working](#)

[with Servers Based on Roles](#)

[Adding Roles and Features](#)





[Removing Roles and Features](#)

[Using Windows PowerShell for Management](#)

[Getting Binaries for Roles and Features](#)

### **[Chapter 3. Configuring Server Settings](#)**

[Getting System Information](#)

[Licensing & Activation](#)

[Configuring Server Naming](#)

[Renaming a Server](#)

[Joining a Domain](#)

[Managing Processor Scheduling](#)

[Allocating Virtual Memory](#)

[Navigating Memory Baselines](#)

[Configuring the Paging File Manually](#)

[Managing the Paging File Automatically](#)

[Configuring Data Execution Prevention](#)

[Using and configuring DEP](#)

[Understanding DEP Compatibility](#)

[Configuring Environment Variables](#)

[Creating Environment Variables](#)

[Editing Environment Variables](#)

[Deleting Environment Variables](#)

[Configuring Startup and Recovery](#)

[Setting Startup Options](#)

[Setting Recovery Options](#)

### **[Chapter 4. Understanding Active](#)**

**[Directory](#)** [Connecting Network Resources](#)

[Working with On-Premises Deployments](#)

[Working with Hybrid Deployments](#)

[Working with Directory Services](#)

[Deploying Directory Services](#)

[Using Active Directory with DNS](#)

[Implementing Directory Services](#)

[Working with Domain Structures](#)

[Navigating Domain and Forest Structures](#)

[Using Organizational Units](#)





[Understanding Sites and Subnets](#)

[Understanding the Directory Structure](#)

[Exploring the Data Store](#)

[Exploring Global Catalogs](#)

[Universal Group Membership Caching](#)

[Replication and Active Directory](#)

[Active Directory and LDAP](#)

[Understanding Operations Master Roles](#)

## **[Chapter 5. Managing Active Directory](#)**

[Working with Active Directory Domains](#)

[Accessing Active Directory](#)

[Navigating Functional Levels](#)

[Tools for Working with Active Directory](#)

[Active Directory Administration Tools](#)

[Installing Domain Controllers](#)

[Setting up a Domain Controller](#)

[Demoting a Domain Controller](#)

[Managing Users, Groups and Computers](#)

[Using Active Directory Users and Computers](#)

[Searching the Directory](#)

[Using Administrative Center and Windows Powershell](#)

[Working with Organizational Units](#)

[Creating Organizational Units](#)

[Viewing and Editing Organizational Unit Properties](#)

[Renaming and Deleting Organizational Units](#)

[Moving Organizational Units](#)

## **[Chapter 6. Maintaining Active Directory](#)**

[Working with Sites](#)

[Creating Sites](#)

[Creating Subnets](#)

[Associating Domain Controllers with Sites](#)

[Configuring Site Links](#)

[Configuring Site Link Bridges](#)

[Maintaining Domain Controllers](#)

[Configuring the Functional Level](#)





[Working with Domainwide Roles](#)

[Working with Forestwide Roles](#)

[Transferring Roles by Using the Command Line](#)

[Configuring Global Catalogs](#)

[Configuring Universal Group Membership Caching](#)

[Using the Active Directory Recycle Bin](#)

[Preparing Schema for the Recycle Bin](#)

[Recovering Deleted Objects](#)

[Using Windows Powershell for Basic and Advanced Recovery](#)

[Using the Enhanced Recycle Bin for Recovery](#)

[Troubleshooting Active Directory](#)

[Using AdsI Edit](#)

[Examining Intersite Topology](#)

[Diagnosing and Resolving Issues](#)

[Seizing Roles By Using the Command Line](#)

## **Chapter 7. Accounts: The Essentials**

[Managing Access to Network Resources](#)

[Navigating Authentication Protocols](#)

[Using Access Controls](#)

[Using Claims-Based Access Controls](#)

[Working with Central Access Policies](#)

[Navigating User and Group Accounts](#)

[Working with User Accounts](#)

[Working with Group Accounts](#)

[Granting Privileges, Rights and Permissions](#)

[Using the Default User Accounts](#)

[Working with Built-in User Accounts](#)

[Working with Predefined User Accounts](#)

[Using Other Default Accounts](#)

[Using the Default Group Accounts](#)

[Working with the Admin Groups](#)

[Working with Implicit Groups](#)

## **Chapter 8. Managing Account Policies**

[Understanding Account Naming Policies](#)

[Using Password and Account Policies](#)





[Using Secure Passwords](#)

[Setting Account Policies](#)

[Configuring Password Policies](#)

[Enforce Password History](#)

[Maximum Password Age](#)

[Minimum Password Age](#)

[Minimum Password Length](#)

[Password Must Meet Complexity Requirements](#)

[Store Passwords Using Reversible Encryption](#)

[Configuring Account Lockout Policies](#)

[Account Lockout Threshold](#)

[Account Lockout Duration](#)

[Reset Account Lockout Counter After](#)

[Configuring Kerberos Policies](#)

[Enforce User Logon Restrictions](#)

[Maximum Lifetime](#)

[Maximum Tolerance](#)

[Configuring Authentication Policies](#)

[Getting Started with Authentication Policies](#)

[Implementing Authentication Policy](#)

[Configuring User Rights Policies](#)

[Configuring Global User Rights](#)

[Configuring Local User Rights](#)

## **[Chapter 9. Creating Accounts](#)**

[Adding User Accounts](#)

[Creating Domain Users](#)

[Creating Local Users](#)

[Adding Groups](#)

[Creating Global Groups](#)

[Creating Local Groups](#)

[Adding Computer Accounts](#)

[Creating Standard and Managed Computer Accounts](#)

[Creating Computer Accounts in Administrative Center](#)

[Managing User Contact Information](#)

[Setting Contact Information](#)







[Searching for Users and Groups](#)

[Managing Group Membership](#)

[Managing Membership in Groups](#)

[Adding or Removing Multiple Accounts](#)

[Setting the Primary Group for Users and Computers](#)

## **[Chapter 10. Working with Managed Accounts](#)**

[Using Managed Accounts](#)

[Implementing Managed Accounts](#)

[Creating Managed Accounts](#)

[Selecting Target Computers](#)

[Installing Managed Accounts](#)

[Configuring Services to Use Managed Accounts](#)

[Maintaining Managed Accounts](#)

[Moving Managed Accounts](#)

[Removing Managed Accounts](#)

[Using Virtual Accounts](#)

## **[Chapter 11. Managing Computers, Users and Groups](#)**

[Maintaining Computer Accounts](#)

[Resetting Locked Computer Accounts](#)

[Moving Computer Accounts](#)

[Disabling and Deleting Computer Accounts](#)

[Joining a Computer to a Domain or Workgroup](#)

[Using Offline Domain Join](#)

[Configuring Paths and Profiles](#)

[Using Environment Variables in Paths](#)

[Using Logon Scripts](#)

[Assigning Home Directories](#)

[Controlling Logon and Access](#)

[Restricting Logon Times](#)

[Restricting Workstation Logons](#)

[Configuring VPN Privileges and Remote Access](#)

[Managing Security Options](#)

[Working with User Profiles](#)

[Using Roaming and Mandatory Profiles](#)

[Restricting Roaming Profiles](#)





[Creating User Profiles](#)

[Managing Local Profiles](#)

[Changing the Profile Type](#)

[Recovering from Profile Problems](#)

[Maintaining Users and Groups](#)

[Copying Domain User Accounts](#)

[Importing and Exporting Accounts](#)

[Renaming Users and Groups](#)

[Managing Multiple Accounts](#)

[Changing and Resetting Passwords](#)

[Enabling User Accounts](#)

[Setting Active Directory Permissions](#)

[Deleting Users and Groups](#)

[Troubleshooting Logon Problems](#)

## **[Chapter 12. Maintaining Your Servers](#)**

[Working with Services](#)

[Navigating Services in Server Manager](#)

[Navigating Services in Computer Management](#)

[Tracking Service Status](#)

[Configuring Service Startup](#)

[Configuring Service Logon](#)

[Configuring Service Recovery](#)

[Disabling Unnecessary Services](#)

[Troubleshooting Services](#)

[Managing System Processes](#)

[Working with Processes](#)

[Troubleshooting Processes](#)

[Tracking System and Security Problems](#)

[Accessing Events in Server Manager](#)

[Accessing Events in Event Viewer](#)

[Filtering Event Logs](#)

[Configuring Event Logging](#)

[Clearing Event Logs](#)

[Archiving the Event Logs](#)

[Using Archiving](#)





[Creating Log Archives](#)

[Viewing Log Archives](#)

## **[Chapter 13. Optimizing Server Performance](#)**

### [Tracking Performance](#)

[Understanding Processor Utilization](#)

[Tracking Memory Allocation and Usage](#) [Understanding Network Utilization](#)

[Tracking Remote User Sessions](#)

### [Preparing for Monitoring](#)

[Creating a Monitoring Plan](#)

[Using the Monitoring Consoles](#)

[Choosing Counters to Monitor](#)

### [Performance Logging](#)

[Creating and Managing Data Collector Sets](#)

[Collecting Counter Data](#)

[Collecting Trace Data](#)

[Collecting Configuration Data](#)

[Viewing Data Collector Reports](#)

[Configuring Performance Counter Alerts](#)

### [Fine-Tuning System Performance](#)

[Tuning Physical and Virtual Memory](#)

[Tuning Processors and Disks](#)

[Tuning Networking](#)

## **[Chapter 14. Using Group Policy](#)**

### [Working with Group Policy](#)

[Understanding Processing Order](#)

[Understanding Policy Application](#)

[Understanding Policy Versioning](#)

### [Managing Local Group Policies](#)

[Using Group Policy Objects](#)

[Accessing the Top-Level Local Policy Settings](#)

[Accessing Tailored Local Group Policies](#)

### [Navigating Enterprise Policies](#)

[Getting Started with GPMC](#)

[Using Policy Editors](#)

[Understanding Default Policies](#)





[Understanding Registry-based Settings](#)

[Editing GPOs](#)

[Using Administrative Templates to Set Policies](#)

[Creating and Linking GPOs](#)

[Using Starter GPOs](#)

[Delegating Group Policy Management](#)

[Checking Delegations for Creating GPOs](#)

[Checking Delegations for Management](#)

[Checking Delegations for a GPO](#)

## **[Chapter 15. Maintaining and Troubleshooting Group Policy](#)**

[Optimizing Group Policy](#)

[Blocking, Overriding, and Disabling Policies](#)

[Refreshing Group Policy](#)

[Configuring the Refresh Interval](#)

[Refreshing Policy Manually](#)

[Modeling Group Policy for Planning](#)

[Determining Current Settings and Refresh Status](#)

[Maintaining Group Policy Objects](#)

[Copying, Pasting, and Importing Policy Objects](#)

[Backing Up and Restoring Policy Objects](#)

[Disabling Part of a GPO](#)

[Changing Processing Preferences](#)

[Configuring Slow-Link Detection](#)

[Removing Links and Deleting GPOs](#)

[Troubleshooting Group Policy](#)

[Diagnosing Group Policy Issues](#)

[Fixing Default Group Policy Objects](#)

## **[Chapter 16. Optimizing Server Security](#)**

[Locking Down Your Servers](#)

[Managing Security Configurations](#)

[Accessing and Preparing the Security Tools](#)

[Optimizing Security Settings](#)

[Changing Settings for Core Policies](#)

[Configuring Restricted Groups](#)

[Restricting System Services](#)





- [Restricting Registry and File System Paths](#)
- [Reviewing, and Applying Security Settings](#)
- [Locking Down Multiple Computers](#)
- [Preparing for Configuration Deployment](#)
- [Deploying Security Configurations](#)

## **Chapter 17. Deploying Windows Server 2016**

- [Preparing for Windows Server 2016](#)
  - [Working with Windows Server 2016](#)
  - [Managing Server Core](#)
  - [Managing Nano Server](#)
- [Performing a Clean Installation](#)
- [Performing an Upgrade Installation](#)
- [Performing Additional Tasks During Installation](#)
  - [Using the Command Line During Installation](#)
  - [Forcing Disk Partition Removal During Installation](#)
  - [Loading Device Drivers During Installation](#)
- [Managing Disk Partitions During Installation](#)
  - [Creating Disk Partitions During Installation](#)
  - [Formatting Disk Partitions During Installation](#)
  - [Deleting Disk Partitions During Installation](#)
  - [Extending Disk Partitions During Installation](#)

## **Chapter 18. Implementing TCP/IP Networking**

- [Navigating Networking Options](#)
  - [Networking Tools & Essentials](#)
  - [Network Interfaces & Teaming](#)
- [Configuring TCP/IP Networking](#)
  - [Configuring Static IP Addresses](#)
    - [Using the Ping Command to Check an Address](#)
    - [Configuring a Static IPv4 or IPv6 Address](#)
  - [Configuring Dynamic IP Addresses and Alternate IP Addressing](#)
  - [Configuring Multiple Gateways](#)
  - [Configuring Networking for Hyper-V](#)
- [Using NIC Teaming](#)
  - [Understanding NIC Teaming](#)
  - [Configuring NIC Teaming](#)





## **Chapter 19. Data Storage: The Essentials**

### Deploying File Services

#### Adding Drives

##### Physical Drives

##### Preparing a Physical Drive for Use

##### The MBR and GPT Partition Styles

##### Legacy and Protective Mbrs

##### Disk Types and File Systems

##### Using Disk Management

##### Using Removable Storage Devices

##### Installing and Checking for a New Drive

##### Understanding Drive Status

#### Working with Disks

##### Disk Management Essentials

##### Special Considerations for Basic and Dynamic Disks

##### Changing Drive Types

##### Converting a Basic Disk to a Dynamic Disk

##### Changing a Dynamic Disk Back to a Basic Disk

##### Reactivating Dynamic Disks

##### Rescanning Disks

##### Moving a Dynamic Disk to a New System

#### Managing Virtual Hard Disks

##### Creating Virtual Hard Disks

##### Working with Virtual Hard Disks

## **Chapter 20. Partitioning and Optimizing Drives**

### Using Basic Disks and Partitions

##### Partitioning Essentials

##### Creating Partitions and Simple Volumes

##### Formatting Partitions

#### Using Volumes and Volume Sets

##### Understanding Volume Basics

##### Understanding Volume Sets

##### Creating Volumes and Volume Sets

##### Deleting Volumes and Volume Sets

##### Managing Volumes





## Compressing Drives and Data

[Compressing Drives](#)

[Compressing Directories and Files](#)

[Expanding Compressed Drives](#)

[Expanding Compressed Directories and Files](#)

## Encrypting Drives and Data

[Understanding Encryption and the Encrypting File System](#)

[Encrypting Directories and Files](#)

[Working with Encrypted Files and Folders](#)

[Configuring Recovery Policies](#)

[Decrypting Files and Directories](#)

## **Chapter 21. Using TPM and BitLocker Drive Encryption**

### Creating Trusted Platforms

[Navigating TPM Essentials](#)

[TPM Management and Policies](#)

[Enabling TPM in Firmware](#)

[Initializing and Preparing a TPM for First Use](#)

[Clearing the TPM](#)

[Changing the TPM Owner Password](#)

### BitLocker Drive Encryption: the Essentials

[Understanding BitLocker Drive Encryption](#)

[Hardware Encryption, Secure Boot, and Network Unlock](#)

[Deploying BitLocker Drive Encryption](#)

### Managing BitLocker Drive Encryption

[Preparing for BitLocker Drive Encryption](#)

[Enabling BitLocker on Nonsystem Volumes](#)

[Enabling BitLocker on USB Flash Drives](#)

[Enabling BitLocker on System Volumes](#)

[Managing and Troubleshooting BitLocker](#)

[Recovering Data Protected by BitLocker Drive Encryption](#)

[Disabling or Turning Off BitLocker Drive Encryption](#)

## **Chapter 22. Using Storage Spaces**

### Getting Started with Storage Spaces

[Grouping Disks into Storage Pools](#)

[Storage Pool Options, Versions and Features](#)





## [Creating Storage Pools](#)

[Using Drives and Volumes in a Storage Pool](#)

[Creating a Storage Pool](#)

## [Creating Virtual Disks and Volumes](#)

[Creating a Virtual Disk in a Storage Space](#)

[Creating a Standard Volume](#)

## [Troubleshooting Storage Spaces](#)

## [\*\*Chapter 23. Using RAID\*\*](#)

### [Implementing RAID](#)

[Implementing RAID-0: Disk Striping](#)

[Implementing RAID-1: Disk Mirroring](#)

[Creating a Mirror Set in Disk Management](#)

[Mirroring an Existing Volume](#)

[Implementing RAID-5: Disk Striping with Parity](#)

[Creating a Striped Set with Parity in Disk Management](#)

### [Managing RAID and Recovering From Failures](#)

[Breaking a Mirrored Set](#)

[Resynchronizing and Repairing a Mirrored Set](#)

[Repairing a Mirrored System Volume to Enable Boot](#)

[Removing a Mirrored Set](#)

[Repairing a Striped Set Without Parity](#)

[Regenerating a Striped Set with Parity](#)

## [\*\*Chapter 24. Maintaining Partitions and Drives\*\*](#)

### [Drive Maintenance Essentials](#)

[Assigning Drive Letters and Paths](#)

[Changing or Deleting the Volume Label](#)

[Resizing Partitions and Volumes](#)

[Deleting Partitions and Drives](#)

### [Converting a Volume to NTFS](#)

[Understanding the Convert Utility Syntax](#)

[Using the Convert Utility](#)

### [Repairing Disks](#)

[Checking Disks Manually](#)

[Running Check Disk Interactively](#)

[Analyzing and Optimizing Disks](#)







## **Chapter 25. Implementing File Sharing**

### Working with SMB

SMB Essentials

SMB Encryption and Bandwidth Limits

SMB Replication

Public Folder Sharing

### Configuring Standard File Sharing

Viewing Existing Shares

Creating Shares in Computer Management

Creating Shares in Server Manager

Changing Share Settings

### Managing Share Permissions

Understanding the Various Share Permissions

Viewing and Configuring Share Permissions

### Connecting to Network Drives

Mapping a Network Drive

Disconnecting a Network Drive

### Managing Existing Shares

Understanding Special Shares

Connecting to Special Shares

Viewing User and Computer Sessions

Managing Sessions and Shares

Managing Open Resources

Stopping File and Folder Sharing

### Configuring NFS Sharing

## **Chapter 26. Using Shadow Copies and Work Folders**

### Using Shadow Copies

Creating Shadow Copies

Restoring a Shadow Copy

Reverting an Entire Volume

Deleting Shadow Copies

Disabling Shadow Copies

### Configuring Synced Sharing

Getting Started with Work Folders

Creating Sync Shares and Enabling SMB Access





[Accessing Work Folders on Clients](#)

## **[Chapter 27. Managing Permissions and Auditing](#)**

[Object Management, Ownership, and Inheritance](#)

[Objects and Object Managers](#)

[Object Ownership and Transfer](#)

[Object Inheritance](#)

[File and Folder Permissions](#)

[Understanding File and Folder Permissions](#)

[Setting Basic File and Folder Permissions](#)

[Setting Special Permissions on Files and Folders](#)

[Setting Claims-Based Permissions](#)

[Auditing System Resources](#)

[Setting Auditing Policies](#)

[Auditing Files and Folders](#)

[Auditing the Registry](#)

[Auditing Active Directory Objects](#)

## **[Chapter 28. Configuring Disk Quotas](#)**

[Using NTFS Disk Quotas](#)

[Understanding NTFS Disk Quotas](#)

[Enabling Quotas Through Group Policy](#)

[Enabling Quotas Through Disk Properties](#)

[Viewing Disk Quota Entries](#)

[Creating Disk Quota Entries](#)

[Deleting Disk Quota Entries](#)

[Exporting and Importing NTFS Disk Quota Settings](#)

[Disabling NTFS Disk Quotas](#)

[Resource Manager Disk Quotas](#)

[Understanding Resource Manager Disk Quotas](#)

[Managing Disk Quota Templates](#)

[Creating Resource Manager Disk Quotas](#)

## **[Chapter 29. Using Group Policy for Administration](#)**

[Managing Automatic Updates in Group Policy](#)

[Configuring Automatic Updates](#)

[Optimizing Automatic Updates](#)

[Using Intranet Update Service Locations](#)





## User and Computer Script Management

Assigning Computer Startup and Shutdown Scripts

Assigning User Logon and Logoff Scripts

## Centrally Managing Special Folders

Redirecting a Special Folder to a Single Location

Redirecting a Special Folder Based on Group Membership Removing

Redirection

## Deploying Software Through Group Policy

Getting to Know Software Installation Policy

Deploying Software Throughout Your Organization

Configuring Software Deployment Options

Updating Deployed Software

Upgrading Deployed Software

## Automatically Configuring Work Folders

## Automatically Enrolling Computer and User Certificates

## **Chapter 30. Implementing Print Services**

### Managing the Print and Document Services Role

Using Print Devices

Printing Essentials

Configuring Print Servers

Enabling and Disabling File and Printer Sharing

### Getting Started with Print Management

#### Installing Printers

Using Autoinstall

Installing Physically-Attached Print Devices

Installing Network-Attached Print Devices

Connecting to Network Printers

#### Managing Installed Printers

Deploying Printer Connections

Configuring Point and Print Restrictions

Moving Printers to a New Print Server

Monitoring Printers and Printer Queues

Solving Spooling Problems

## **Chapter 31. Configuring and Maintaining Print Services**

### Configuring Print Server Properties





[Locating the Spool Folder and Enabling Printing](#)

[Managing High-Volume Printing](#)

[Enabling Print Job Error Notification](#)

[Managing Printer Drivers](#)

[Updating a Printer Driver](#)

[Configuring Drivers for Clients](#)

[Configuring Printer Properties](#)

[Adding Comments and Locations](#)

[Listing Printers in Active Directory](#)

[Setting Separator Pages and Print Device Modes](#)

[Changing the Printer Port](#)

[Starting and Stopping Printer Sharing](#)

[Setting Printer Access Permissions](#)

[Setting Document Defaults](#)

[Scheduling and Prioritizing Print Jobs](#)

[Scheduling Printer Availability and Priority](#)

[Configuring Print Spooling](#)

[Managing Print Jobs](#)

[Viewing Printer Queues and Print Jobs](#)

[Pausing the Printer and Resuming Printing](#)

[Emptying the Print Queue](#)

[Pausing, Resuming, and Restarting Document Printing](#)

[Removing Documents and Canceling Print Jobs](#)

[Checking Document Properties](#)

[Setting Document Priority](#)

[Scheduling Document Printing](#)

[Auditing Print Jobs](#)

## **Chapter 32. Implementing DHCP**

[Understanding DHCP](#)

[Using Dynamic IPv4 Addressing and Configuration](#)

[Using Dynamic IPv6 Addressing and Configuration](#)

[Checking IP Address Assignment](#)

[Understanding Scopes](#)

[Installing a DHCP Server](#)

[Installing DHCP Components](#)



|||||

[Starting and Using the DHCP Console](#)

[Connecting to Remote DHCP Servers](#)

[Starting and Stopping a DHCP Server](#)

[Authorizing a DHCP Server in Active Directory](#)

[Configuring DHCP Servers](#)

[Configuring Server Bindings](#)

[Updating DHCP Statistics](#)

[Auditing and Troubleshooting DHCP](#)

[Integrating DHCP and DNS](#)

[Avoiding IP Address Conflicts](#)

[Saving and Restoring the DHCP Configuration](#)

## **[Chapter 33. Managing and Maintaining DHCP](#)**

[Creating and Managing Scopes](#)

[Creating Normal Scopes for IPv4 Addresses](#)

[Creating Normal Scopes for IPv6 Addresses](#)

[Creating Multicast Scopes](#)

[Setting Scope Options](#)

[Assigning Server Options](#)

[Assigning Scope Options](#)

[Assigning Reservation Options](#)

[Modifying Scopes](#)

[Activating and Deactivating Scopes](#)

[Enabling the Bootstrap Protocol](#)

[Removing a Scope](#)

[Configuring Multiple Scopes on a Network](#)

[Creating and Managing Failover Scopes](#)

[Creating Failover Scopes](#)

[Modifying or Removing Failover Scopes](#)

[Creating and Managing Superscopes](#)

[Creating Superscopes](#)

[Adding Scopes to a Superscope](#)

[Removing Scopes From a Superscope](#)

[Activating and Deactivating a Superscope](#)

[Deleting a Superscope](#)

[Managing Addressing](#)

|||||



[Viewing Scope Statistics](#)

[Enabling and Configuring MAC Address Filtering](#)

[Setting a New Exclusion Range](#)

[Reserving DHCP Addresses](#)

[Modifying Reservation Properties](#)

[Deleting Leases and Reservations](#)

[Backing Up and Restoring the DHCP Database](#)

[Backing Up the DHCP Database](#)

[Restoring the DHCP Database From Backup](#)

[Using Backup and Restore to Move the DHCP Database to a New Server](#)

[Forcing the DHCP Server Service to Regenerate the DHCP Database](#)

[Reconciling Leases and Reservations](#)

## **Chapter 34. Implementing DNS**

[Understanding DNS](#)

[Integrating Active Directory and DNS](#)

[Enabling DNS on the Network](#)

[Configuring Name Resolution on DNS Clients](#)

[Installing DNS Servers](#)

[Installing and Configuring the DNS Server Service](#)

[Configuring a Primary DNS Server](#)

[Configuring a Secondary DNS Server](#)

[Configuring Reverse Lookups](#)

[Configuring Global Names](#)

## **Chapter 35. Managing and Maintaining DNS**

[Managing DNS Servers](#)

[Adding and Removing Servers to Manage](#)

[Starting and Stopping a DNS Server](#)

[Using DNSSEC and Signing Zones](#)

[Creating Child Domains Within Zones](#)

[Creating Child Domains in Separate Zones](#)

[Deleting a Domain or Subnet](#)

[Managing DNS Records](#)

[Adding Address and Pointer Records](#)

[Adding a PTR Record Later](#)

[Adding DNS Aliases with CNAME](#)





[Adding Mail Exchange Servers](#)

[Adding Name Servers](#)

[Viewing and Updating DNS Records](#)

[Updating Zone Properties and the SOA Record](#)

[Modifying the SOA Record](#)

[Allowing and Restricting Zone Transfers](#)

[Notifying Secondaries of Changes](#)

[Setting the Zone Type](#)

[Enabling and Disabling Dynamic Updates](#)

[Managing DNS Server Configuration and Security](#)

[Enabling and Disabling IP Addresses for a DNS Server](#)

[Controlling Access to DNS Servers Outside the Organization](#)

[Creating Nonforwarding and Forwarding-Only Servers](#)

[Creating Forwarding Servers](#)

[Configuring Conditional Forwarding](#)

[Enabling and Disabling Event Logging](#)

[Using Debug Logging to Track DNS Activity](#)

[Monitoring a DNS Server](#)

[\*\*Index\*\*](#)









# How to Use This Guide

This guide combines the complete text of Windows Server 2016: Essentials for Administration and Windows Server 2016: Server Infrastructure, providing 35 chapters and 250,000 words of in-depth insights into Microsoft's latest server operating system. The first chapter, Welcome to Windows Server 2016, provides the essential context for working with the product. Chapters detailing how to put Windows Server 2016 to work follow.

William Stanek has been developing expert solutions and writing professionally for Microsoft Windows Server since 1995. In this book, William shares his extensive knowledge of the product, delivering ready answers for day-to-day management and zeroing in on core commands and techniques.

This book is written especially for IT professionals working with, supporting Windows Server 2016, including nano server, core and full GUI deployments.





## Print Readers

Print editions of this book include an index and some other elements not available in the digital edition. Updates to this book are available online. Visit <http://www.williamrstanek.com/server2016admin/> to get any updates. This content is available to all readers.







## Digital Book Readers

Digital editions of this book are available at all major retailers, at libraries upon request and with many subscription services. If you have a digital edition of this book that you downloaded elsewhere, such as a file sharing site, you should know that the author doesn't receive any royalties or income from such downloads. Already downloaded this book or others? Donate here to ensure William can keep writing the books you need:

[https://www.paypal.com/cgi-bin/webscr?cmd=\\_sxclick&hosted\\_button\\_id=CPSBGLZ35AB26](https://www.paypal.com/cgi-bin/webscr?cmd=_sxclick&hosted_button_id=CPSBGLZ35AB26)





## Support Information

Every effort has been made to ensure the accuracy of the contents of this book. As corrections are received or changes are made, they will be added to the online page for the book available at:

<http://www.williamrstanek.com/server2016admin/>

If you have comments, questions, or ideas regarding the book, or questions that are not answered by visiting the site above, send them via e-mail to:

[williamstanek@aol.com](mailto:williamstanek@aol.com)

Other ways to reach the author:

Facebook: <http://www.facebook.com/William.Stanek.Author>

Twitter: <http://twitter.com/williamstanek>

It's important to keep in mind that Microsoft software product support is not offered. If you have questions about Microsoft software or need product support, please contact Microsoft.

Microsoft also offers software product support through the Microsoft Knowledge Base at:

<http://support.microsoft.com/>

## Conventions & Features

This book uses a variety of elements to help keep the text clear and easy to follow. You'll find code terms and listings in monospace, except when I tell you to actually enter or type a command. In that case, the command appears in **bold**. When I introduce and define a new term, I put it in italics.

The first letters of the names of menus, dialog boxes, user interface elements, and commands are capitalized. Example: the Add Roles And Features Wizard. This book also has notes, tips and other sidebar elements that provide additional details on points that need emphasis.

Keep in mind that throughout this book, where William has used click, right-click and double-click, you also can use touch equivalents: tap, press and hold, and double tap. Also, when using a device without a physical keyboard, you are able to enter text by using the onscreen keyboard. If a device has no physical keyboard, simply touch an input area on the screen to display the onscreen keyboard.





## Share & Stay in Touch

The marketplace for technology books has changed substantially over the past few years. In addition to becoming increasingly specialized and segmented, the market has been shrinking rapidly, making it extremely difficult for books to find success. If you want William to be able to continue writing and write the books you need for your career, raise your voice and support his work.

Without support from you, the reader, future books by William will not be possible. Your voice matters. If you found the book to be useful, informative or otherwise helpful, please take the time to let others know by sharing about the book online.

To stay in touch with William, visit him on Facebook or follow him on Twitter. William welcomes messages and comments about the book, especially suggestions for improvements and additions. If there is a topic you think should be covered in the book, let William know.





# Chapter 1. Welcome to Windows Server 2016

Windows Server 2016 is the latest version of the venerable server operating system from Microsoft. The operating system includes many features and options for deployment and management that weren't previously available, including:

- [Windows Server Containers](#)
- [Hyper-V Containers](#)
- [Nano Server](#)
- [Storage Spaces Direct](#)
- [Storage Replica](#)
- [And much more](#)

The way you approach Windows Server 2016 will depend on your background and your implementation plans. This chapter covers getting started with the operating system and explores the extent to which the changes affect how you work with and manage servers. Although this book focuses on administration, the tips and techniques presented can help anyone who supports, develops for, or works with the operating system.




# Windows Server 2016 and Windows 10

Windows Server 2016 and Windows 10 share a number of common features. These features have a common code base and extend across many areas of the operating systems, including management, security, networking, and storage. Because of this, you can apply much of what you know about Windows 10 to Windows Server 2016.

In Windows 10 and Windows Server 2016, features with common code bases have identical management interfaces. This is true for both Settings and Control Panel. Of course, exceptions exist in some cases for standard default settings. For example, Windows Server 2016 does not use Sleep or related states and Windows servers do not have sleep, hibernate, or resume functionality. Additionally, as you typically do not want to use extended power management options on servers, Windows Server 2016 has a limited set of power options.

Like Windows 10, Windows Server 2016 uses User Account Control (UAC). Through UAC, all applications are run by using either standard user or administrator privileges, and a security prompt is displayed by default whenever you run an application that requires administrator privileges. However, security prompts are not displayed by default when you make changes to Windows settings, which is different from the default settings in Windows 10.



As with Windows 10, Microsoft is continuing to transition the management of system properties from Control Panel to the Settings utility. This means you'll increasingly



rely on Settings to configure system properties. You can open Settings by clicking Start and then clicking the Settings button.



When working with Settings, it's important to remember that certain options aren't available when you are logged in using a standard user account. Although it would be handy to be able to start Settings with administrator privileges, you can't and there isn't a workaround to allow you to do so at the time of this writing.



That said, some options that aren't available in Settings when you are using a standard user account continue to be available in Control Panel. For example, although you can only work with accounts in Settings when you are logged in with an administrator account, Control Panel continues to provide account management options regardless of whether you are using a standard user or administrator account.



# Getting to Know Windows Server 2016

The Windows Server 2016 operating system is available in several different editions, including Standard and Datacenter. All Windows Server 2016 editions have the same basic hardware requirements:

- A 64-bit processor that supports NX and DEP
- 512 MB RAM for a minimal installation, with at least 2 GB preferred for a standard deployment
- 32 GB of disk space for the system partition
- A 1-Gigabit Ethernet adapter
- A DVD drive (if you plan to install the operating system using DVD media)
- TPM 2.0 (if you plan to use TPM or features that require TPM, such as BitLocker Drive Encryption)





## Using Servers in Domains and Workplaces

Because the various server editions support the same core features and administration tools, you can use the techniques discussed in this book regardless of which Windows Server 2016 edition you're using. When you install a Windows Server 2016 system, you configure the system according to its role on the network, as the following guidelines describe:

- Servers are generally assigned to be part of a workgroup or a domain.
- Workgroups are loose associations of computers in which each individual computer is managed separately.
- Domains are collections of computers you can manage collectively by means of domain controllers, which are Windows Server 2016 systems that manage access to the network, to the directory database, and to shared resources.


**NOTE** In this book, Windows Server 2016 and Windows Server 2016 family refer to all editions of Windows Server 2016. The various server editions support the same core features and administration tools.

Windows 10 and Windows Server 2016 also support a workplace configuration. A workplace is a loose association of computers that grants access to certain internal network resources and business apps. Workplaces have specific benefits:


- If users have Windows 10 devices from which they want to access corporate resources, those devices can use a workplace configuration to remotely connect to a computer at work. Users can then change their network passwords and connect to internal websites. Users also can use workplaces as an alternative way to access Outlook Web Access (OWA). Here, users connect to the workplace and then access OWA by using the internal URL (rather than an external URL).
- If administrators have Windows 10 devices from which they want to access corporate resources, those devices can use a workplace configuration to remotely connect to servers and perform assigned tasks by using administration links.

You implement workplaces by installing the Windows Server Essentials Experience role on servers running Windows Server 2016. Because the Windows Server Essentials Experience role also helps you connect servers to cloud-based applications and services, you can use the role to extend the functionality of your servers to the cloud and integrate servers with Office 365 and other Microsoft Online services.






Windows 10 and Windows Server 2016 also support Work Folders, a similar but distinctly different feature from workplaces. Work Folders allow users to synchronize their corporate data to their devices and vice versa. Those devices can be joined to the corporate domain or a workplace. To deploy Work Folders, you add the File And Storage Services > File and iSCSI Services > Work Folders role service to a server and then configure Work Folders by using Server Manager.



## Navigating Configuration Options

Before you deploy Windows Server 2016, you should carefully plan the server architecture and look closely at the deployment options and feature enhancements to determine whether and how they can be put to use in your organization. Windows Server 2016 is designed to provide optimal performance for server-related tasks and can be used in two basic configurations:

- **Server Core installation** A minimal installation option that provides a fixed subset of roles but does not include the Desktop Experience. You can configure a Server Core installation with a limited set of roles. A limited user interface is provided for managing the server, and most management is done locally at a command prompt with Windows PowerShell or remotely by using management tools. This installation option is ideally suited to situations in which you want to dedicate servers to a specific server role or combination of roles. Because additional functionality is not installed, the overhead caused by other services is reduced, providing more resources for the dedicated role or roles. Generally, the limited interface also is inherently more secure than other installation types.
- **Server With Desktop Experience** An installation option that provides full functionality—also referred to as a full-server installation . You can configure a server to have any allowed combination of roles, role services, and features, and a full user interface is provided for managing the server locally. The full installation option provides the most dynamic solution and is recommended for deployments of Windows Server 2016 in which the server role might change over time.



You choose the installation type during installation of the operating system. However, in a significant change from earlier releases of Windows Server, you cannot change the installation type after you've installed a server.



While user interfaces and graphical tools are not installed with a core installation, you can manage a core installation remotely using graphical administration tools. For local or remote administration, you also can use the command prompt with Windows PowerShell.


With a full-server installation, the full server graphical shell is installed, including:

- User interfaces
- The graphical management tools
- Ink and handwriting services
- Server media foundation

When you are working with a full-server installation the Desktop Experience is installed as well. The Desktop Experience provides Windows desktop functionality on the server. Windows features added include Windows Media Player, desktop themes, Video for Windows (AVI support), Disk Cleanup, Sync Center, Sound Recorder, Character Map, and Snipping Tool. Although these features allow a server to be used like a desktop computer, they can reduce the server's overall performance.

Whether you are working with a full or core installation, Windows Defender is installed automatically to help protect your servers against viruses, spyware, rootkits, and other types of malware. Rootkit detection helps to safeguard computers from malware that inserts itself into non-Microsoft drivers. If Windows Defender detects that a nonMicrosoft driver has been infected, it prevents the driver from starting. Microsoft drivers are protected at startup as part of other security features.

When you are working with a full-server installation, the graphical interface for Windows Defender is installed as a feature and cannot be uninstalled. On core installations, Windows Defender operates without the graphical interface and protects that server in the same way a full-server installation is protected. In either configuration, Windows Defender runs as a background process with malware definitions and other updates downloaded and installed via Windows Update.





You can manage Windows Defender using Settings. In Settings, click Update & Security and then click Windows Defender. You can then manage basic protection settings and exclusions.



For more control over protection settings, you'll need to use Windows PowerShell. With PowerShell, you have access to all Windows Defender settings and can manage them easily. You can list all the related cmdlets by entering the following command:

```
Get-Command -Module Defender
```

Set-MpPreference and Get-MpPreference are two cmdlets you'll use the most. You use the Set-MpPreference cmdlet to control the way Windows Defender works and the types of scans it performs. You use Get-MpPreference to list the current configuration.

Windows Defender is but one example of a program, feature or option that is better managed at the command prompt with Windows PowerShell than with the graphical interface. But it's not an all or nothing proposition when it comes to working with a fullserver or core server installation. You can manage either installation remotely using Server Manager, Microsoft Management Console (MMC) or any of the Remote Server Administration Tools that use MMC. It's just the local administration that changes.

Indeed, when you are deciding between a full server or core server installation, the required roles may be the determining factor. With a core installation, a server can only use a specific set of roles. These roles include:

- Active Directory Certificate Services (AD CS)
- Active Directory Domain Services (AD DS)
- Active Directory Lightweight Directory Services (AD LDS)
- Active Directory Rights Management Server (AD RMS)
- DHCP Server
- DNS Server
- File Services including File Server Resource Manager (FSRM)
- Hyper-V
- Print and Document Services
- Routing and Remote Access Server with the Remote Desktop Services Connection Broker, Licensing and Virtualization role services. Streaming Media Services
- Web Server including some ASP.NET components)
- Windows Server Update Server (WSUS)

Thus, if you need to install other roles on a server, you'll have to use a full server installation rather than a server core installation.

## Using Nano Server and Containers

Beyond the basics, Windows Server 2016 has a number of other deployment and configuration options. Nano server is one such option you should familiarize yourself with as you work with the operating system. Containers are another.





## Nano Server Essentials

Nano server is a minimum footprint operating system installation that's optimized to run containers. Although similar to a server core installation, a nano server is designed to be managed remotely and has only a minimal recovery console available when working locally. As nano server is meant to be used in very specific scenarios, the configuration supports a very limited set of roles, including:

- [Hyper-V](#)
- [Failover Clustering](#)
- [File Server](#)
- [DNS Server](#)
- [Web Server](#)

This makes nano server ideal for use as:

- [A storage host for file services](#)
- [A web server hosting Internet Information Services \(IIS\)](#)
- [An application server](#)
- [A DNS server](#)
- [A compute host for Hyper-V virtual machines](#)

When you are working with nano server, the only local machine management option is the Recovery Console, which is text-based and navigated using the keyboard. Here, you use the arrow keys or Tab to navigate between options and Enter to select. To go back a screen or a page, you press Esc. Although the console is meant to help you diagnose and resolve networking issues that prevent you from connecting to the server, you can use those same features at any time to view and configure network adapters, TCP/IP settings and server naming. With a keyboard and monitor connected to the server, you can log in to the recovery console using an administrator account.

## Deploying Nano Server

You can deploy nano server on a physical computer or in a virtual machine. Typically, you do so using a Windows Image (WIM) for nano server that is converted to a Virtual Hard Disk (VHD). You can also create your own custom images.


The Windows Server 2016 installation media includes a base image for Nano Server, which is stored in the \NanoServer folder. Within this folder is a subfolder called Packages that contains the package files used to add roles and features to Nano Server images.



# Container Essentials

Containers are another type of virtualization. While previous virtualization efforts focused on virtual machines that simulate the underlying physical machine, run their own operating systems and manage their own virtual memory, processors and resources, containers are all about virtualizing the operating system layer so that the need for a pseudo machine is removed entirely.

Instead of building a virtual machine with its own operating system, memory, processors and resources, containers make use of the existing hardware by building a virtualization layer on top of it. Every container behaves as if its running its own copy of the operating system, with its own file system, registry, namespaces and network addresses.



With the need for an all-inclusive machine architecture removed, containers are lightweight, flexible and dynamic. With containers, you don't have to start an operating system, allocate resources and then start processes and applications, you only have to start processes and applications on top of an existing operating system and allocate a portion of available resources. This means containers start in seconds, instead of minutes as with virtual machines.

Containers provide a fully self-contained abstraction mechanism with all the runtime dependencies needed to support a particular working environment, whether for a single application or an entire service stack. Each instance includes the application or applications in the necessary architecture and a minimum set of operating system components needed to run it.



Nano servers aren't the only server installations that can use containers. Any installation of Windows Server 2016 can use containers. In a way, containers are like Lego bricks, as you can stack containers on top of each other and build complex structures using

basic building blocks.

Every container implementation starts with a Windows base container. Typically, this base layer runs Server Core or Nano Server. You can then stack other containers on top of the base container. For example, you could have one container layer for a particular application and another layer on top of it to provide authentication. The entire container stack is then managed with a manifest that describes what ships with the application and how to build the application.


Because the container stack accounts for all the dependencies and components needed to run an application or service, the container stack has seamless portability and can be moved consistently and reliably between target environments, such from development to test to stage to production, or even back and forth between on-premises servers and hosted Azure servers.



# Deploying Containers

Before you can use containers on a server you must install the Containers feature. Windows Server 2016 supports two container variants:

- [Windows Server Containers](#)
- [Hyper-V Containers](#)



Windows Server containers maximize density and performance by sharing the operating system kernel with the host server and other Windows Server containers. The shared data and programming interfaces enhance efficiencies while processes and resources remain as isolated from each other as possible.


In contrast, Hyper-V containers provide high isolation and control by maintaining their own copy of the Windows kernel and assigning memory directly. As there are no shared programming interfaces or data, Hyper-V containers provide the same level of isolation found in virtual machines. If you want to use Hyper-V containers on a server, you must install the Hyper-V role in addition to the Containers feature.

Windows Server containers are best used with trusted applications while Hyper-V containers are best used when applications must remain isolated from each other. With Windows Containers, a namespace violation or resource governance flaw could cause problems on the host or other hosted containers. With Hyper-V containers, applications in one container are fully isolated from applications in other containers. As with virtual machines, the host only exposes a constrained interface to a Hyper-V container for communication and sharing of host resources.


The choice as to whether to use Windows Server containers or Hyper-V containers isn't one that must be made in development. Rather, it is a deployment decision that can be made and changed as necessary at any time. If a containerized application is deployed on-premises, you may want to use Windows Server containers in production. If the application is later moved to multi-tenant hosted servers, Hyper-V containers could then be used.

# Getting Started with Windows Server 2016


As with Windows 10, the Start menu is back in Windows Server 2016 and the operating system supports both universal apps and traditional software programs. Universal apps can be purchased in the Microsoft store or developed in house and deployed separately.




Apps that have tiles on Start can be started by clicking the tile. When you right-click a tile, an options panel usually is displayed, allowing you to unpin the tile, resize the tile or perform additional tasks. If an app can be opened with administrator or alternate user privileges, the related options are listed when you right-click the app tile and then select More.



Any text entered while on Start is entered into the Search box in the Search panel. By default, the Search box is focused on Apps, Settings, Documents, and the Web, which is meant to allow you to quickly find installed apps, settings, and files. However, unlike a Windows desktop, a standard deployment of Windows server doesn't use indexing services or build the indexes required to facilitate settings-related searches (as this could impact server performance). This means searches on servers primarily returns results only for apps and the Web.




If you want to search for system settings, your best option is to do so using Settings or Control Panel. Both apps have a search box, which when used matches your search text to any heading or other text used within the app.




From Start, you access Control Panel by clicking the Control Panel tile. From either Start or the desktop, you can display Settings by pressing Windows key+I.

Additionally, because File Explorer is pinned to the desktop taskbar by default, you can access Control Panel by following these steps:

1. Open File Explorer by clicking the taskbar icon.
2. Click the leftmost arrow button in the address list.
3. Click Control Panel.




Start and Desktop have a convenient menu that you can display by right-clicking the lower-left corner of the Start screen or the desktop. Alternatively, you can press Windows key + X.



By default, the command prompt and the administrator command prompt are options on the shortcut menu. The alternative is for the PowerShell prompt and the administrator PowerShell prompt to be displayed on this menu. To configure which options are available, right-click the taskbar and then select Properties. In the Taskbar And Navigation Properties dialog box, on the Navigation tab, select or clear the Replace Command Prompt With Windows PowerShell... checkbox as appropriate.






Shut Down and Restart are options of Power settings now. This means to shut down or restart a server, you follow these steps:

1. Click Start () , and then click Power ().
2. Click Shut Down or Restart as appropriate.

Alternatively, if configured as a power option, you can press the server's physical power button to initiate an orderly shutdown by logging off and then shutting down. If you are using a desktop-class system and the computer has a sleep button, the sleep button is disabled by default, as are closing the lid options for portable computers. Additionally, most other power options are disabled on servers except for display turn off. By default, servers are configured to turn off the display after 10 minutes of inactivity.





Windows 10 and Windows Server 2016 support the Advanced Configuration and Power Interface (ACPI) 5.0 specification. Windows uses ACPI to control system and device power state transitions, putting devices in and out of full-power (working), low-power, and off states to reduce power consumption.

The power settings for a computer come from the active power plan. You can access power plans in Control Panel by clicking System And Security and then clicking Power Options. Windows Server 2016 includes the Power Configuration (Powercfg.exe) utility for managing power options from the command prompt. At a command prompt, you can view the configured power plans by entering **powercfg /l**. The active power plan is marked with an asterisk.

The default, active power plan in Windows Server 2016 is called Balanced. The Balanced plan is configured to do the following:

- Never turn off hard disks (as opposed to turning off hard disks after a specified amount of idle time).
- Disable timed events to wake the computer (as opposed to enabling wake on timed events).



- Enable USB selective suspend (as opposed to disabling selective suspend). Use
- moderate power savings for idle PCI Express links (as opposed to maximum power savings being on or off).
- Use active system cooling by increasing the fan speed before slowing processors (as opposed to using passive system cooling to slow the processors before increasing fan speed).
- Use minimum processor and maximum processor states if supported (as opposed to using a fixed state).

NOTE Power consumption is an important issue, especially as organizations try to become more earth friendly. Saving power also can save your organization money and, in some cases, allow you to install more servers in your data centers. If you install Windows Server 2016 on a laptop—for testing or for your personal computer, for example—your power settings will be slightly different, and you'll also have settings for when the laptop is running on battery.





## Chapter 2. Working with Windows Servers

When you install Windows Server 2016 on a new system, you can configure the server to be a member server, a domain controller, or a stand-alone server. Member servers are part of a domain but don't store directory information. Domain controllers are distinguished from member servers because they store directory information and provide authentication and directory services for the domain. Stand-alone servers aren't part of a domain and authenticate logon requests independently.

Although you can manage servers by logging on locally, you'll more typically perform management tasks from a management computer using Windows PowerShell, Server Manager and the Remote Server Administration Tools (RSAT). At times, you might also need to use Remote Desktop Connection (mstsc.exe) to establish remote sessions with servers. Remote Desktop Connection allows you to manage remote systems in much the same way as you would if you were logged on locally.

Whether remote management and remote desktop are enabled is configured as part of your initial server setup. You also must ensure Windows Firewall or any other firewall that might be between your management computer and the remote server has exceptions that allow you to work remotely.

## Preparing Networking


In Windows 10 and Windows Server 2016, networks are identified as one of the following network types:

- **Domain** A network in which computers are connected to the corporate domain to which they are joined.
- **Work** A private network in which computers are configured as members of a workgroup and are not connected directly to the public Internet.
- **Workplace** A private network in which computers are configured as members of a workplace to which devices can connect over the public Internet. (Windows 10 only)
- **Home** A private network in which computers are configured as members of a homegroup and are not connected directly to the public Internet.
- **Public** A public network in which computers are connected to a network in a public place, such as a coffee shop or an airport, rather than an internal network.

These network types are organized into three categories: private, domain, and public. Each network category has an associated network profile. Because a computer saves sharing and firewall settings separately for each network category, you can use different block and allow settings for each network category. When you connect to a network, a



dialog box is displayed in which you can specify the network category. If you select Private, and the computer determines that it is connected to the corporate domain to which it is joined, the network category is set as Domain Network.




## Setting Network Category and Sharing

Based on the network category, Windows Server configures settings that turn network discovery on or off. When network discovery settings are enabled and a server is connected to a network, the server can see other network computers and devices and is visible on the network. Otherwise, the server cannot discover other computers and devices on the network and other computers on the network cannot discover the server. Other settings, including those for Windows Firewall, depend on the network category as well.

You can confirm the network category for active adapters in Network And Sharing Center. In Control Panel, click Network And Internet and then click Network And Sharing Center. If the category is set incorrectly, one way to fix the problem is to click Troubleshoot Problems and then follow the prompts to allow the network troubleshooter to detect the problem and help you resolve it. Another way to resolve the problem is to use Windows PowerShell to make the change. The default alias for the first network connection on a device typically is “Ethernet”. If so, you can use an administrator PowerShell prompt and the following code to change the “Ethernet” connection from public to private:

```
Get-NetConnectionProfile -InterfaceAlias "Ethernet" | Set-NetConnectionProfile -NetworkCategory Private
```

Using Advanced Sharing Settings in Network And Sharing Center, you can enable discovery and file sharing. In Control Panel, click Network And Internet, click Network And Sharing Center and then click Change Advanced Sharing Settings.



Discovery and file sharing are blocked by default on a public network, which enhances security by preventing computers on the public network from discovering other computers and devices on that network. When discovery and file sharing are disabled, files and printers you have shared from a computer cannot be accessed from the network. Additionally, some programs might not be able to access the network.

# Working with Networking Protocols

Windows Server uses TCP/IP as the default wide area network (WAN) protocol. Normally, networking is installed during installation of the operating system. You can also install TCP/IP networking through local area connection properties.

The TCP and IP protocols make it possible for computers to communicate across various networks and the Internet by using network adapters. Windows has a dual IPlayer architecture in which both Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6) are implemented and share common transport and network layers. IPv4 has 32-bit addresses and is the primary version of IP used on most networks. IPv6, alternatively, has 128-bit addresses and is the next-generation version of IP.

**NOTE** DirectAccess clients send only IPv6 traffic across the DirectAccess connection to the DirectAccess server. Thanks to the NAT64/DNS64 support, DirectAccess clients can initiate communications with IPv4-only hosts on the corporate intranet. NAT64/DNS64 work together to translate incoming connection traffic from an IPv6 node to IPv4 traffic. The NAT64 translates the incoming IPv6 traffic to IPv4 traffic and performs the reverse translation for response traffic. The DNS64 resolves the name of an IPv4-only host to a translated IPv6 address.

**REAL WORLD** The TCP Chimney Offload enables the networking subsystem to offload the processing of a TCP/IP connection from the computer's processors to its network adapter as long as the network adapter supports TCP/IP offload processing. Both TCP/IPv4 connections and TCP/IPv6 connections can be offloaded. For Windows 7 and later releases of Windows, TCP connections are offloaded by default on 10 gigabits per second (Gbps) network adapters, but they are not offloaded by default on 1-Gbps network adapters. To offload TCP connections on a 1-Gbps or 10-Gbps network adapter, you must enable TCP offloading by entering the following command at an administrator command prompt: **netsh int tcp set global chimney=enabled** . You can check the status of TCP offloading by entering **netsh int tcp show global** . Although TCP offloading works with Windows Firewall, TCP offloading won't be used with IPsec, Windows virtualization (Hyper-V), network load balancing, or the Network Address Translation (NAT) service. To determine whether TCP offloading is working, enter **netstat-t** and check the offload state. The offload state is listed as offloaded or inhost .

Windows also uses receive-side scaling and network direct memory access (NetDMA). You can enable or disable receive-side scaling by entering **netsh int tcp set global rss=enabled** or **netsh int tcp set global rss=disabled** , respectively. To

check the status of RSS, enter **netsh int tcp show global** . You can enable or disable NetDMA by setting a DWord value under the EnableTCPA registry entry to **1** or **0** , respectively. This registry entry is found under HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters.

IPv4's 32-bit addresses are commonly expressed as four separate decimal values, such as 127.0.0.1 or 192.168.10.52. The four decimal values are referred to as octets because each represents 8 bits of the 32-bit number. With standard unicast IPv4 addresses, a variable part of the IP address represents the network ID and a variable part of the IP address represents the host ID. A host's IPv4 address and its media access control (MAC) address used by the host's network adapter have no correlation.

IPv6's 128-bit addresses are divided into eight 16-bit blocks delimited by colons. Each

16-bit block is expressed in hexadecimal form, such as FEC0:0:0:02BC:FF:BECE:FE4F:961D. With standard unicast IPv6 addresses, the first 64 bits represent the network ID, and the last 64 bits represent the network interface. Because many IPv6 address blocks are set to 0, a contiguous set of 0 blocks can be expressed as "::", a notation referred to as double-colon notation . Using double-colon notation, the two 0 blocks in the previous address can be compressed as FEC0::02BC:FF:BECE:FE4F:961D. Three or more 0 blocks would be compressed in the same way. For example, FFE8:0:0:0:0:0:0:1 becomes FFE8::1.

When networking hardware is detected during installation of the operating system, both IPv4 and IPv6 are enabled by default; you don't need to install a separate component to enable support for IPv6. Group Policy Preferences also have been updated to support IPv6 addresses for TCP/IP printers and VPN connections. Item-level Targeting options also allow you to configure IPv6 address ranges.

You can manage TCP/IP settings for a server using Network Connections. In Control Panel, click Network And Internet, click Network And Sharing Center and then click Change Adapter Settings. You'll then see a list of available network adapters.





To manage a network adapter, right-click the adapter name and then select Properties. This opens the properties dialog box for the adapter. With the properties dialog box open, you can configure IPv4 and IPv6 settings using the following techniques:

- Select Internet Protocol Version 4 (TCP/IPv4) and then click Properties. Use the Internet Protocol Version 4 Properties dialog box to configure the IPv4 settings.
- Select Internet Protocol Version 6 (TCP/IPv6) and then click Properties. Use the Internet Protocol Version 6 Properties dialog box to configure the IPv6 settings.

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

☐ Obtain an IP address automatically

☒ Use the following IP address:

IP address: 10 . 0 . 0 . 172

Subnet mask: 255 . 255 . 255 . 0

Default gateway: 10 . 0 . 0 . 1

☐ Obtain DNS server address automatically

☒ Use the following DNS server addresses:

Preferred DNS server: 10 . 0 . 0 . 171

Alternate DNS server: 75 . 75 . 75 . 75

☐ Validate settings upon exit

Advanced...



OK Cancel

Most modern server hardware has multiple network adapters installed to support NIC Teaming, which allows multiple network adapters to be grouped together for sharing bandwidth, traffic failover or both. If a network adapter is teamed in software (vs. hardware-based teaming), you'll see that the Microsoft Network Adapter Multiplexor Protocol is enabled while the IPv4 and IPv6 protocols aren't selected. This is because you configure TCP/IP settings for teamed adapters using the primary NIC in the team, which is a multiplexed adapter. Thus, if adapters are teamed edit the settings for the primary NIC rather than the team members.

You configure NIC Teaming in Server Manager. When you are working with the All Servers node, right-click a server entry and then select Configure NIC Teaming. Next,



use the options provided in the NIC Teaming console to manage the team configuration.






# Navigating Management Options

When you want to manage server configurations, Server Manager and the Server Configuration utility will likely be your tools of choice.



# Getting Started with Server Manager

Server Manager is a graphical administration tool that is available whenever you are working with Windows Server 2016. Not only can you use Server Manager to add or remove roles, role services, and features, but you can also use Server Manager to configure basic settings and options.




Server Manager's default view is the dashboard. The dashboard provides quick links for adding roles and features to local and remote servers, adding servers to manage, and creating server groups. You'll find similar options are on the Manage menu:

- **Add Roles And Features** Starts the Add Roles And Features Wizard, which you can use to install roles, role services, and features on the server.
- **Add Servers** Opens the Add Servers dialog box, which you can use to add servers you want to manage. Added servers are listed when you select the All Servers node. Right-click a server in the Servers pane of the All Servers node to display a list of management options, including Restart Server, Manage As, and Remove Server.
- **Create Server Group** Opens the Create Server Group dialog box, which you can use to add servers to server groups for easier management. Server Manager creates role-based groups automatically. For example, domain controllers are listed under AD DS, and you can quickly find information about any domain controllers by selecting the related node.


Windows Server 2016 automatically starts Server Manager whenever you log on by default. If you don't want the console to start each time you log on, you can change the default settings. Click Manage, and then click Server Manager Properties. In the Server

Manager Properties dialog box, select Do Not Start Server Manager Automatically At Logon and then click OK.



**NOTE** Group Policy also can be used to control automatic start of Server Manager. Enable or disable the Do Not Display Server Manager Automatically At Logon policy setting within System\Server Manager in the Administrative Templates for Computer Configuration.

In Server Manager's left pane (also referred to as the console tree ), you'll find options for accessing the dashboard, the local server, all servers added for management, and server groups. When you select Local Server in the console tree, you can manage the basic configuration of the server.



Information about the local server is organized into several main headings, each with an associated management panel:

- **Best Practices Analyzer** Allows you to run the Best Practices Analyzer on the server and review the results. To start a scan, click Tasks and then click Start BPA Scan.

- **Events** Provides summary information about warning and error events from the server's event logs. Click an event to display more information about the event.
- **Performance** Allows you to configure and view the status of performance alerts for CPU and memory usage. To configure performance alerts, click Tasks and then click Configure Performance Alerts.
- **Properties** Shows the computer name, domain, network IP configuration, time zone, and more. Each property can be clicked to quickly display a related management interface.
- **Roles And Features** Lists the roles and features installed on the server, in the approximate order of installation. To remove a role or feature, right-click it and then select Remove Role Or Feature.
- **Services** Lists the services running on the server by name, status, and start type. Right-click a service to manage its run status.

The Properties panel is where you perform much of your initial server configuration. Properties available for quick management include the following:

- **Computer Name/Domain** Shows the computer name and domain. Click either of the related links to display the System Properties dialog box with the Computer Name tab selected. You can then change a computer's name and domain information by clicking Change, providing the computer name and domain information, and then clicking OK. By default, servers are assigned a randomly generated name and are configured as part of a workgroup called WORKGROUP. In the Small Icons or Large Icons view of Control Panel, you also can display the System Properties dialog box with the Computer Name tab selected by clicking System and then clicking Change Settings under Computer Name, Domain, And Workgroup Settings.
- **Customer Experience Improvement Program** Shows whether the server is participating in the Customer Experience Improvement Program (CEIP). Click the related link to change the participation settings. Participation in CEIP allows Microsoft to collect information about the way you use the server. Microsoft collects this data to help improve future releases of Windows. No data collected as part of CEIP personally identifies you or your company. If you elect to participate, you can also provide information about the number of servers and desktop computers in your organization and in your organization's general industry. If you opt out of CEIP by turning this feature off, you miss the opportunity to help improve Windows.
- **Ethernet** Shows the TCP/IP configuration of wired Ethernet connections. Click the related link to display the Network Connections console. You can then configure network connections by double-tapping or double-clicking the connection you want



to work with and then clicking Properties to open the Properties dialog box. By default, servers are configured to use dynamic addressing for both IPv4 and IPv6. You can also display the Network Connections console by clicking Change Adapter Settings under Tasks in Network And Sharing Center.

- **IE Enhanced Security Configuration** Shows the status of Microsoft Internet Explorer Enhanced Security Configuration (IE ESC). Click the related link to enable or disable IE ESC. If you click the link for this option, you can turn this feature on or off for administrators, users, or both. IE ESC is a security feature that reduces the exposure of a server to potential attacks by raising the default security levels in Internet Explorer security zones and changing default Internet Explorer settings. By default, IE ESC is enabled for both administrators and users.

**REAL WORLD** In most cases, you should enable IE ESC on a server for both users and administrators. However, enabling IE ESC reduces the functionality of Internet Explorer. When IE ESC is enabled, security zones are configured as follows: the Internet zone is set to Medium-High, the Trusted Sites zone is set to Medium, the Local Intranet zone is set to Medium-Low, and the Restricted zone is set to High. When IE ESC is enabled, the following Internet settings are changed: the Enhanced Security Configuration dialog box is on, third-party browser extensions are off, sounds in web pages are off, animations in web pages are off, signature checking for downloaded programs is on, server certificate revocation is on, encrypted pages are not saved, temporary Internet files are deleted when the browser is closed, warnings for secure and nonsecure mode changes are on, and memory protection is on.

- **NIC Teaming** Shows the status and configuration of network interface card (NIC) teaming. Click the related link to add or remove teamed interfaces and to manage related options.
- **Product ID** Shows the product identifier for Windows Server. Click the related link to enter a product key and activate the operating system over the Internet.
- **Remote Desktop** Click the related link to display the System Properties dialog box with the Remote tab selected. You can then configure Remote Desktop by selecting the configuration option you want to use and clicking OK. By default, no remote connections to a server are allowed. In the Small Icons or Large Icons view of Control Panel, you can display the System Properties dialog box with the Remote tab selected by double-tapping or double-clicking System and then clicking Remote Settings in the left pane.






- **Remote Management** Shows whether remote management of this server from other servers is enabled. Click the related link to enable or disable remote management.
- **Time Zone** Shows the current time zone for the server. Click the related link to display the Date And Time dialog box. You can then configure the server's time zone by clicking Change Time Zone, selecting the appropriate time zone, and then clicking OK twice. You can also display the Date And Time dialog box by right-clicking the clock on the taskbar and then selecting Adjust Date/Time. Although all servers are configured to synchronize time automatically with an Internet time server, the time synchronization process does not change a computer's time zone.
- **Windows Error Reporting** Shows the status of Windows Error Reporting (WER). Click the related link to change the participation settings for WER. In most cases, you'll want to enable WER for at least the first 60 days following installation of the operating system. With WER enabled, your server sends descriptions of problems to Microsoft, and Windows notifies you of possible solutions to those problems. You can view problem reports and possible solutions by using Action Center. To open Action Center, click the Action Center icon in the notification area of the taskbar, and then click Open Action Center.
- **Windows Firewall** Shows the status of Windows Firewall. If Windows Firewall is active, this property displays the name of the firewall profile that currently applies and the firewall status. Click the related link to display the Windows Firewall utility. By default, Windows Firewall is enabled. In the Small Icons or Large Icons view of Control Panel, you can display Windows Firewall by clicking the Windows Firewall option.
- **Windows Update** Shows the current configuration of Windows Update. Click the related link to display the Windows Update utility in Control Panel, which you can then use to enable automatic updating (if Windows Update is disabled) or to check for updates (if Windows Update is enabled). In the Small Icons or Large Icons view of Control Panel, you can display Windows Update by selecting the Windows Update option.

**NOTE** I've provided this summary of options as an introduction and quick reference. I'll discuss the related configuration tasks and technologies in more detail throughout this and other chapters in the book.

## Getting Started with Sconfig

When you log on to a Server Core installation, you can use the Server Configuration utility (Sconfig) to manage the server. Sconfig is a text-based utility with a menu system that allows you to configure domain and workgroup membership, computer

naming, network settings, and more. You also can use Sconfig to log off, restart, and shut down a server.



To start Sconfig, simply enter **sconfig** at the command prompt. You can then choose menu options and follow the prompts to configure the server. Options available for quick management include the following:

- **Add Local Administrator** Adds a domain or local user to the local administrators group on the server. For a domain account, be sure to specify the account name in domain\username format, such as imaginedl\williams.
- **Computer Name** Shows the server name. Select this option to change the name of the server.
- **Configure Remote Management** Shows whether remote management is enabled. Provides options for enabling or disabling remote management as well as configuring server response to ping.
- **Date And Time** Displays the date and time control panel, which allows you to configure date, time and time zone. Using Internet Time settings, you also can specify an Internet Time server and related settings. If your organization has a time server, use Group Policy to configure time services instead.
- **Domain/Workgroup** Shows the domain or work group membership. Select this option to join the server to a domain or workgroup.
- **Download And Install Updates** Performs an immediate download and install of either all available updates or only recommended updates.
- **Network Settings** Provides options for managing TCP/IP settings using either DHCP or static IP addressing. You manage available adapters according to their NIC index number. When you configure IP addressing, you can specify the IP address,





subnet mask and default gateway. Additional options are provided for configuring DNS server settings.

- **Remote Desktop** Shows whether remote desktop connections are allowed.

Provides options for enabling or disabling remote desktop connections. •

- **Windows Activation** Provides options for displaying licensing information, activating Windows and entering a product key.

- **Windows Update Settings** Shows the status of Windows Update. Provides options changing Windows Update to automatic or manual.



# Getting Started with Windows PowerShell

For additional flexibility when working locally or remotely, you might want to use Windows PowerShell. Windows PowerShell is a full-featured command shell that can use built-in commands called cmdlets, built-in programming features, and standard command-line utilities. A command console and a graphical environment are available.

Although the Windows PowerShell console and the graphical scripting environment are installed by default, several other PowerShell features are not installed by default. They include the Windows PowerShell 2.0 engine, which is provided for backward compatibility, and Windows PowerShell Web Access, which lets a server act as a web gateway for managing the server remotely by using PowerShell and a web client.




You can open a Windows PowerShell console by clicking the related tile on Start or in the Apps list. You also can start Windows PowerShell from a Windows command shell (Cmd.exe) by typing **powershell** and pressing Enter.

After starting PowerShell, you can enter the name of a cmdlet at the prompt, and the cmdlet will run in much the same way as a command-line command. You also can execute cmdlets in scripts. Cmdlets are named by using verb-noun pairs. The verb tells you what the cmdlet does in general. The noun tells you what specifically the cmdlet works with. For example, the Get-Variable cmdlet gets all Windows PowerShell environment variables and returns their values, or it gets a specifically named environment variable and returns its value. The common verbs associated with cmdlets are as follows:

- **Get-** Queries a specific object or a subset of a type of object, such as a specified performance counter or all performance counters
- **Set-** Modifies specific settings of an object
- **Enable-** Enables an option or a feature
- **Disable-** Disables an option or a feature

- New-** Creates a new instance of an item, such as a new event or service
- **Remove-** Removes an instance of an item, such as an event or event log



At the PowerShell prompt, you can get a complete list of cmdlets by entering **get-help \*-\*** . To get Help documentation on a specific cmdlet, enter **get-help** followed by the cmdlet name, such as **get-help get-variable** . Windows PowerShell V3 and later use online and updatable Help files. Because of this, you might see only basic syntax for cmdlets and functions. To get full Help details, you'll have to either use online Help or download the Help files to your computer. For online Help, add the **-online** option to your **get-help** command, such as:

```
get-help get-variable -online
```

Use the **Update-Help** cmdlet to download and install the current Help files from the Internet. Without parameters, **Update-Help** updates the Help files for all modules installed on the computer. However, **Update-Help** does the following:

- Downloads files only once a day
- Installs files only when they are newer than the ones on the computer
- Limits the total size of uncompressed Help files to 1 GB

You can override these restrictions by using the **-Force** parameter. You can save Help files to the local computer by using **Save-Help**.

All cmdlets also have configurable aliases that act as shortcuts for executing a cmdlet. To list all aliases available, enter **get-alias** at the Windows PowerShell prompt. You can create an alias that invokes any command by using the following syntax. `new-item -path alias: AliasName -value: FullCommandPath`

Here **AliasName** is the name of the alias to create, and **FullCommandPath** is the full path to the command to run, such as the following.

```
new-item -path alias:sm -value:c:\windows\system32\compmgmtlauncher.exe
```



This example creates the alias `sm` for starting Server Manager. To use this alias, you simply type `sm` and then press Enter when you are working with Windows PowerShell. It's important to note that Windows PowerShell 3 and later versions automatically import required modules the first time you use a related command. With PowerShell 2, you needed to explicitly import a module before you could run any of its commands.

Generally, anything you can enter at a command prompt also can be entered at the PowerShell prompt. This is possible because PowerShell looks for external commands and utilities as part of its usual processing. As long as the external command or utility is found in a directory specified by the `PATH` environment variable, the command or utility is run as appropriate. However, keep in mind that the PowerShell execution order could affect whether a command runs as expected. For PowerShell, the execution order is:

1. Alternate built-in or profile-defined aliases.
2. Built-in or profile-defined functions.
3. Cmdlets or language keywords
4. Scripts with the `.ps1` extension
5. External commands, utilities, and files.

Following this, if any element in 1 through 4 of the execution order has the same name as a command utility or app, that element will run instead of the expected utility or app.





# Managing Your Servers Remotely

You can use Server Manager, Microsoft Management Consoles (MMCs), and Windows PowerShell to perform many management tasks on remote computers. As long as the servers are in the same domain or you are working in a workgroup and have added the remote computers as trusted hosts, you can connect to the servers and remotely manage them regardless of whether they are running full-server or Server Core installations. On the computer you want to use for managing remote computers, you should be running a current Windows desktop or server operating system with the Remote Server Administration Tools installed.

With Windows Server 2016, the Remote Server Administration Tools are installed as a feature by using the Add Roles And Features Wizard. If the binaries for the tools have been removed, you need to install the tools by specifying a source, as discussed in the “Getting Binaries for Roles and Features” section later in the chapter.

You can get the Remote Server Administration Tools for Windows 10 as a download from the Microsoft Download Center (<http://download.microsoft.com>). Different versions are available for x64 and x86 systems.



# Remote Management Essentials

By default, remote management is enabled for servers running Windows Server 2016 for two types of applications and commands:

- Applications and commands that use Windows Remote Management (WinRM) and Windows PowerShell remote access for management
- Applications and commands that use Windows Management Instrumentation (WMI) and Distributed Component Object Model (DCOM) remote access for management

These types of applications and commands are permitted for remote management because of exceptions configured in Windows Firewall, which are enabled by default. In Windows Firewall, exceptions for allowed apps that support remote management include the following:


- Windows Management Instrumentation
- Windows Remote Management
- Windows Remote Management (Compatibility)

In Windows Firewall With Advanced Security, there are inbound rules that correspond to the standard firewall allowed apps:

- For WMI, the inbound rules are Windows Management Instrumentation (WMI-In), Windows Management Instrumentation (DCOM-In), and Windows Management Instrumentation (ASync-In).
- For WinRM, the matching inbound rule is Windows Remote Management (HTTP-In).
- For WinRM compatibility, the matching inbound rule is Windows Remote Management - Compatibility Mode (HTTP-In).

You manage these exceptions or rules in either the standard Windows Firewall or in Windows Firewall With Advanced Security, not both. If you want to allow remote management by using Server Manager, MMCs, and Windows PowerShell, you typically want to permit WMI, WinRM, and WinRM compatibility exceptions in Windows Firewall.

When you are working with Server Manager, you can select Local Server in the console tree to view the status of the remote management property. If you don't want to allow remote management of the local server, click the related link. In the Configure Remote Management dialog box, clear Enable Remote Management Of This Server From Other Computers, and then click OK.



When you clear Enable Remote Management Of This Server From Other Computers and then click OK, Server Manager performs several background tasks that disable Windows Remote Management and Windows PowerShell remote access for management on the local server. One of these tasks is to turn off the related exception that allows apps to communicate through Windows Firewall by using Windows Remote Management. The exceptions for Windows Management Instrumentation and Windows Remote Management (Compatibility) aren't affected.

You also can disable or enable remote management at an administrator command prompt. Enable remote management by entering **configure-SMRemoting.exe –enable** . Disable remote management by entering **configure-SMRemoting.exe –disable** .

**NOTE** If you want to make it possible to remotely manage a computer running Windows 10 by using the WS-Management protocol, enter **winrm quickconfig** at an elevated command prompt. Then each time you are prompted to make configuration changes, enter **Y** . This will start the Windows Remote Management service, configure WinRM to accept WS-Management requests on any IP address, create a Windows Firewall exception for Windows Remote Management, and configure LocalAccountTokenFilterPolicy to grant appropriate administrative rights for remote management.

Many other types of remote management tasks depend on other exceptions for Windows Firewall. Keep the following in mind:

- Remote Desktop is enabled or disabled separately from remote management. To allow someone to connect to the local server by using Remote Desktop, you must allow related connections to the computer and configure access. When working with a local server in Server Manager, click the Enabled or Disabled link for Remote Desktop and then either allow or don't allow remote connections.
- Remote Service Management must be configured as an allowed app in Windows Firewall to remotely manage a computer's services. In the advanced firewall, several related rules allow management via named pipes (NP) and remote procedure calls (RPC).
- Remote Event Log Management must be configured as an allowed app in Windows Firewall to remotely manage a computer's event logs. In the advanced firewall, there are several related rules that allow management via NP and RPC.





- Remote Volume Management must be configured as an allowed app in Windows Firewall to remotely manage a computer's volumes. In the advanced firewall, there are several related rules that allow management of the Virtual Disk Service and Virtual Disk Service Loader.
- Remote Scheduled Task Management must be configured as an allowed app in Windows Firewall to remotely manage a computer's scheduled tasks. In the advanced firewall, there are several related rules that allow management of scheduled tasks via RPC.

Only Remote Service Management is enabled by default.







# Configuring Remote Management and Trusted Hosts

Remoting features are supported by the WS-Management protocol and the Windows Remote Management (WinRM) service that implements WS-Management in Windows. If you want to manage servers remotely, you need to be sure that WinRM is installed and that the server has a WinRM listener enabled. You can verify the availability of WinRM by following these steps:

- 1. Open an administrator PowerShell prompt. One way to do this is to click Start, right-click the Windows PowerShell tile, point to More and then select Run As Administrator.
- 2. The WinRM service is configured for manual start by default. You must change the startup type to Automatic and start the service on each computer you want to work with. At the Windows PowerShell prompt, you can verify that the WinRM service is running by using the following command:

```
get-service winrm
```

As shown in the following example, the value of the Status property in the output should be Running :

| Status  | Name  | DisplayName               |
|---------|-------|---------------------------|
| -----   | ----  | -----                     |
| Running | WinRM | Windows Remote Management |

If the service is stopped, enter the following command to start the service, and then configure it to start automatically in the future:

```
set-service -name winrm -startuptype automatic -status running
```

- 3. To configure Windows PowerShell for remoting, enter the following command:

```
Enable-PSRemoting -force
```


In many cases, you are able to work with remote computers in other domains. However, if the remote computer is not in a trusted domain or you are working in a workgroup and want to access a computer in a domain, the remote computer might not be able to authenticate your credentials. To enable authentication, you need to add the remote computer to the list of trusted hosts for the local computer in WinRM. To do so, enter the following.

```
winrm set winrm/config/client '@{TrustedHosts="RemoteComputer"}'
```

Here RemoteComputer is the name of the remote computer, such as the following.

```
winrm set winrm/config/client '@{TrustedHosts="CorpServer56"}'
```





```
PS C:\Windows\system32> enable-psremoting -force
PS C:\Windows\system32> winrm set winrm/config/client @{TrustedHosts="corpserver32"}
Client
NetworkDelay = 5000
URLPrefix = wsman
AllowUnencrypted = false
Auth
Basic = true
Digest = true
Kerberos = true
Negotiate = true
Certificate = true
CredSSP = false
DefaultPorts
HTTP = 5985
HTTPS = 5986
TrustedHosts = CorpServer32
PS C:\Windows\system32>
```

When you are working with computers in workgroups or homegroups, you must use HTTPS as the transport or add the remote machine to the TrustedHosts configuration settings. If you cannot connect to a remote host, verify that the service on the remote host is running and is accepting requests by running the following command on the remote host. `winrm quickconfig`

This command analyzes and configures the WinRM service. If the WinRM service is set up correctly, you'll get output similar to the following.

WinRM already is set up to receive requests on this machine.


WinRM already is set up for remote management on this machine.

If the WinRM service is not set up correctly, you receive errors and need to respond affirmatively to several prompts that allow you to automatically configure remote management. When this process is complete, WinRM should be set up correctly.

## Adding Remote Servers for Management


Using Server Manager, you can connect to and manage remote servers, provided that you've added the server for management. If the computer you are logged on to is a member of an Active Directory domain, you can add another member server by completing these steps:

1. Open Server Manager. In the left pane, select All Servers to view the servers that have been added for management already. If the server you want to work with isn't listed, select Add Servers on the Manage menu to display the Add Servers dialog box.
2. On the Active Directory panel of the Add Servers dialog box, type the computer name or fully qualified domain name of a remote Windows server that is in a trusted domain. After you enter a name, click Find Now.
3. In the Name list, double-click the server to add it to the Selected list.
4. Repeat steps 2 and 3 to add others servers, and then click OK.



If the computer you are logged on to is a member of a workgroup, you can add other trusted servers by completing these steps:

1. Open Server Manager. In the left pane, select All Servers to view the servers that have been added for management already. If the server you want to work with isn't listed, select Add Servers on the Manage menu to display the Add Servers dialog box.
2. On the DNS panel of the Add Servers dialog box, enter the server name or IP address, and then click the Search button.




3. In the Name list, double-click the server to add it to the Selected list.
4. Repeat steps 2 and 3 to add others servers, and then click OK.

Just because you can add a trusted server doesn't mean the account you are using has the appropriate permissions for management. If the server is online but access to the

server is denied in Server Manager, you may need to connect to the server by using alternate credentials:

1. Right-click the server with the access problem in the All Servers node and then select Manage As.



The screenshot shows a Windows Server Manager console window with a red header bar that says 'Refresh failed'. Below the header is a table of servers. The table has columns for Server Name, IPv4 Address, Manageability, Last Update, and Windows Activation. Two servers are listed: CORPSEVER17 and CorpServer32. CorpServer32 is highlighted in blue and shows an 'Access denied' status.

| Server Name  | IPv4 Address | Manageability                             | Last Update          | Windows Activation  |
|--------------|--------------|-------------------------------------------|----------------------|---------------------|
| CORPSEVER17  | 10.0.0.172   | Online - Performance counters not started | 5/23/2016 3:07:25 PM | 00133-32500-01416-A |
| CorpServer32 | -            | Online - Access denied                    | 5/23/2016 3:07:41 PM | -                   |

2. In the Windows Security dialog box, enter your alternate credentials. By default, credentials you provide are cleared when you exit Server Manager. To save the credentials and use them each time you log on, select Remember My Credentials in the Windows Security dialog box.
3. Click OK. You need to repeat this procedure any time you change the password associated with the alternate credentials.



The screenshot shows the 'Windows Security' dialog box with the title 'Enter User Account and Password'. It contains instructions on how to enter the user name and password, and a checkbox for 'Remember my credentials'. The user name 'imaginedlands\williams' is entered in the first text box, and the password is masked with dots in the second text box. The 'Remember my credentials' checkbox is checked. The 'OK' button is highlighted with a blue border.

Windows Security

Enter User Account and Password

Enter the name and password of an account with user rights on the selected servers.

For example, user, user@example.contoso.com, or domain\user name.

Entering a blank user name and password deletes cached credentials.


imaginedlands\williams

.....

☒ Remember my credentials

OK Cancel

Connecting to a trusted host with the appropriate credentials should clear any access denied errors.




**SERVERS**  
All servers | 2 total

Filter [ ] [ ] [ ] [ ]


| Server Name | IPv4 Address | Manageability                             | Last Update          | Windows Activation  |
|-------------|--------------|-------------------------------------------|----------------------|---------------------|
| CORPSEVER17 | 10.0.0.172   | Online - Performance counters not started | 5/23/2016 3:17:25 PM | 00133-32500-01416-A |
| CORPSEVER32 | 10.0.0.171   | Online - Performance counters not started | 5/23/2016 3:17:26 PM | 00133-32500-01416-A |

Server Manager also provides an import process for adding many servers for management. To add servers using an import, follow these steps:

1. Create a text file that has one host name, fully qualified domain name, or IP address per line.



2. In Server Manager, select Add Servers on the Manage menu. In the Add Servers dialog box, select the Import panel.



3. Click the options button to the right of the File box, and then use the Open dialog box to locate and open the server list.
4. In the Add Servers dialog box, the Computers list shows the imported list of servers. Add servers you want to manage to the Selected list by selecting them and then clicking the Add button.

**NOTE** To select all the servers, click the first name in the list, hold Shift and then click the last name in the list. Finally, click the Add button to add all the servers to the Selected list.



## Grouping Servers for Management

Server Manager shows the host name of each added server in the All Servers view. The

All Servers view also lists the Manageability status. If a server is listed as Not Accessible, you typically need to log on locally to resolve the problem. For example, you may need to enable remote management.

In the All Servers view, added servers are listed in the Servers pane so that you can manage them each time you work with Server Manager. Server Manager tracks services,


events, and more for each server, and each server is added automatically to the appropriate server groups based on the roles and features installed.

Automatically created server groups make it easier to manage the various roles and features that are installed on your servers. If you select the AD DS group, as an example, you get a list of the domain controllers you added for management in addition to any critical or warning events for these servers and the status of services the role depends on.

If you want to group servers by department, geographic location, or otherwise, you can create your own server groups. When you create groups, the servers you want to work with don't have to be added to Server Manager already. You can add servers by searching Active Directory or Domain Name System (DNS), or by importing a list of servers by host name, fully qualified domain names, or IP addresses. Any server you add to a custom group is added automatically for management, too.

To create a server group, complete these steps:

1. [Open Server Manager. Select Create Server Group on the Manage menu to display the Create Server Group dialog box.](#)



2. Enter a descriptive name for the group. Use the panels and options provided to add servers to the group. Keep the following in mind:


- The Server Pool panel, selected by default, lists servers that have been added for management already. If a server you want to add to your group is listed here, add it to the group by double-tapping or double-clicking it.
- The Active Directory panel allows you to enter the computer name or fully qualified domain name of the remote server that is running Windows Server. After you enter a name, click Find Now. In the Name list, double-click a server to add it to the Selected list.
- The DNS panel allows you to add servers by computer name or IP address. After you enter the name or IP address, click the Search button. In the Name list, double-click a server to add it to the Selected list.
- The Import panel allows you to import a list of servers. Click the options button to the right of the File box, and then use the Open dialog box to locate and open the server list. In the Computer list, double-click a server to add it to the Selected list.

3. Click OK to create the server group.




## Working with Added Servers

You can view information about servers in the Servers pane of a server group or in the All Servers view. When you right-click a server name, you display an extended list of management options that allow you to add roles and features, restart the server, and more.



Any management option you choose opens the related tool with the selected server in focus. If alternate credentials are required and you saved your credentials previously, the saved credentials typically are used for authentication. Otherwise, you are prompted for the necessary credentials.


To run executables for administration tools, the related remote server administration tool must be installed on the management computer. If a tool you select isn't installed, you'll see a command prompt window with an error stating "XYZ.exe is not recognized as an internal or external command, operable program or batch file."



When you choose Windows PowerShell as an option, Server Manager opens an administrator PowerShell prompt and then establishes an interactive remote session with the server. Once the session is established, you can enter commands as if you



were using Windows PowerShell locally. End the interactive session by closing the PowerShell window.



If you are working from your management computer and don't have Server Manager open, you can establish similar remote sessions by opening an administrator PowerShell prompt and entering the following command:

```
enter-psession ComputerName -credential UserName
```

where ComputerName is the name of the remote computer and UserName is the name of a user who is a member of the Administrators group on the remote computer or in the domain of which the remote computer is a member. If necessary, the user name can be specified in domain\username format or as an email address for a Microsoft account.

You will then be prompted for the authorized user's password. Type the password and then press Enter. The following example opens an interactive remote session with Server23 using the credentials of WilliamS in the ImaginedLands domain:

```
enter-psession server23 -credential imaginedlands\williams
```

The command prompt changes to show that you are connected to the remote computer, as shown in the following example:

```
[Server23]: PS C:\Users\wrstane.k.imaginedlands\Documents>
```

Once the session is established, you can enter commands as if you were using Windows PowerShell locally. End the interactive session by using the command **ExitPSSession** or entering **exit**.

## Managing Roles and Features

Server Manager automatically creates server groups based on the roles of the servers added for management. As an example, the first time you add a domain controller, Server Manager might create AD DS, DNS, and File And Storage Services groups to help you more easily track the roles of the domain controllers.




## Working with Servers Based on Roles

When you select a role-based group in the left pane, the Servers pane shows the servers you added for management that have this role. The details for the selected server group provide the following information:


- Summary information about events. Server Manager lists recent warning and error events. If you click an event, you can get more information about the event.
- Summary information about the status of related system services. You can right-click a service to manage its run status.

You can manage a service by right-clicking the service and then clicking Stop Services, Start Services, Pause Services, Resume Services, or Restart Services as appropriate. In many cases, if a service isn't running as you think it should, you can use the Restart option to resolve the issue by stopping and then starting the service.



By default, Server Manager refreshes details every 10 minutes. You can refresh the details manually by clicking the Refresh button on the Tasks toolbar. If you want to set a

different default refresh interval, click Manage and then click Server Manager Properties. Next, set the new refresh interval in minutes and then click OK.



## Adding Roles and Features

The Manage menu has two key options for working with roles and features:

- **Add Roles And Features** Starts the Add Roles And Features Wizard, which you can use to install roles and features on a server added for management.
- **Remove Roles And Features** Starts the Remove Roles And Features Wizard, which you can use to uninstall roles and features on a server added for management.




With Windows Server 2016, you can install roles and features on running servers (whether physical machines or virtual) in addition to virtual hard disks. Servers you want to work with must be added for management in Server Manager, and they must be online. Virtual hard disks that you want to work with don't have to be online, but they must be selectable when you are browsing for them. Because of this, you might need to map a network drive to access a network share. With this in mind, you can add a server role or feature by following these steps:

1. In Server Manager, click **Add Roles And Features** on the **Manage** menu. This starts the **Add Roles And Features Wizard**. If the wizard displays the **Before You Begin** page, read the introductory text, and then click **Next**. You can avoid displaying the **Before You Begin** page the next time you start this wizard by selecting the **Skip This Page By Default** check box before clicking **Next**.
2. On the **Select Installation Type** page, **Role-Based Or Feature-Based Installation** is selected by default. Click **Next**.




3. On the **Select Destination Server** page, you can choose to install roles and features on running servers or virtual hard disks. Either select a server from the server pool or select a server from the server pool on which to mount a virtual hard disk (VHD). If you are adding roles and features to a VHD, click **Browse** and then use the **Browse For Virtual Hard Disks** dialog box to locate the VHD. When you are ready to continue, click **Next**.



4. On the Select Server Roles page, select the role or roles to install. If additional features are required to install a role, an additional dialog box will be displayed. Click Add Features to close the dialog box and add the required features to the server installation. Click Next to continue.




5. On the Select Features page, select the feature or features to install. If additional features are required to install a feature you selected, an additional dialog box will be displayed. Click Add Features to close the dialog box and add the required features to the server installation. When you are ready to continue, click Next.




**NOTE** With some roles, an extra wizard page will be displayed, and this provides additional information about using and configuring the role. You might also have the opportunity to install additional role services as part of a role.

6. On the Confirm Installations Selections page, click the Export Configuration Settings link to generate an installation report that can be displayed in Internet Explorer. If roles or features that you are installing require a server restart to complete, you can allow this by selecting the Restart option.




7. If the server on which you want to install roles or features doesn't have all the required binary source files, the server gets the files via Windows Update by default or from a location specified in Group Policy. You also can specify an alternate path for the source files. To do this, click the Specify An Alternate Source Path link, enter that alternate path in the box provided, and then click OK. For example, if you mounted a Windows image and made it available on the local server as discussed in the section "Getting Binaries for Roles and Features" later in this chapter, you could enter the alternate path as **c:\mountdir\windows\winsxs** . For network shares, enter the UNC path to the share, such as **\\CorpServer82\WinServer2016\** . For mounted Windows images, enter the WIM path prefixed with **WIM:** and including the index of the image to use, such as **WIM:\\CorpServer82\WinServer2016\install.wim:4** .
8. After you review the installation options and save them as necessary, click Install to begin the installation process. The Installation Progress page tracks the progress of the installation. If you close the wizard, click the Notifications icon in Server Manager, click Task Details and then click the Action link provided to reopen the wizard.



9. When the wizard finishes installing the server with the roles and features you selected, the Installation Progress page will be updated to reflect this. Review

the installation details to ensure that all phases of the installation were completed successfully.

Note any additional actions that might be required to complete the installation, such as restarting the server or performing additional installation tasks. Click the Notifications icon to get additional information.




If any portion of the installation failed, note the reason for the failure. Review the Server Manager entries for installation problems and take corrective actions as appropriate.

## Removing Roles and Features


You can remove roles and features from a server by following these steps:

1. In Server Manager, click Remove Roles And Features on the Manage menu. This starts the Remove Roles And Features Wizard. If the wizard displays the Before You Begin page, read the introductory text and then click Next. You can avoid displaying the Before You Begin page the next time you start this wizard by selecting the Skip This Page By Default check box before clicking Next.
2. On the Select Destination Server page, you can choose to remove roles and features from running servers or virtual hard disks (VHDs). Either select a server from the server pool or select a server from the server pool on which to mount a VHD. If you are removing roles and features from a VHD, click Browse, and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, click Next.






- On the Remove Server Roles page, clear the check box for the role you want to remove. If you try to remove a role that another role or feature depends on, a warning prompt appears stating that you cannot remove the role unless you also remove the other role. If you click the Remove Features button, the wizard also removes the dependent roles and features. Note that if you want to keep related management tools, you should clear the Remove Management Tools check box prior to clicking the Remove Features button and then click Continue. Click Next.



- On the Remove Features page, the currently installed features are selected. To remove a feature, clear the related check box. If you try to remove a feature that another feature or role depends on, a warning prompt appears stating that you cannot remove the feature unless you also remove the other feature or role. If you click the Remove Features button, the wizard also removes the dependent roles and features. Note that if you want to keep related management tools, you should clear the Remove Management Tools check box and then click Continue prior to clicking the Remove Features button. Click Next.



- On the Confirm Removal Selections page, review the related components that the wizard will remove based on your previous selections. If roles or features that you are uninstalling require a server restart to complete, you can allow this by selecting the Restart option.




- Click Remove. The Removal Progress page tracks the progress of the removal. If you close the wizard, click the Notifications icon in Server Manager, click Task Details and then click the Action link provided to reopen the wizard.



- When the wizard finishes modifying the server configuration, the Removal Progress page is displayed. Review the modification details to ensure that all phases of the removal process were completed successfully.

Note any additional actions that might be required to complete the removal, such as restarting the server or performing additional removal tasks. Click the Notifications icon to get additional information.








If any portion of the removal failed, note the reason for the failure. Review the Server Manager entries for removal problems and take corrective actions as appropriate.



# Using Windows PowerShell for Management

Server Manager's command-line counterpart is the ServerManager module for Windows PowerShell. At a Windows PowerShell prompt, you can obtain a detailed list

of a server's current state with regard to roles, role services, and features by entering **get-windowsfeature** . Each installed role, role service, and feature is highlighted and marked as such, and a management naming component in brackets follows the display name of each role, role service, and feature.



| Display Name                                        | Name                    | Install state    |
|-----------------------------------------------------|-------------------------|------------------|
| [ ] Active Directory Certificate Services           | AD-Certificate          | Available        |
| [ ]   Certificate Authority                         | ADCS-Cert-Authority     | Available        |
| [ ]   Certificate Enrollment Policy Web Service     | ADCS-Enroll-web-Pol     | Available        |
| [ ]   Certificate Enrollment Web Service            | ADCS-Enroll-web-Svc     | Available        |
| [ ]   Certificate Authority Web Enrollment          | ADCS-Web-Enrollment     | Available        |
| [ ]   Network Device Enrollment Service             | ADCS-Device-Enrollment  | Available        |
| [ ]   Online Responder                              | ADCS-Online-Cert        | Available        |
| [ ] Active Directory Domain Services                | AD-Domain-Services      | Available        |
| [ ] Active Directory Federation Services            | ADFS-Federation         | Available        |
| [ ] Active Directory Lightweight Directory Services | AD LDS                  | Available        |
| [ ] Active Directory Rights Management Services     | AD RMS                  | Available        |
| [ ]   Active Directory Rights Management Server     | AD RMS-Server           | Available        |
| [ ]   Identity Federation Support                   | AD RMS-Identity         | Available        |
| [ ] Device Health Attestation                       | DeviceHealthAttestat... | Available        |
| [ ] DHCP Server                                     | DHCP                    | UninstallPending |
| [ ] DNS Server                                      | DNS                     | Available        |
| [ ] Fax Server                                      | Fax                     | Available        |
| [X] File and Storage Services                       | FileAndStorage-Services | Installed        |
| [X]   File and ISCSI Services                       | File-Services           | Installed        |
| [X]   File Server                                   | FS-FileServer           | Installed        |

By using `Install-WindowsFeature` or `Uninstall-WindowsFeature` followed by the management name, you can install or uninstall a role, role service, or feature. For example, you can install Network Load Balancing by entering **install-windowsfeature nlb** . You can add **-includeallsubfeature** when installing components to add all subordinate role services or features. Management tools are not included by default. To add the management tools, add **-includemanagementtools** when installing components.

**REAL WORLD** Modules are imported implicitly by default. Otherwise, you need to import the module before you can use the cmdlets it provides. You import the ServerManager module by entering **Import-Module ServerManager** at the Windows PowerShell prompt. After the module is imported, you can use it with the currently running instance of Windows PowerShell. The next time you start Windows PowerShell, you need to import the module again if you want to use its features.

## Getting Binaries for Roles and Features

Binaries needed to install roles and features are referred to as payloads . With Windows

Server 2016, payloads are stored in subfolders of the

%SystemDrive%\Windows\WinSXS folder. Not only can you uninstall a role or feature, but you also can uninstall and remove the payload for a feature or role by using the – Remove parameter of the Uninstall-WindowsFeature cmdlet. Subcomponents of the role or feature are also removed. Add the - **includeallmanagementtools** parameter to also remove management tools.

When you want to install a role or feature, you can install the related components and restore any removed payloads for these components by using the Install-WindowsFeature cmdlet. By default, when you use Install-WindowsFeature, payloads are restored via Windows Update.

In the following example, you restore the Active Directory Directory Services (AD DS) binaries and all related subfeatures via Windows Update:

```
install-windowsfeature -name ad-domain-services -includeallsubfeature
```

You can use the –Source parameter to restore a payload from a Windows Imaging (WIM) format mount point. For example, if your enterprise has a mounted Windows Image for the edition of Windows Server 2016 you are working with available at the network path \\ImServer18\WinS16DC, you could specify the source as follows:

```
install-windowsfeature -name ad-domain-services -includeallsubfeature -  
source \\imserver18\wins16dc
```

Keep in mind that the path you specify is used only if required binaries are not found in the Windows Side-By-Side folder on the destination server. Although many large enterprises might have standard images that can be mounted by using network paths, you also can mount the Windows Server 2016 distribution media and use the Windows\WinSXS folder from the installation image as your source. To do this, follow these steps:

1. Insert the installation disc into the server's CD/DVD drive, and then create a folder to mount the Installation image by entering the following command: **mkdir c:\mountdir .**
2. Locate the index number of the image you want to use by entering the following command at an elevated command prompt, where e: is the drive designator of the server's disc drive: **dism /get-wiminfo /wimfile:e:\sources\install.wim.**
3. Mount the installation image by entering the following command at an elevated command prompt, where e: is the drive designator of the server's disc drive, 2 is the index of the image to use, and c:\mountdir is the mount directory: **dism /mount-wim /wimfile:e:\sources\install.wim /index:2**

**/mountdir:c:\mountdir /readonly** . Mounting the image might take several minutes.

4. Use Install-WindowsFeature at a Windows PowerShell command prompt with the source specified as **c:\mountdir\windows\winsxs** , as shown in this example:

```
install-windowsfeature -name ad-domain-services -includeallsubfeature -  
source c:\mountdir\windows\winsxs
```

Group Policy can be used to control whether Windows Update is used to restore payloads and to provide alternate source paths for restoring payloads. The policy you want to work with is Specify Settings For Optional Component Installation And Component Repair, which is under System in the Administrative Templates for Computer Configuration. This policy also is used for obtaining payloads needed to repair components.

If you enable this policy, you can do the following:

- Specify the alternate source file path for payloads as a network location. For network shares, enter the UNC path to the share, such as \\CorpServer82\WinServer2016\ . For mounted Windows images, enter the WIM path prefixed with **WIM:** and including the index of the image to use, such as WIM:\\CorpServer82\WinServer2016\install.wim:4.
- Specify that Windows Update should never be used to download payloads. If you enable the policy and use this option, you do not have to specify an alternate path. In this case, payloads cannot be obtained automatically and administrators will need to explicitly specify the alternate source path.
- Specify that Windows Update should be used for repairing components rather than Windows Server Update Services.

Specify settings for optional component installation and component repair

Specify settings for optional component installation and component repair

Previous Setting

Next Setting

☐ Not Configured

☒ Enabled

☐ Disabled

Comment:

Supported on:

At least Windows Server 2012, Windows 8 or Windows RT

Options:

Help:

Alternate source file path

WIM:\Server23\Server2016\install.wim:4

☒ Never attempt to download payload from Windows Update

☐ Download repair content and optional features directly from Windows Update instead of Windows Server Update Services (WSUS)

This policy setting specifies the network locations that will be used for the repair of operating system corruption and for enabling optional features that have had their payload files removed.

If you enable this policy setting and specify the new location, the files in that location will be used to repair operating system corruption and for enabling optional features that have had their payload files removed. You must enter the fully qualified path to the new location in the "Alternate source file path" text box. Multiple locations can be specified when each path is separated by a semicolon.





## Chapter 3. Configuring Server Settings


When it comes to system properties, Microsoft is still in the process of moving the furniture around, with the Settings app increasingly taking over tasks previously performed in Control Panel. That said, many essential administration tasks are still performed using the System console. You use the System console to view system information and perform basic configuration tasks. To access the System console, click System And Security in Control Panel and then click System. To access system-related options in Settings, click System and then select About.



# Getting System Information

The System console is divided into four basic areas that provide links for performing common tasks and a system overview:

- **Windows Edition** Shows the operating system edition and version, and lists any service packs you applied.
- **System** Lists the processor, memory, and type of operating system installed on the computer as well as the process architecture
- **Computer Name, Domain, And Workgroup Settings** Provides the computer name, description, domain, and workgroup details. If you want to change any of this information, click Change Settings, and then click Change in the System Properties dialog box.
- **Windows Activation** Shows whether you have activated the operating system and the product key. If Windows Server 2016 isn't activated yet, click the link provided to start the activation process and then follow the prompts.



When you're working in the System console, links in the left pane provide quick access to key support tools, including the following:



- **Device Manager**
- **Remote Settings**
- **Advanced System Settings**

In Settings, you can get similar information using by selecting System and then clicking About.



# Licensing & Activation

Although volume-licensed versions of Windows Server 2016 might not require activation or product keys, retail versions of Windows Server 2016 require both activation and product keys. If Windows Server 2016 has not been activated, you can use System to activate the operating system by selecting Activate Windows Now under Windows Activation. You also can activate Windows by entering **slmgr –ato** at a command prompt.



You can change the product key provided during installation of Windows Server 2016 to stay in compliance with your licensing plan or upgrade the server edition. In Settings, click Update & Security and then click Activation. Next, click Change Product Key and then enter the new product key.

At a command prompt, type **slmgr –ipk** followed by the product key you want to use, and then press Enter. When Windows finishes validating the product key, you need to reactivate the operating system.




NOTE The Windows Software Management Licensing tool has many other options, including options for offline activation by using a confirmation identifier.

To view this and other options, enter **slmgr** at a command prompt.



# Configuring Server Naming


In Settings, you can display and modify a server's network identification using the options on the About page. The About page shows the host name of the server and the organization membership. If a server is a member of a domain, the domain name is listed. Otherwise, the workgroup membership is listed.



## Renaming a Server

With a domain controller or a certificate authority, you should only change the computer name after removing the related role from the computer. Otherwise, you can change the computer name by following these steps:

1. In Settings, click System and then click About.
2. Click Rename PC.
3. Type the new name of the server and then click Next.
4. To finalize the name change, you must restart the server. Click either Restart Now or Restart Later as appropriate.



Use Rename-Computer to rename a server using Windows PowerShell. The basic syntax is:

```
Rename-computer -NewName NewName {AddtlParams}  
{AddtlParams}  
-ComputerName CurrentName -DomainCredential Domain\UserName -Restart
```

where NewName is the new computer name, CurrentName is the name of the remote server that you want to reconfigure, and Domain\UserName is an account with administrator permissions. Following this, you could change the name of the server you are logged on to and restart it by entering:

```
Rename-computer -NewName FileServer82 -Restart
```

Or the name of a remote server using:

```
Rename-computer -NewName FileServer18 -ComputerName CorpServer25 -Restart
```

Add a user account to use alternate credentials. You'll then be prompted for the account password.

## Joining a Domain


By default, servers are members of a workgroup named WORKGROUP. If needed, you

also can use Settings to join a server to a domain or change a server's domain membership status. To join a computer to a domain follow these steps:

1. In Settings, click System and then click About.
2. Click Join A Domain. When prompted, type the fully qualified name of the domain to join and then click Next.



3. Enter the name and password of an account with permission to add the computer to the previously specified domain and then click OK.



You will need to restart the computer for the domain membership change to be applied. If you have problems joining a computer to a domain:

- Check the TCP/IP settings for the primary network adapter and ensure the IP address, subnet, gateway and DNS server details are correct.
- Ensure connections aren't being blocked by the Windows Firewall running on the computer or any other firewall between the computer you are working with and the domain controller the computer is trying to communicate with.

**REAL WORLD** You also may want to determine whether the computer was manually assigned a primary DNS suffix. Why? At the time of this writing, an issue in software kept a server with a primary DNS suffix from joining a domain and the solution was to remove the primary DNS suffix, which is a property you can change in the System console but not in Settings.

Use Add-Computer to add a computer to a domain using Windows PowerShell. The basic syntax is:

```
Add-computer -Domain DomainName {AddtlParams}  
{AddtlParams}  
-ComputerName RemoteComputer -Credential Domain\UserName -Restart
```

where DomainName is the name of the domain to join, RemoteComputer is the name of the remote server that you want to reconfigure, and Domain\UserName is an account with administrator permissions. Following this, you could add FileServer82 to the ImaginedLands domain and restart it by entering::

```
Add-computer -ComputerName FileServer82 -DomainName ImaginedLands -Restart
```


Add a user account to use alternate credentials. You'll then be prompted for the account password.

# Managing Processor Scheduling

Processor scheduling determines the responsiveness of applications you are running interactively (as opposed to background applications that might be running on the system as services). For production servers, you'll typically want processor scheduling optimized for background services, which is the default setting. However, on development servers or when you are using Windows Server 2016 as your desktop operating system, you'll want to optimize processor scheduling for programs.

You configure processor scheduling by following these steps:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter. (Or, in Control Panel, click System And Security, click System, and then click Advanced System Settings.)
2. Click Settings in the Performance panel. This opens the Performance Options dialog box.
3. In the Performance Options dialog box, click the Advanced tab.
4. In the Processor Scheduling panel, you have the following options:
  - **Programs** Use this option to give the active application the best response time and the greatest share of available resources.
  - **Background Services** Use this option to give background applications a better response time than the active application.
5. Click OK.



## Allocating Virtual Memory

With virtual memory, you can use disk space to extend the amount of memory available on a system by using part of the hard disk as part of system memory. This feature writes



RAM to disks by using a process called paging . With paging, a set amount of RAM, such as 8192 megabytes (MB), is written to the disk as a paging file. The paging file can be accessed from the disk when needed in place of physical RAM.

An initial paging file is created automatically for the drive containing the operating system. By default, other drives don't have paging files, so you must create these paging files if you want them. When you create a paging file, you set an initial size and a maximum size. Paging files are written to the volume as a file named Pagefile.sys.




## Navigating Memory Baselines

Windows Server 2016 allocates virtual memory much differently than early releases of Windows Server. These changes have occurred for a variety of reasons, including because modern servers have physical memory that is significantly faster than physical disks and because Windows itself has changed. By default, Windows Server allocates virtual memory in an amount that is large enough to accommodate:


- **In Use Memory** This is memory used by processes, drivers or the operating system.
- **Standby Memory** This is memory that contains cached data and code that is not actively in use.
- **Modified Memory** This is memory whose contents must be written to disk before it can be reused.

To get details about the total memory available on a server as well as memory usage, you can use Task Manager, which is opened by right-clicking Start and then selecting Task Manager on the shortcut menu.



After opening Task Manager, click More Details to expand the view and get detailed information about processes, performance, users and services. On the Performance tab, you'll then be able to view resource availability and usage for the processor, memory and network adapters.





When you select Memory on the Performance tab, the total physical memory is listed in the upper right corner. In this example, the server has 16 GB of physical memory. The Memory Usage graph shows you how memory is being used over time and the Memory Composition graph shows you how memory is allocated.

If you move the mouse pointer over the various graphed allocations for memory composition, you'll see the exact in-use, standby and modified memory allocations. These allocations are also reflected in the In Use, Available, Committed and Cached values. As both standby memory and modified memory are a type of cached memory, the Cached value shows their combined allocation.

Virtual memory usage is reflected in the Committed, Cached, Paged Pool and Nonpaged Pool memory allocations. The Committed value is shown as a pair of values with the total virtual memory listed on the right. Total virtual memory reflects the combined value of total physical memory on the server and the total memory paged to disk.

Once you know more about how a server is using physical and virtual memory over time, you'll be better prepared to optimize the memory configuration. Generally, you want to allocate virtual memory in an amount that accommodates the workloads of the server. The virtual memory allocation can be given a minimum size and a maximum size. The typical default allocation allows the paging file to grow and shrink dramatically, as shown in the following example:



|                                       |         |
|---------------------------------------|---------|
| Total paging file size for all drives |         |
| Minimum allowed:                      | 16 MB   |
| Recommended:                          | 2938 MB |
| Currently allocated:                  | 2432 MB |

To ensure that paging files don't become fragmented, which can result in poor system performance, you can:

- Use a fixed virtual memory size rather than a dynamically allocated size. To do this, set the initial size and the maximum size of the virtual memory to the same value. This ensures that the paging file is consistent and can be written to a single contiguous file (if possible, given the amount of space on the volume).
- Use an initial virtual memory size that reflect baseline usage and a maximum size that allows the virtual memory to expand if needed. This ensures that the baseline paging file is consistent while allowing the server to better handle peak usage scenarios.

For industrial-strength enterprise servers, hardware vendors and your Microsoft support team typically will provide guidelines to help you manage the memory configuration. For other servers, here are some basic recommendations:


- In most cases, for computers with 8 gigabytes (GB) of RAM or less, I recommend setting the total paging file size so that it's twice the amount of physical RAM on the system. For instance, on a computer with 8 GB of RAM, you would ensure that the Total Paging File Size For All Drives setting is at least 16,384 MB.
- On systems with more than 8 GB of RAM, you should follow the hardware manufacturer's guidelines for configuring the paging file. Typically, this means setting the paging file to be the same size as physical memory.

## Configuring the Paging File Manually

You can configure virtual memory by following these steps:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter. (Or, in Control Panel, click System And Security, click System, and then click Advanced System Settings.)
2. Click Settings in the Performance panel. This opens the Performance Options dialog box.

3. In the Performance Options dialog box, click the Advanced tab, and then click Change to display the Virtual Memory dialog box.



The following information is provided:

- **Paging File Size For Each Drive** Provides information about the currently selected drive, and allows you to set its paging file size. Space Available indicates how much space is available on the drive.
  - **Drive [Volume Label] and Paging File Size** Show how virtual memory is currently configured on the system. Each volume is listed with its associated paging file (if any). The paging file range shows the initial and maximum size values set for the paging file.
  - **Total Paging File Size For All Drives** Provides a recommended size for virtual RAM on the system, and tells you the amount currently allocated. If this is the first time you're configuring virtual RAM, notice that the recommended amount has already been given to the system drive (in most instances).
4. By default, Windows Server manages the paging file size for all drives. If you want to configure virtual memory manually, clear the Automatically Manage Paging File Size For All Drives check box.
  5. In the Drive list, select the volume you want to work with.
  6. Select Custom Size, and then enter values in the Initial Size and Maximum Size boxes.
  7. Repeat steps 4–6 for each volume you want to configure.
  8. Click Set to save the changes.

**NOTE** The paging file is also used for debugging purposes when a Stop error occurs on the system. If the paging file on the system drive is smaller than the minimum amount required to write the debugging information to the paging file, this feature is disabled. If you want to use debugging, you should set the minimum size to equal the amount of RAM on the system. For example, a system with 8 GB of RAM would need a paging file of 8096 MB on the system drive.


9. Click OK three times to close the open dialog boxes. When you close the System utility, you'll get a prompt asking if you want to restart the system. Click Restart.



## Managing the Paging File Automatically

You can have Windows Server 2016 automatically manage virtual memory by following these steps:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter. (Or, in Control Panel, click System And Security, click System, and then click Advanced System Settings.)
2. Click Settings in the Performance panel. This opens the Performance Options dialog box.
3. In the Performance Options dialog box, click the Advanced tab, and then click Change to display the Virtual Memory dialog box.
4. Select the Automatically Manage Paging File Size For All Drives check box.



5. Click OK three times to close the open dialog boxes. When you close the System utility, you'll get a prompt asking if you want to restart the system. Click Restart.



## Configuring Data Execution Prevention

Windows Server 2016 requires a processor that includes hardware-based Data Execution Prevention (DEP) support. DEP is a memory-protection technology. DEP tells the computer's processor to mark all memory locations in an application as nonexecutable unless the location explicitly contains executable code. If code is executed from a memory page marked as nonexecutable, the processor can raise an exception and prevent the code from executing. This process prevents malicious code such as a virus from inserting itself into most areas of memory, because only specific areas of memory are marked as having executable code.

In Windows Server, memory randomization and other enhancements also prevent malware from inserting itself into startup and running processes. Windows Server uses address space layout randomization (ASLR) to randomly determine how and where important data is stored in memory, which makes it much more difficult for malware to find the specific locations in memory to attack.

**NOTE** Windows also supports DEP as implemented by Advanced Micro Devices, Inc. (AMD) processors that provide the no-execute page-protection (NX) processor feature. Such processors support the related instructions and must be running in physical address extension (PAE) mode.

# Using and configuring DEP

You can determine whether a computer supports DEP by using the System utility. If a computer supports DEP, you can also configure it by following these steps:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter.  
(Alternatively, in Control Panel, click System And Security and then click System. In the left pane, click Advanced System Settings.)
2. Click Settings in the Performance panel. This opens the Performance Options dialog box.
3. In the Performance Options dialog box, click the Data Execution Prevention tab. The text at the bottom of this tab indicates whether the computer supports execution protection.
4. If a computer supports execution protection and is configured appropriately, you can configure DEP by using the following options:
  - **Turn On DEP For Essential Windows Programs And Services Only** Enables DEP only for operating system services, programs, and components. This is the default and recommended option for computers that support execution protection and are configured appropriately.
  - **Turn On DEP For All Programs Except Those I Select** Configures DEP, and allows for exceptions. Select this option, and then click Add to specify programs that should run without execution protection. With this option, execution protection will work for all programs except those you select.
5. If you turned on DEP and allowed exceptions, you can add or remove 32-bit programs as exceptions:
  - To add a 32-bit program as an exception, click Add. Use the Open dialog box to find the executable file for the program you are configuring as an exception, and then click Open. The program is then added as an exception.
  - To remove a program previously added as an exception, click the program name, and then click Remove.
  - To temporarily disable a program as an exception (this might be necessary for troubleshooting), clear the check box next to the program name.
6. Click OK to save your settings.





## Understanding DEP Compatibility

To be compatible with DEP, applications must be able to mark memory explicitly with the Execute permission. Applications that cannot do this will not be compatible with the NX processor feature. If you experience memory-related problems running applications, you should determine which applications are having problems and configure them as exceptions rather than disable execution protection completely. This way, you still get the benefits of memory protection and can selectively disable memory protection for programs that aren't running properly with the NX processor feature.

Execution protection is applied to both user-mode and kernel-mode programs. A usermode execution protection exception results in a STATUS\_ACCESS\_VIOLATION exception. In most processes, this exception will be an unhandled exception, resulting in termination of the process. This is the behavior you want, because most programs violating these rules, such as a virus or worm, will be malicious in nature.


You cannot selectively enable or disable execution protection for kernel-mode device drivers the way you can with applications. Furthermore, on compliant 32-bit systems, execution protection is applied by default to the memory stack. On compliant 64-bit systems, execution protection is applied by default to the memory stack, the paged pool, and the session pool. A kernel-mode execution protection access violation for a device driver results in an ATTEMPTED\_EXECUTE\_OF\_NOEXECUTE\_MEMORY exception.

## Configuring Environment Variables

Windows uses environment variables to track important strings, such as a path where files are located or the logon domain controller host name. Environment variables defined for use by Windows—called system environment variables—are the same no matter who is logged on to a particular computer. Environment variables defined for use by users or programs—called user environment variables—are different for each user of a particular computer.

You configure system and user environment variables by means of the Environment Variables dialog box. To access this dialog box, open the System Properties dialog box, click the Advanced tab, and then click Environment Variables.






## Creating Environment Variables

You can create an environment variable by following these steps:


1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter.  
(Alternatively, in Control Panel, click System And Security and then click System. In the left pane, click Advanced System Settings.)
2. Click Environment Variables to open the Environment Variables dialog box.
3. Click New under User Variables or under System Variables, whichever is appropriate. This opens the New User Variable dialog box or the New System Variable dialog box, respectively.
4. In the Variable Name box, enter the variable name. In the Variable Value box, enter the variable value.
5. Click OK.



## Editing Environment Variables


You can edit most environment variables by following these steps:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter.  
(Alternatively, in Control Panel, click System And Security and then click System. In the left pane, click Advanced System Settings.)
2. Click Environment Variables to open the Environment Variables dialog box.
3. Select the variable in the User Variables or System Variables list.
4. Click Edit under User Variables or under System Variables, whichever is appropriate. The Edit User Variable dialog box or the Edit System Variable dialog box opens.
5. Enter a new value in the Variable Value box, and then click OK.



Some variables, such as Path, are exceptions. The Path variable controls the execution path for the operating system and to help prevent the variable from being updated incorrectly, Windows Server opens an editing dialog box that allows you to more easily manipulate the individual paths. Here you can:

- Click a path entry and then click Edit to modify that specific path value.
- Click a path entry and then click Move Up or Move Down to change the execution order.
- Click New to add a path entry to the end of the current path.
- Click Edit Text to edit the path variable directly as you do other variables.



**NOTE** When you create or modify environment variables, most of the variables are valid immediately after they are created or modified. With system variables, some changes take effect after you restart the computer. With user variables, some changes take effect the next time the user logs on to the system.

## Deleting Environment Variables


To delete an environment variable:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter.  
(Alternatively, in Control Panel, click System And Security and then click System. In the left pane, click Advanced System Settings.)
2. Click Environment Variables to open the Environment Variables dialog box.
3. Select the variable and then click Delete.

# Configuring Startup and Recovery

Startup and recovery options control the default operating system, the display of recovery options, and what actions the server performs in the event of catastrophic failure. You configure related options for servers using the Startup And Recovery dialog box. To access this dialog box:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter. (Or, in Control Panel, click System And Security, click System, and then click Advanced System Settings.)
2. Click Settings in the Startup And Recovery panel.




The screenshot shows the 'Startup and Recovery' dialog box with the following settings:

- System startup**
  - Default operating system: Windows Server 2016
  - ☒ Time to display list of operating systems: 30 seconds
  - ☐ Time to display recovery options when needed: 30 seconds
- System failure**
  - ☒ Write an event to the system log
  - ☒ Automatically restart
  - Write debugging information**
    - Automatic memory dump
    - Dump file: %SystemRoot%\MEMORY.DMP
    - ☒ Overwrite any existing file
    - ☐ Disable automatic deletion of memory dumps when disk space is low

Buttons: OK, Cancel

## Setting Startup Options

The System Startup area of the Startup And Recovery dialog box controls system startup. To specify the default operating system for a computer with multiple bootable operating systems, select one of the operating systems listed in the Default Operating System list. These options change the configuration settings used by the Windows Boot Manager.



Upon startup of a computer with multiple bootable operating systems, Windows Server displays the startup configuration menu for 30 seconds by default. You can change this by performing either of the following actions:

- Boot immediately to the default operating system by clearing the Time To Display List Of Operating Systems check box.
- Display the available options for a specific amount of time by selecting the Time To Display List Of Operating Systems check box and then setting a time delay in seconds.


On most systems, you'll generally want to use a value of 3 to 5 seconds. This is long enough for you to make a selection, yet short enough to expedite the system startup process.

When the system is in a recovery mode and booting, a list of recovery options might be displayed. As you can with the standard startup options, you can configure recovery startup options in one of two ways. You can set the computer to boot immediately by using the default recovery option by clearing the Time To Display Recovery Options When Needed check box, or you can display the available options for a specific amount of time by selecting Time To Display Recovery Options When Needed and then setting a time delay in seconds.

## Setting Recovery Options

You control system recovery with the System Failure and Write Debugging Information areas of the Startup And Recovery dialog box. Administrators use recovery options to control precisely what happens when the system encounters a non-recoverable system error (also known as a Stop error). The available options for the System Failure area are as follows:

- **Write An Event To The System Log** Logs the error in the system log, allowing administrators to review the error later using Event Viewer.
- **Automatically Restart** Select this option to have the system attempt to reboot when a fatal system error occurs.



**NOTE** Configuring automatic reboots isn't always a good thing. Sometimes you might want the system to halt rather than reboot to ensure that the system gets proper attention. Otherwise, you would know that the system rebooted only when you viewed the system logs or if you happened to be in front of the system's monitor when it rebooted.

You use the Write Debugging Information list to choose the type of debugging information you want to write to a dump file. You can use the dump file to diagnose system failures. The options are as follows:

- **None** Use this option if you don't want to write debugging information. **Active**
- **Memory Dump** Use this option to limit the dump to active physical memory. The dump file size depends on the amount of in-use physical memory.
- **Automatic Memory Dump** Use this option to let Windows determine which type of memory dump is best, and then create the dump file accordingly.

- **Complete Memory Dump** Use this option to dump all physical memory. The dump file size depends on the amount of physical memory being used, up to a maximum file size equal to the total physical RAM on the server.
- **Kernel Memory Dump** Use this option to dump the physical memory area being used by the Windows kernel. The dump file size depends on the size of the Windows kernel.
- **Small Memory Dump** Use this option to dump the physical memory segment in which the error occurred. This dump is 256 kilobytes (KB) in size.

If you elect to write to a dump file, you must also set a location for it. The default dump locations are %SystemRoot%\Minidump for small memory dumps and - %SystemRoot%\Memory.dmp for all other memory dumps. You'll also usually want to select Overwrite Any Existing File. Selecting this option ensures that any existing dump files are overwritten if a new Stop error occurs.

**BEST PRACTICES** You can create the dump file only if the system is properly configured. The system drive must have a sufficiently large memory-paging file (as set for virtual memory on the Advanced tab), and the drive the dump file is written to must have sufficient free space. For example, my server has 8 GB of RAM and requires a paging file on the system drive of the same size—8 GB. In establishing a baseline for kernel memory usage, I found that the server uses between 892 and 1076 MB of kernel memory. Because the same drive is used for the dump file, the drive must have at least 9 GB of free space to create a dump of debugging information. (That's 8 GB for the paging file and about 1 GB for the dump file.)

By default, memory dumps are deleted when disk space is low. Although this typically is a good idea, you may want to disable automatic deletion on a server experiencing system failures. This will ensure the dump file is available for troubleshooting. Select Disable Automatic Deletion of Memory Dumps to prevent the dump file from being removed when disk space is low.

## Chapter 4. Understanding Active Directory

Windows Server 2016 is designed to provide the services required for users and computers to work with network resources. Although Microsoft is increasingly expanding the role of the operating system into the cloud-integrated networks of the future, Windows servers currently are limited to specific deployment scenarios, all of which require Active Directory Domain Services (AD DS). In this chapter, you'll learn Active Directory fundamentals, regarding deployment options, domain structures and directory structure.





# Connecting Network Resources

Windows Server 2016 can be used with:

- On-premises deployments
- Hybrid deployments







## Working with On-Premises Deployments

An on-premises deployment requires Active Directory Domain Services (AD DS). AD DS provides the essential directory services necessary for establishing a domain, including the data store that stores information about objects on the network and makes that information available to users.

AD DS uses domain controllers to manage access to network resources. After users authenticate themselves by logging on to a domain, their stored credentials can be used to access resources on the network. Because AD DS is the heart of Active Directory and is required for directory-enabled applications and technologies, I typically refer to it simply as Active Directory rather than Active Directory Domain Services or AD DS.





## Working with Hybrid Deployments

A hybrid deployment requires extending Active Directory services to the cloud using Azure Active Directory Connect. By connecting to Azure, you integrate your onpremises directories into Microsoft's online services, including Office 365, allowing users to work with Office 365 and Azure as discussed in Office 365 & Exchange Online: Essentials for Administration (Stanek & Associates, 2016).

Azure AD Connect can be used with or without Active Directory Federation Services (AD FS) . Use Azure AD Connect without AD FS to simplify the architecture for streamlined deployments. Use Azure AD Connect with AD FS when you need to customize sign in options or work with multiple forests.

AD FS complements the authentication and access-management features of AD DS by extending them to the Internet and the World Wide Web. AD FS uses web agents to provide users with access to web applications and proxies to manage client access. After AD FS is configured, users can use their digital identities to authenticate themselves over the web.



# Working with Directory Services

Regardless of whether you use an on-premises or hybrid configuration, you must deploy Active Directory to establish the internal network structure—the basic components of which are:

- Domains, which provide the logical structure and are simply a group of computers that share a common directory database.
- Sites, which provide the physical structure and are simply a group of computers in one or more IP subnets.

Active Directory supports a multimaster replication model with servers dedicated to managing domain operations. These servers are called domain controllers .

In a multimaster replication model, any domain controller can process directory changes and then replicate those changes to other domain controllers automatically. Windows Server distributes an entire directory of information, called a data store . Inside the data store are sets of objects representing user, group, and computer accounts in addition to shared resources such as servers, files, and printers.

Domains that use Active Directory are referred to as Active Directory domains . Although Active Directory domains can function with only one domain controller, you can and should configure multiple domain controllers in the domain. This way, if one domain controller fails, you can rely on the other domain controllers to handle authentication and other critical tasks.

Windows Server also supports read-only domain controllers (RODCs) and restartable AD DS. An RODC is an additional domain controller that hosts a read-only replica of a domain's Active Directory data store.

RODCs are ideally suited to the needs of branch offices, where a domain controller's physical security cannot be guaranteed. Except for passwords, RODCs store the same objects and attributes as writable domain controllers. These objects and attributes are replicated to RODCs through unidirectional replication from a writable domain controller that acts as a replication partner.

Any domain user can be delegated as a local administrator of an RODC without granting any other rights in the domain. An RODC can act in the role of a global catalog but cannot act in the role of an operations master.

**NOTE** Because RODCs by default do not store passwords or credentials other than for their own computer account and the Kerberos Target (Krbtgt) account, RODCs pull user and computer credentials from a writable domain controller. If

allowed by a password replication policy that is enforced on the writable domain controller, an RODC retrieves and then caches credentials as necessary until the credentials change. Because only a subset of credentials is stored on an RODC, this limits the number of credentials that can possibly be compromised.

Restartable AD DS is a feature that allows an administrator to start and stop AD DS. In the Services console, the Active Directory Domain Services service is available on domain controllers, allowing you to easily stop and restart AD DS in the same way as for any other service that is running locally on the server. While AD DS is stopped:

- You can perform maintenance tasks that would otherwise require restarting the server, such as performing offline defragmentation of the Active Directory database, applying updates to the operating system, or initiating an authoritative restore.
- Other domain controllers can handle authentication and logon tasks. Cached credentials, smart cards, and biometric logon methods continue to be supported. If no other domain controller is available and none of these logon methods applies, you can still log on to the server by using the Directory Services Restore Mode account and password.

All domain controllers support restartable AD DS—even RODCs. Because of restartable AD DS, domain controllers have three possible states:

- **Active Directory Started** Active Directory is started, and the domain controller has the same running state as a domain controller running Windows 2000 Server or Windows Server 2003. This allows the domain controller to provide authentication and logon services for a domain.
- **Active Directory Stopped** Active Directory is stopped, and the domain controller can no longer provide authentication and logon services for a domain. This mode shares some characteristics of both a member server and a domain controller in Directory Services Restore Mode. As with a member server, the server is joined to the domain. Users can log on interactively by using cached credentials, smart cards, and biometric logon methods. Users can also log on over the network by using another domain controller for domain logon. As with Directory Services Restore Mode, the Active Directory database (Ntds.dit) on the local domain controller is offline. This means you can perform offline AD DS operations, such as defragmentation of the database and application of security updates, without having to restart the domain controller.
- **Directory Services Restore Mode** Active Directory is in restore mode. The domain controller has the same restore state as a domain controller running



Windows Server 2003. This mode allows you to perform an authoritative or nonauthoritative restore of the Active Directory database.

When working with AD DS in the Stopped state, you should keep in mind that dependent services are also stopped. This means that File Replication Service (FRS), Kerberos Key Distribution Center (KDC), and Intersite Messaging are stopped before Active Directory is stopped, and that even if they are running, these dependent services are restarted when Active Directory restarts. Further, you can restart a domain controller in Directory Services Restore Mode, but you cannot start a domain controller in the Active Directory Stopped state. To get to the Stopped state, you must first start the domain controller in the customary way and then stop AD DS.





# Deploying Directory Services

For many years now, Active Directory has been the heart of Windows–based domains.

Just about every administrative task you perform affects Active Directory in some way. Active Directory technology is based on standard Internet protocols and is designed to help you clearly define your network's structure.





# Using Active Directory with DNS


DNS is an integral part of Active Directory technology—so much so that you must configure DNS on the network before you can install Active Directory. DNS is a standard Internet service that organizes groups of computers into domains. DNS domains are organized into a hierarchical structure. The DNS domain hierarchy is defined on an Internet-wide basis, and the different levels within the hierarchy identify computers, organizational domains, and top-level domains. DNS is also used to map host names to numeric TCP/IP addresses. Through DNS, an Active Directory domain hierarchy can also be defined on an Internet-wide basis, or the domain hierarchy can be separate from the Internet and private.

When you refer to computer resources in a DNS domain, you use a fully qualified domain name (FQDN), such as zeta.imagedlands.com. Here, zeta represents the name of an individual computer, imagedlands represents the organizational domain, and com is the top-level domain. Top-level domains (TLDs) are at the base of the DNS hierarchy. TLDs are organized geographically by using two-letter country codes, such as CA for Canada; by organization type, such as com for commercial organizations; and by function, such as mil for U.S. military installations.

Normal domains, such as imagedlands.com, are also referred to as parent domains because they're the parents of an organizational structure. You can divide parent domains into subdomains, which you can then use for different offices, divisions, or geographic locations. For example, the FQDN for a computer at the Seattle office for Microsoft could be designated as jacob.seattle.imagedlands.com. Here, jacob is the computer name, seattle is the subdomain, and imagedlands.com is the parent domain. Another term for a subdomain is a child domain .

# Implementing Directory Services

With Windows Server 2016, you install Active Directory in a two-part process. First, you start the process in Server Manager by clicking Manage and then clicking Add Roles And Features. This runs the Add Roles And Features Wizard, which you use to specify that you want to add the AD DS role to the server. The wizard installs binaries needed for the role, and the progress of this process is shown on the Installation Progress page.




**REAL WORLD** Binaries needed to install roles and features are referred to as payloads. With Windows Server 2016, not only can you uninstall a role or feature, but you also can uninstall and remove the payload for that feature or role by using the – Remove parameter of the Uninstall-WindowsFeature cmdlet.

You can restore a removed payload by using Install-WindowsFeature. Payloads are restored via Windows Update by default. Use the –Source parameter to restore a payload from a WIM mount point. In the following example, you restore the AD DS binaries and all related subfeatures via Windows Update.

```
install-windowsfeature -name ad-domain-services -includeallsubfeature
```

When the installation completes, you start the Active Directory Domain Services Configuration Wizard by clicking the Promote This Server To A Domain Controller link on the Installation Progress page, and then you use this wizard to configure the role.



**NOTE** If you close the wizard, click the Notifications icon in Server Manager, and then click the click Task Details and then click the Promote This Server To A Domain Controller link provided for the Post-Deployment Configuration task.

As part of any Active Directory deployment, the metadata associated with directory structures may need to be updates. These metadata elements are part of the directory schema and the utility that performs schema updates is Adprep.exe. The Active Directory Domain Services Configuration Wizard will also run Adprep.exe to prepare schema as appropriate.

If you have not run Adprep.exe separately previously, and you are installing the first domain controller that runs Windows Server 2016 in an existing domain or forest, the wizard will prompt you to supply credentials to run Adprep commands. To prepare a forest, you need to provide credentials for a member of the Enterprise Admins group, the Schema Admins group, and the Domain Admins group in the domain that hosts the schema master. To prepare a domain, you need to provide credentials for a member of the Domain Admins group. If you are installing the first read-only domain





controller (RODC) in a forest, you need to provide credentials for a member of the Enterprise Admins group.

As discussed previously, Active Directory uses DNS for name resolution. If DNS isn't already installed, you are prompted to install it as part of installing AD DS. If no domain exists, the wizard helps you create a domain and configure Active Directory in the new domain. The wizard can also help you add child domains to existing domain structures. To verify that a domain controller is installed correctly, do the following:

- Check the Directory Service event log for errors.
- Ensure that the SYSVOL folder is accessible to clients.
- Verify that name resolution is working through DNS.
- Verify the replication of changes to Active Directory.

**NOTE** In the rest of this chapter, I often use the terms directory and domains to refer to Active Directory and Active Directory domains, respectively, except when I need to distinguish Active Directory structures from DNS or other types of directories.

Keep in mind that when you use Server Manager for Windows Server 2016, any necessary preparations are done automatically when you deploy a domain controller. This means the Configuration Wizard automatically updates the Active Directory schema for the forest and domain so that the schema is compatible with Windows Server 2016 as necessary.

Any domain controller running Windows Server can be configured as a read-only domain controller (RODC). When you install the DNS Server service on an RODC, the RODC can act as a read-only DNS (RODNS) server. In this configuration, the following conditions are true:

- The RODC replicates the application directory partitions that DNS uses, including the ForestDNSZones and DomainDNSZones partitions. Clients can query an RODNS server for name resolution. However, the RODNS server does not support client updates directly, because the RODNS server does not register resource records for any Active Directory–integrated zone that it hosts.
- When a client attempts to update its DNS records, the server returns a referral. The client can then attempt to update against the DNS server that is provided in the referral. Through replication in the background, the RODNS server then attempts to retrieve the updated record from the DNS server that made the update. This replication request is only for the changed DNS record. The entire list of data changed in the zone or domain is not replicated during this special request.



It's important to point out that the first domain controller installed in a forest or domain cannot be an RODC. However, you can configure subsequent domain controllers as read-only.





# Working with Domain Structures

Active Directory provides both logical and physical structures for network components. Logical structures help you organize directory objects and manage network accounts and shared resources. Logical structures include the following:

- **Organizational units** A subgroup of domains that often mirrors the organization's business or functional structure.
- **Domains** A group of computers that share a common directory database.
- **Domain trees** One or more domains that share a contiguous namespace.
- **Domain forests** One or more domain trees that share common directory information.

Physical structures serve to facilitate network communication and to set physical boundaries around network resources. Physical structures that help you map the physical network structure include the following:

- **Subnets** A network group with a specific IP address range and network mask.
- **Sites** One or more subnets. Sites are used to configure directory access and replication.




# Navigating Domain and Forest Structures

Domains and forests are integral parts of any directory. As discussed previously, an Active Directory domain is simply a group of computers that share a common directory database. Active Directory domain names must be unique. For example, you can't have two microsoft.com domains, but you can have a parent domain microsoft.com, with the child domains seattle.microsoft.com and ny.microsoft.com.

If the domain is part of a private network, the name assigned to a new domain must not conflict with any existing domain name on the private network. If the domain is part of the Internet, the name assigned to a new domain must not conflict with any existing domain name throughout the Internet. To ensure uniqueness on the Internet, you must register the parent domain name before using it. You can register a domain through any designated registrar. You can find a current list of designated registrars at InterNIC ( [www.internic.net](http://www.internic.net) ).

Each domain has its own security policies and trust relationships with other domains. Domains can also span more than one physical location, which means that a domain can consist of multiple sites and those sites can have multiple subnets. Within a domain's directory database, you'll find objects defining accounts for users, groups, computers, and shared resources such as printers and folders.



Domains have an operating mode which determines that availability of various features.


This operating mode is referred to as the domain functional level.

Domain functions are limited and controlled by the domain functional level. Several domain functional levels are available, including the following:


- **Windows Server 2008** Supports domain controllers running Windows Server 2008 and later.
- **Windows Server 2008 R2** Supports domain controllers running Windows Server 2008 R2 and later.
- **Windows Server 2012** Supports domain controllers running Windows Server 2012 and later.
- **Windows Server 2012 R2** Supports domain controllers running Windows Server 2012 R2 and later.
- **Windows Server 2016** Supports domain controllers running Windows Server 2016.

Each Active Directory domain has a DNS domain name, such as imaginedlands.com. One or more domains sharing the same directory data are referred to as a forest. The domain names within this forest can be noncontiguous or contiguous in the DNS naming hierarchy.


When domains have a contiguous naming structure, they're said to be in the same domain tree. In the example that follows, the root domain tvpress.com has two child domains: seattle.tvpress.com and ny.tvpress.com. These domains, in turn, have subdomains. All the domains are part of the same tree because they have the same root domain.



If the domains in a forest have noncontiguous DNS names, they form separate domain trees within the forest. A domain forest can have one or more domain trees. In this example, the tvpress.com and reagentpress.com domains form the roots of separate domain trees in the same forest.



You can access domain structures by using Active Directory Domains And Trusts. Start this snap-in for the Microsoft Management Console (MMC) from the Tools menu in Server Manager. When working with Active Directory Domains And Trusts, you'll find separate entries for each root domain. Right-click a domain entry and select Properties to get additional information about the domain.



Like domains, forests have an operating mode which determines that availability of various features. This operating mode is referred to as the forest functional level.

Forest functions are limited and controlled by the forest functional level. Several forest functional levels are available, including the ones listed here:

- **Windows Server 2008** Supports domain controllers running Windows Server 2008 and later
- **Windows Server 2008 R2** Supports domain controllers running Windows Server 2008 R2 and later



- **Windows Server 2012** Supports domain controllers running Windows Server 2012 and later
- **Windows Server 2012 R2** Supports domain controllers running Windows Server 2012 R2 and later
- **Windows Server 2016** Supports domain controllers running Windows Server 2016

Generally, you cannot lower the domain or forest functional level after you raise it. However, when you raise the domain or forest functional level to Windows Server 2012 or higher, you can lower it. If you are using Windows Server 2016 domain or forest functional level, you can lower it to Windows Server 2012 R2, Windows Server 2012 or Windows Server 2008 R2. If you are using Windows Server 2012 R2 domain or forest functional level, you can lower it to Windows Server 2012 or Windows Server 2008 R2. Additionally, with either scenario, if Active Directory Recycle Bin has not been enabled, you can lower the domain or forest functional level to Windows Server 2008.



## Using Organizational Units

Organizational units (OUs) are subgroups within domains that often mirror an organization's functional or business structure. You can also think of OUs as logical containers into which you place accounts, shared resources, and other OUs. For example, you could create OUs named HumanResources, IT, Engineering, and Marketing for the microsoft.com domain. You could later expand this scheme to include child units. Child OUs for Marketing could include OnlineSales, ChannelSales, and PrintSales.

Objects placed in an OU can come only from the parent domain. For example, OUs associated with seattle.microsoft.com can contain objects for this domain only. You can't add objects from ny.microsoft.com to these containers, but you could create separate OUs to mirror the business structure of seattle.microsoft.com.


OUs are helpful in organizing objects to reflect a business or functional structure. Still, this isn't the usingonly reason to use OUs. Other reasons include the following:

- OUs allow you to assign group policies to a small set of resources in a domain without applying the policies to the entire domain. This helps you set and manage group policies at the appropriate level in the enterprise.
- OUs create smaller, more manageable views of directory objects in a domain. This helps you manage resources more efficiently.
- OUs allow you to delegate authority and to easily control administrative access to domain resources. This helps you control the scope of administrator privileges in the domain. You could grant user A administrative authority for one OU and not for others. Meanwhile, you could grant user B administrative authority for all OUs in the domain.

OUs are represented as folders in Active Directory Users And Computers. You can start this snap-in for MMC from the Tools menu in Server Manager. Right-click a domain or OU node to create a new directory object in the domain or OU.







## Understanding Sites and Subnets

As discussed previously, a site is a group of computers in one or more IP subnets. You use sites to map your network's physical structure. Site mappings are independent of logical domain structures, so there's no necessary relationship between a network's physical structure and its logical domain structure. With Active Directory, you can create multiple sites within a single domain or create a single site that serves multiple domains. The IP address ranges used by a site and the domain namespace also have no connection.

You can think of a subnet as a group of network addresses. Unlike sites, which can have multiple IP address ranges, subnets have a specific IP address range and network mask. Subnet names are shown in the form `network / bits-masked`, such as `192.168.19.0/24`. In this case, the network address `192.168.19.0` and network mask `255.255.255.0` are combined to create the subnet name `192.168.19.0/24`.


**NOTE** Don't worry, you don't need to know how to create a subnet name. In most cases, you enter the network address and the network mask, and then Windows Server generates the subnet name for you.

Computers are assigned to sites based on their location in a subnet or a set of subnets. If computers in subnets can communicate efficiently with one another over the network, they're said to be well connected. Ideally, sites consist of subnets and computers that are all well connected. If the subnets and computers aren't well connected, you might need to set up multiple sites. Being well connected gives sites several advantages:

- When clients log on to a domain, the authentication process first searches for domain controllers that are in the same site as the client. This means that local domain controllers are used first, if possible, which localizes network traffic and can speed up the authentication process.

- Directory information is replicated more frequently within sites than between sites. This reduces the network traffic load caused by replication while ensuring that local domain controllers get up-to-date information quickly. You can also use site links to customize how directory information is replicated between sites. A domain controller that is designated to perform intersite replication is called a bridgehead server . By designating a bridgehead server to handle replication between sites, you place the bulk of the intersite replication burden on a specific server rather than on any available server in a site.

You access sites and subnets through Active Directory Sites And Services. Because this is a snap-in for the MMC, you can add it to any updateable console. You can also open Active Directory Sites And Services from the Tools menu in Server Manager.





# Understanding the Directory Structure

Active Directory has many components and is built on many technologies. Directory data is made available to users and computers through data stores and global catalogs. Although most Active Directory tasks affect the data store, global catalogs are equally important because they're used during logon and for information searches. In fact, if the global catalog is unavailable, standard users can't log on to the domain. The only way to change this behavior is to cache universal group membership locally. As you might expect, caching universal group membership has advantages and disadvantages, which I'll discuss in a moment.

You access and distribute Active Directory data by using directory access protocols and replication. Directory access protocols allow clients to communicate with computers running Active Directory. Replication is necessary to ensure that updates to data are distributed to domain controllers. Although multimaster replication is the primary technique you use to distribute updates, some changes to data can be handled only by individual domain controllers called operations masters. Application directory partitions also change the way multimaster replication works.

With application directory partitions, enterprise administrators (those belonging to the Enterprise Admins group) can create replication partitions in the domain forest. These partitions are logical structures used to control the replication of data within a domain forest. For example, you could create a partition to strictly control the replication of DNS information within a domain, thereby preventing other systems in the domain from replicating DNS information.

An application directory partition can appear as a child of a domain, a child of another application partition, or a new tree in the domain forest. Replicas of the application directory partition can be made available on any Active Directory domain controller, including those acting as global catalog servers. Although application directory partitions are useful in large domains and forests, they add overhead in terms of planning, administration, and maintenance.

## Exploring the Data Store

The data store contains information about objects, such as accounts, shared resources, OUs, and group policies. Another name for the data store is the directory, which refers to Active Directory itself.

Domain controllers store the directory in a file called Ntds.dit. This file's location is set when Active Directory is installed, and it should be on an NTFS file system drive formatted for use with Windows Server. You can also save directory data separately





from the main data store. This is true for group policies, scripts, and other types of public information stored on the shared system volume (SYSVOL).

Sharing directory information is called publishing . For example, you publish information about a printer by sharing the printer over the network. Similarly, you publish information about a folder by sharing the folder over the network.

Domain controllers replicate most changes to the data store in multimaster fashion. Administrators for small or medium-size organizations rarely need to manage replication of the data store. Replication is handled automatically, but you can customize it to meet the needs of large organizations or organizations with special requirements.

Not all directory data is replicated. Instead, only public information that falls into one of the following three categories is replicated:

- **Domain data** Contains information about objects within a domain. This includes objects for accounts, shared resources, organizational units, and group policies. •
- **Configuration data** Describes the directory's topology. This includes a list of all domains, domain trees, and forests, and also the locations of the domain controllers and global catalog servers.
- **Schema data** Describes all objects and data types that can be stored in the directory. The default schema provided with Windows Server describes account objects, shared resource objects, and more. You can extend the default schema by defining new objects and attributes or by adding attributes to existing objects.

## Exploring Global Catalogs

When universal group membership isn't cached locally, global catalogs enable network logon by providing universal group membership information when a logon process is initiated. Global catalogs also enable directory searches throughout the domains in a forest. A domain controller designated as a global catalog stores a full replica of all objects in the directory for its host domain and a partial replica for all other domains in the domain forest.

**NOTE** Partial replicas are used because only certain object properties are needed for logon and search operations. Partial replication also means that less information needs to be circulated on the network, reducing the amount of network traffic.

By default, the first domain controller installed on a domain is designated as the global catalog. If only one domain controller is in the domain, the domain controller





and the global catalog are the same server. Otherwise, the global catalog is on the domain controller you've configured as such. You can also add global catalogs to a domain to help improve response time for logon and search requests. The recommended technique is to have one global catalog per site within a domain.

Domain controllers hosting the global catalog should be well connected to domain controllers acting as infrastructure masters. The role of infrastructure master is one of the five operations master roles you can assign to a domain controller. In a domain, the infrastructure master is responsible for updating object references. The infrastructure master does this by comparing its data with that of a global catalog. If the infrastructure master finds outdated data, it requests updated data from a global catalog. The infrastructure master then replicates the changes to the other domain controllers in the domain. For more information about operations master roles, see the "Understanding Operations Master Roles" section later in this chapter.

When only one domain controller is in a domain, you can assign the infrastructure master role and the global catalog to the same domain controller. When two or more domain controllers are in the domain, however, the global catalog and the infrastructure master must be on separate domain controllers. If they aren't, the infrastructure master won't find out-of-date data and will never replicate changes. The only exception is when all domain controllers in the domain host the global catalog. In this case, it doesn't matter which domain controller serves as the infrastructure master.

One of the key reasons to configure additional global catalogs in a domain is to ensure that a catalog is available to service logon and directory search requests. Again, if the domain has only one global catalog and the catalog isn't available, and there's no local caching of universal group membership, standard users can't log on, and those who are logged on can't search the directory. In this scenario, the only users who can log on to the domain when the global catalog is unavailable are members of the Domain Admins group.

Searches in the global catalog are very efficient. The catalog contains information about objects in all domains in the forest. This allows directory search requests to be resolved in a local domain rather than in a domain in another part of the network. Resolving queries locally reduces the network load and allows for quicker responses in most cases.

**TIP** If you notice slow logon or query response times, you might want to configure additional global catalogs. But more global catalogs usually means more replication data being transferred over the network.





## Universal Group Membership Caching

In a large organization, having global catalogs at every office location might not be practical. Not having global catalogs at every office location presents a problem, however, if a remote office loses connectivity with the main office or a designated branch office where global catalog servers reside. If this occurs, standard users won't be able to log on; only members of Domain Admins will be able to log on. This happens because logon requests must be routed over the network to a global catalog server at a different office, and this isn't possible with no connectivity.

As you might expect, you can resolve this problem in many ways. You can make one of the domain controllers at the remote office a global catalog server by following the procedure discussed in the "Configuring Global Catalogs" section in Chapter 6. The disadvantage of this approach is that the designated server or servers will have an additional burden placed on them and might require additional resources. You also have to manage more carefully the up time of the global catalog server.

Another way to resolve this problem is to cache universal group membership locally. Here, any domain controller can resolve logon requests locally without having to go through a global catalog server. This allows for faster logons and makes managing server outages much easier because your domain isn't relying on a single server or a group of servers for logons. This solution also reduces replication traffic. Instead of replicating the entire global catalog periodically over the network, only the universal group membership information in the cache is refreshed. By default, a refresh occurs every eight hours on each domain controller that's caching membership locally.

Universal group membership caching is site-specific. Remember, a site is a physical directory structure consisting of one or more subnets with a specific IP address range and network mask. The domain controllers running Windows Server and the global catalog they're contacting must be in the same site. If you have multiple sites, you need to configure local caching in each site. To learn how to configure caching, see the "Configuring Universal Group Membership Caching" section in Chapter 6.

## Replication and Active Directory

Active Directory stores three types of information: domain data, schema data, and configuration data. Domain data is replicated to all domain controllers within a particular domain. Schema and configuration data are replicated to all domains in the domain tree or forest. In addition, all objects in an individual domain and a subset of object properties in the domain forest are replicated to global catalogs.

This means that domain controllers store and replicate the following:





- Schema information for the domain tree or forest
- Configuration information for all domains in the domain tree or forest
- All directory objects and properties for their respective domains

However, domain controllers hosting a global catalog store and replicate schema information for the forest and configuration information for all domains in the forest. They also store and replicate a subset of the properties for all directory objects in the forest that's replicated only between servers hosting global catalogs and all directory objects and properties for their respective domain:

- Schema information for the forest
- Configuration information for all domains in the forest
- A subset of the properties between global catalogs
- All directory objects and properties for their domain

To get a better understanding of replication, consider the following scenario, in which you're installing a new network:

1. Start by installing the first domain controller in domain A. The server is the only domain controller and also hosts the global catalog. No replication occurs because no other domain controllers are on the network.
2. Install a second domain controller in domain A. Because there are now two domain controllers, replication begins. To make sure that data is replicated properly, assign one domain controller as the infrastructure master and the other as the global catalog. The infrastructure master watches for updates to the global catalog and requests updates to changed objects. The two domain controllers also replicate schema and configuration data.
3. Install a third domain controller in domain A. This server isn't a global catalog. The infrastructure master watches for updates to the global catalog, requests updates to changed objects, and then replicates those changes to the third domain controller. The three domain controllers also replicate schema and configuration data.
4. Install a new domain, domain B, and add domain controllers to it. The global catalog hosts in domain A and domain B begin replicating all schema and configuration data in addition to a subset of the domain data in each domain. Replication within domain A continues as previously described. Replication within domain B begins.



# Active Directory and LDAP

The Lightweight Directory Access Protocol (LDAP) is a standard Internet communications protocol for TCP/IP networks. LDAP is designed specifically for accessing directory services with the least amount of overhead. LDAP also defines operations that can be used to query and modify directory information.

Active Directory clients use LDAP to communicate with computers running Active Directory whenever they log on to the network or search for shared resources. You can also use LDAP to manage Active Directory.

LDAP is an open standard that many other directory services use. This makes interdirectory communications easier and provides a clearer migration path from other directory services to Active Directory. You can also use Active Directory Service Interfaces (ADSI) to enhance interoperability. ADSI supports the standard application programming interfaces (APIs) for LDAP that are specified in Internet standard Request for Comments (RFC) 1823. You can use ADSI with scripting and Windows PowerShell to create and manage objects in Active Directory.





# Understanding Operations Master Roles

Operations master roles accomplish tasks that are impractical to perform in multimaster fashion. Five operations master roles are defined, and you can assign these roles to one or more domain controllers. Although certain roles can be assigned only once in a domain forest, other roles must be defined once in each domain.

Every Active Directory forest must have the following roles:

- **Schema master** Controls updates and modifications to directory schema. To update directory schema, you must have access to the schema master. To determine which server is the current schema master for the domain, start a command prompt and enter **dsquery server -hasfsmo schema** .
- **Domain naming master** Controls the addition or removal of domains in the forest. To add or remove domains, you must have access to the domain naming master. To determine which server is the current domain naming master for the domain, start a command prompt and enter **dsquery server -hasfsmo name** .

These forestwide roles must be unique in the forest. This means that you can assign only one schema master and one domain naming master in a forest.

Every Active Directory domain must have the following roles:

- **Relative ID master** Allocates relative IDs to domain controllers. Whenever you create a user, group, or computer object, domain controllers assign a unique security ID to the related object. The security ID consists of the domain's security ID prefix and a unique relative ID allocated by the relative ID (RID) master. To determine which server is the current relative ID master for the domain, start a command prompt and enter **dsquery server -hasfsmo rid** .
- **PDC emulator** When you use mixed-mode or interim-mode operations, the PDC emulator acts as a Windows NT PDC. Its job is to authenticate Windows NT logons, process password changes, and replicate updates to BDCs. The PDC emulator is the default time server and, as such, also performs time synchronization in a domain. To determine which server is the current PDC emulator for the domain, start a command prompt and enter **dsquery server -hasfsmo pdc** .
- **Infrastructure master** Updates object references by comparing its directory data with that of a global catalog. If the data is outdated, the infrastructure master requests updated data from a global catalog and then replicates the changes to the other domain controllers in the domain. To determine which server is the current infrastructure operations master for the domain, start a command prompt and enter **dsquery server hasfsmo infr** .



These domainwide roles must be unique in each domain. This means that you can assign only one relative ID master, one PDC emulator, and one infrastructure master in each domain.

Operations master roles are usually assigned automatically, but you can reassign them. When you install a new network, the first domain controller in the first domain is assigned all the operations master roles. If you later create a child domain or a root domain in a new tree, the first domain controller in the new domain is automatically assigned operations master roles, too. In a new domain forest, the domain controller is assigned all operations master roles. If the new domain is in the same forest, the assigned roles are relative ID master, PDC emulator, and infrastructure master. The schema master and domain naming master roles remain in the first domain in the forest.

When a domain has only one domain controller, that computer handles all the operations master roles. If you're working with a single site, the default operations master locations should be sufficient. As you add domain controllers and domains, however, you'll probably want to move the operations master roles to other domain controllers.

When a domain has two or more domain controllers, you should configure two domain controllers to handle operations master roles. In this case, you make one domain controller the operations master and designate the second as your standby operations master. The standby operations master can then be used if the primary one fails. Be sure that the domain controllers are direct replication partners and are well connected.

As the domain structure grows, you might want to split up the operations master roles and place them on separate domain controllers. This can improve the responsiveness of the operations masters. Pay particular attention to the current responsibilities of the domain controller you plan to use.

**BEST PRACTICES** Two roles you should not separate are schema master and domain naming master. Always assign these roles to the same server. For the most efficient operations, you also usually want the RID master and PDC emulator to be on the same server. But you can separate these roles if necessary. For example, on a large network where peak loads are causing performance problems, you probably want to place the RID master and PDC emulator on separate domain controllers. Additionally, you usually shouldn't place the infrastructure master on a domain controller hosting a global catalog. See the "Exploring global catalogs" section earlier in this chapter for details.





## Chapter 5. Managing Active Directory

Core Active Directory management focuses on key tools that you use routinely and the core tasks you perform with those tools. In this chapter, you'll learn about the tools you can use to manage Active Directory and about specific techniques for managing domain controllers and organizational units (OUs).





# Working with Active Directory Domains

Although you must configure both Active Directory and DNS on a Windows Server network, Active Directory domains and DNS domains have different purposes. Active Directory domains help you manage accounts, resources, and security. DNS domains establish a domain hierarchy that is primarily used for name resolution. Windows Server uses DNS to map host names, such as zeta.imaginedlands.com, to numeric TCP/IP addresses, such as 192.16.18.8.





## Accessing Active Directory

User computers running professional or business editions of Windows can make full use of Active Directory. These computers access the network as Active Directory clients and have full use of Active Directory features. As clients, these systems can use transitive trust relationships that exist within the domain tree or forest. A transitive trust is one that isn't established explicitly. Rather, the trust is established automatically based on the forest structure and permissions set in the forest. These relationships allow authorized users to access resources in any domain in the forest.

Server computers provide services to other systems and can act as domain controllers or member servers. A domain controller is distinguished from a member server because it runs Active Directory Domain Services.

Domains can have one or more domain controllers. When a domain has multiple domain controllers, the controllers automatically replicate directory data with one another by using a multimaster replication model. This model allows any domain controller to process directory changes and then replicate those changes to other domain controllers.

Because of the multimaster domain structure, all domain controllers have equal responsibility by default. You can, however, give some domain controllers precedence over others for certain tasks, such as specifying a bridgehead server that has priority in replicating directory information to other sites. In addition, some tasks are best performed by a single server. A server that handles this type of task is called an operations master. There are five flexible single master operations (FSMO) roles, and you can assign each to a different domain controller. For more information, see the “Understanding Operations Master Roles” section in Chapter 4.

Every Windows-based computer that joins a domain has a computer account. Like other resources, computer accounts are stored in Active Directory as objects. You use computer accounts to control access to the network and its resources. A computer accesses a domain by using its account, which is authenticated before the computer can access the network.

**REAL WORLD** Domain controllers use Active Directory's global catalog to authenticate both computer and user logons. If the global catalog is unavailable, only members of the Domain Admins group can log on to the domain because the universal group membership information is stored in the global catalog, and this information is required for authentication. In Windows Server 2003 and later servers, you have the option of caching universal group membership locally,





which solves this problem. For more information, see the “Understanding the Directory Structure” section in Chapter 4.

## Navigating Functional Levels

Before deploying domain controllers running Windows Server 2016, you may want to consider the network impact. When you install a domain controller using a newer operating system than previously available in a domain, the schema for Active Directory typically has to be updated in the forest and in the affected domain. This upgrade process can generate a lot of network traffic while information is being replicated. Sometimes the entire upgrade process can take 15 minutes or longer. During this time, you might experience delayed responsiveness when communicating with servers and higher latency on the network, so you might want to schedule the upgrade outside normal business hours. It’s also a good idea to thoroughly test compatibility with existing applications (especially legacy applications) beforehand.

When you configure a domain or forest, you set the functional level. As there is one functional level available for each current version of Windows Server, the available functional levels are as follows:

- [Windows Server 2008](#)
- [Windows Server 2008 R2](#)
- [Windows Server 2012](#)
- [Windows Server 2012 R2](#)
- [Windows Server 2016](#)

Although you set domain and forest functional levels separately, the forest functional level determines the minimum operating mode for every domain in the forest. You cannot set a domain functional level that is lower than the functional level of the forest. For example, if the forest is set to Windows Server 2012 mode, the minimum functional level of any domain in the forest is Windows Server 2012, though domains could operate at a higher level, such as in Windows Server 2016 mode.

The functional level also determines what releases of Windows Server can operate in the domain or forest as domain controllers. The minimum server version required is the same as the functional level. Thus, if the domain functional level is set to Windows Server 2012, all controllers in the domain must be running Windows Server 2012 or later, and if the forest functional level is set to Windows Server 2012, all controllers in each domain in the forest must be running Windows Server 2012 or later.



Finally, the functional level also determines which Active Directory features are available. Today, most domains and forests operate at the Windows Server 2008 or higher domain and forest functional level. These operating levels allow you to take advantage of the many Active Directory enhancements that improve manageability, performance, and supportability, including the following:

- **Active Directory Administrative Center** Provides a task-orientated interface for managing Active Directory. In Server Manager, click Tools, and then click Active Directory Administrative Center.
- **Active Directory module for Windows PowerShell** Provides cmdlets for managing Active Directory when you are working with Windows PowerShell.
- **Active Directory Recycle Bin** Allows administrators to undo the accidental deletion of Active Directory objects in much the same way as they can recover deleted files from the Windows Recycle Bin. For more information, see “Using the Active Directory Recycle Bin” later in Chapter 6.
- **Active Directory Web Services** Introduces a web service interface for Active Directory domains.
- **Authentication Mechanism Assurance** Improves the authentication process by allowing administrators to control resource access based on whether a user logs on by using a certificate-based logon method. Thus, an administrator can specify that a user has one set of access permissions when logged on by using a smart card and a different set of access permissions when not logged on by using a smart card.
- **Managed service accounts** Introduces a special type of domain user account for managed services that reduces service outages and other issues by having Windows manage the account password and related service principal names (SPNs) automatically. For more information, see “Implementing managed accounts” in Chapter 10, “Working with Managed Accounts.”
- **Managed virtual accounts** Introduces a special type of local computer account for managed services that provides the ability to access the network by using a computer identity in a domain environment. For more information, see “Using Virtual Accounts” in Chapter 10.
- **Offline domain join** Allows administrators to preprovision computer accounts in the domain to prepare operating systems for deployment. This allows computers to join a domain without having to contact a domain controller.

When you raise the domain and forest functional level to Windows Server 2012 or higher, you can take advantage of additional enhancements, including:



- **Active Directory–based activation** Allows you to use Active Directory to automatically activate clients running current versions of Windows or Windows Server. Any client connected to the service is activated. Requires volume licensing. Active Directory schema must be updated for at least Windows Server 2012. Keys are set using the Volume Activation server role or the command line.
- **Authentication policies** Allow you to specify access control conditions that restrict the devices that can request Kerberos tickets for user, service, and computer accounts. Applies only to users, computers, managed services accounts, and group managed service accounts. Accounts must be members of the Protected Users group.
- **Claims-based policy controls** Allow access and audit policies to be defined flexibly. Claims policy must be enabled for Default Domain Controllers Policy and your file servers also must run at least Windows Server 2012.
- **Deferred index creation** Allows deferring of index creation within the directory until UpdateSchemaNow is received or the domain controller is rebooted.
- **Enhanced Fine-Grained Password Policy** Allows administrators to use Active Directory Administrative Center to create and manage Password Settings objects (PSOs).
- **Enhanced Recycle Bin** Allows administrators to recover deleted objects by using Active Directory Administrative Center. Domains must have Recycle Bin enabled to use the feature.
- **Group Managed Service Accounts** Allow multiple services to share a single managed service account.
- **Kerberos constrained delegation across domains** Allows managed service accounts to act on behalf of users across domains and forests.
- **Kerberos with Armoring** Improves domain security; allows a domain-joined client and domain controller to communicate over a protected channel. On clients, enable Require FAST policy; on domain controllers, enable Support CBAC And Kerberos Armoring policy.
- **Off-premises domain join** Allows a computer to be domain-joined over the Internet. Domains must be Direct Access–enabled to use the feature.
- **Protected Users security group** Provides additional protections against authentication threats by requiring accounts that are members of this group to use only Kerberos for authentication.
- **Relative ID (RID) soft ceiling and warnings** Adds warnings as global RID space is used up. Adds a soft ceiling of 900 million RIDs used that prevents RIDs from being issued until administrator overrides.



- **Server Manager integration** Allows you to perform all the steps required to deploy local and remote domain controllers.
- **Virtual domain controller cloning** Allows you to safely deploy virtualized replicas of domain controllers. Also helps maintain domain controller state.

You can raise domain and forest functional levels using the Active Directory Administration tools and lower functional levels using Windows PowerShell. The options for raising or lowering the functional level depend on whether Active Directory Recycle Bin has been enabled. If Active Directory Recycle Bin has been enabled, the range for raising or lowering functional levels is between Windows Server 2008 R2 mode and Windows Server 2016 mode. If Active Directory Recycle Bin has not been enabled, the range for raising or lowering functional levels is between Windows Server 2008 mode and Windows Server 2016 mode.

## Tools for Working with Active Directory

You can manage Active Directory using graphical administration tools, command-line tools, and Windows PowerShell cmdlets.


# Active Directory Administration Tools

The main Active Directory administration tools are provided as snap-ins for the Microsoft Management Console (MMC). These tools include:

- **Active Directory Administrative Center** For performing management tasks.
- **Active Directory Domains And Trusts** For working with domains, domain trees, and domain forests.
- **Active Directory Module For Windows PowerShell** For managing Active Directory when you are working with Windows PowerShell.
- **Active Directory Sites And Services** For managing sites and subnets.
- **Active Directory Users And Computers** For managing users, groups, computers, and organizational units.
- **ADSI Edit** Opens and edits the Active Directory Services Interface for domain, schema, and configuration containers.
- **DNS Server** Manages properties of Domain Name System (DNS) servers, zones, and resource records.
- **Group Policy Management** For managing the way Group Policy is used in the organization. It provides access to Resultant Set of Policy (RSOP) for modeling and logging.

**REAL WORLD** Windows Firewall can affect remote administration with some MMC snap-ins. If Windows Firewall is enabled on a remote computer and you receive an error message stating that you don't have appropriate rights, the network path isn't found, or access is denied, you might need to configure an exception on the remote computer for incoming TCP port 445. To resolve this problem, you can enable the Windows Firewall: Allow Inbound Remote Administration Exception policy setting within Computer Configuration\Policies\Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile. Alternatively, enter the following at a command prompt on the remote computer: **netsh firewall set portopening tcp 445 smb enable** .

You can access the Active Directory administration tools from the Tools menu in Server Manager, or add them to any updateable MMC. If you're using another computer with access to a Windows Server domain, the tools won't be available until you install them. One technique for installing these tools is to use the Add Roles And Features Wizard to add the Remote Server Administration Tools feature for AD DS. You'll then be able to select the individual administration tools to install.



Other support tools for configuring, managing, and troubleshooting Active Directory include:

- **Active Directory Administration Tool (Ldp.exe)** Performs Lightweight Directory Access Protocol (LDAP) operations on Active Directory
- **Directory Services Access Control Lists Utility (Dscls.exe)** Manages access control lists (ACLs) for objects in Active Directory
- **Replication Diagnostics Tool (Repadmin.exe)** Manages and monitors replication by using the command line
- **Windows Domain Manager (Netdom.exe)** Allows domain and trust relationships management from the command line

Repadmin and Netdom aren't the only command-line tools available. Other tools for managing Active Directory from the command line include:

- **Adprep** Allows you to manually prepare a Windows forest or domain for installation of Windows domain controllers (DCs). To prepare a forest or a domain, use **adprep /forestprep** and **adprep /domainprep**, respectively. If you plan to install any read-only domain controllers, you should also run **adprep /rodcprep** for the forest.

**REAL WORLD** Windows Server 2016 automatically prepares forests and domains for you. However, you must use an account with appropriate permissions. For forest and RODC prep to succeed, you need to use an administrator account that is a member of Enterprise Admins, Schema Admins, or Domain Admins in the forest root domain. For domain prep to succeed, you need to use an account that is a member of the Domain Admins group in an applicable domain.



You can run Adprep on any domain-joined server running a 64-bit version of Windows Server 2008 R2 or later. The server needs network connectivity to the schema master for the forest and the infrastructure master of the domain where you want to add the domain controller.

- **Dsadd** Adds computers, contacts, groups, organizational units, and users to Active Directory. Enter **dsadd objectname / ?** at a command prompt to display Help information about using the command, such as **dsadd computer /? .**
- **Dsget** Displays properties of computers, contacts, groups, organizational units, users, sites, subnets, and servers registered in Active Directory. Enter **dsget objectname / ?** at a command prompt to display Help information about using the command, such as **dsget subnet /? .**
- **Dsmmod** Modifies properties of computers, contacts, groups, organizational units, users, and servers that exist in Active Directory. Enter **dsmmod objectname / ?** at a command prompt to display Help information about using the command, such as **dsmmod server /? .**
- **Dsmove** Moves a single object to a new location within a single domain or renames the object without moving it. Enter **dsmove /?** at a command prompt to display Help information about using the command.
- **Dsquery** Uses search criteria to find computers, contacts, groups, organizational units, users, sites, subnets, and servers in Active Directory. Enter **dsquery /?** at a command prompt to display Help information about using the command.
- **Dsrm** Removes objects from Active Directory. Enter **dsrm /?** at a command prompt to display Help information about using the command.
- **Ntdsutil** Allows the user to view site, domain, and server information; manage operations masters; and perform database maintenance of Active Directory. Enter **ntdsutil /?** at a command prompt to display Help information about using the command.

Whereas Adprep is located in the \support\adprep folder on the Windows Server 2016 installation media, the other tools become available when you install the Remote Server Management Tools for AD DS.

## Installing Domain Controllers

To establish a server as a domain controller, you add and configure Active Directory Domain Services (AD DS). Later, if you don't want the server to handle controller tasks, you can demote the server by removing AD DS.





Every Active Directory installation requires DNS. When you install AD DS, you can include DNS server installation, if it is needed. To add the first domain controller that runs Windows Server 2016 to an existing Active Directory infrastructure, the Active Directory Installation Wizard automatically runs Adprep.exe as needed for the forest and domain. Preparing the forest and domain includes updating the Active Directory schema as needed, creating new objects and containers as needed, and modifying security descriptors and access control lists as needed. For forest prep, the account you use must be a member of the Schema Admins group, the Enterprise Admins group, and the Domain Admins group of the domain that hosts the schema master, which is, by default, the forest root domain. For domain prep, you use an account that can log on to the infrastructure master and is a member of the Domain Admins group.



# Setting up a Domain Controller

Active Directory can be deployed in a number of different configurations. You can:

- [Install a new Active Directory forest](#)
- [Add a new child domain or domain tree in an existing forest](#)
- [Add a domain controller to an existing domain](#)

To install a new Active Directory forest, you must be logged on as a local administrator for the computer and the computer should be in a stand-alone configuration. To install a new child domain or new domain tree in an existing forest, you must be logged on as a member of the Enterprise Admins group and the computer should be a member of a domain in that forest. To install an additional domain controller in an existing domain, you must be logged on as a member of the Domain Admins group and the computer should be a member of that domain.

Before you install AD DS to set up a domain controller, you should also ensure the server's network settings are configured appropriately. Domain controllers must be able to communicate with other servers on the network.

**REAL WORLD** If you are adding a domain controller to an existing domain, the new domain controller will get its directory data via replication from another controller in the domain. An alternative technique for installing domain controllers is to use backup media. To install a domain controller from backup media, create a backup of the system state data of a domain controller, and then restore it on a different server. When you create a domain controller from backup media, you eliminate the need to replicate the entire directory database over the network to the new domain controller. This can really save the day when you have bandwidth limitations or the directory database is very large.

Once you've ensured the basic requirements are met, you can install a domain controller by following these steps:

1. [In Server Manager, the local server is added automatically for management. If you want to install AD DS on another server, you need to add the server for management by using the Add Servers option.](#)

**TIP** Using Server Manager for remote management requires the configuration discussed in "Managing Your Servers Remotely" in Chapter 2 and a minimum set of permissions. Typically, you must have Domain Admin or other explicit permissions to add a server and remotely manage it.

2. In Server Manager, click Manage, and then click Add Roles And Features. This starts the Add Roles And Features Wizard. If the wizard displays the Before You Begin page, read the Welcome message, and then click Next.

The screenshot shows the 'Before You Begin' page of the Add Roles and Features Wizard. The left sidebar contains a list of steps: 'Before You Begin' (selected), 'Installation Type', 'Server Selection', 'Server Roles', 'Features', 'Confirmation', and 'Results'. The main content area contains the following text:

This wizard helps you install roles, role services, or features. You determine which roles, role services, or features to install based on the computing needs of your organization, such as sharing documents, or hosting a website.

To remove roles, role services, or features:  
Start the Remove Roles and Features Wizard

Before you continue, verify that the following tasks have been completed:

- The Administrator account has a strong password
- Network settings, such as static IP addresses, are configured
- The most current security updates from Windows Update are installed

If you must verify that any of the preceding prerequisites have been completed, close the wizard, complete the steps, and then run the wizard again.

To continue, click Next.

3. On the Select Installation Type page, select Role-Based Or Feature-Based Installation, and then click Next.

The screenshot shows the 'Select installation type' page of the Add Roles and Features Wizard. The left sidebar contains a list of steps: 'Before You Begin', 'Installation Type' (selected), 'Server Selection', 'Server Roles', 'Features', 'Confirmation', and 'Results'. The main content area contains the following text:

Select the installation type. You can install roles and features on a running physical computer or virtual machine, or on an offline virtual hard disk (VHD).

☒ **Role-based or feature-based installation**  
Configure a single server by adding roles, role services, and features.

☐ **Remote Desktop Services installation**  
Install required role services for Virtual Desktop Infrastructure (VDI) to create a virtual machine-based or session-based desktop deployment.

4. On the Select Destination Server page, the server pool shows servers you've added for management. Click the server you are configuring, and then click Next.

The screenshot shows the 'Select destination server' page of the Add Roles and Features Wizard. The left sidebar contains a list of steps: 'Before You Begin', 'Installation Type', 'Server Selection' (selected), 'Server Roles', 'Features', 'Confirmation', and 'Results'. The main content area contains the following text:

Select a server or a virtual hard disk on which to install roles and features.

☒ **Select a server from the server pool**

☐ **Select a virtual hard disk**

Server Pool

Filter:


| Name                      | IP Address | Operating System              |
|---------------------------|------------|-------------------------------|
| CorpServer17.imaginedl... | 10.0.0.172 | Microsoft Windows Server 2016 |

5. On the Select Server Roles page, select Active Directory Domain Services, and then click Next three times.





6. On the Confirm Installation Selections page, confirm that the Active Directory Domain Services is the role that is being installed. As the server will need to be restarted to complete the second part of the installation, you may want to select the Restart The Destination Server checkbox. Then when prompted to confirm, click Yes to allow the necessary automatic restart.
7. Click Install. This starts the installation.




8. When the initial installation task completes, you need to click Promote This Server To A Domain Controller to start the Active Directory Domain Services Configuration Wizard. If you closed the Add Roles And Features Wizard window, you need to click the Notifications icon, and then click Promote This Server To A Domain Controller.



The configuration wizard takes you through the steps needed to install Active Directory, which might include automatically preparing the directory schema in the forest and domain for Windows Server 2016. You need to specify whether this is a domain controller for a new domain or an additional domain controller for an existing domain.



The image shows the 'Deployment Configuration' window in Windows Server. The left sidebar has a menu with 'Deployment Configuration' selected, followed by 'Domain Controller Options', 'Additional Options', 'Paths', 'Review Options', 'Prerequisites Check', 'Installation', and 'Results'. The main area is titled 'Select the deployment operation' and contains three radio buttons: 'Add a domain controller to an existing domain' (selected), 'Add a new domain to an existing forest', and 'Add a new forest'. Below this is a section 'Specify the domain information for this operation' with a 'Domain:' label and a text box containing 'imaginedlands.local', followed by a 'Select...' button. At the bottom, there is a 'Supply the credentials to perform this operation' section with a text box containing 'IMAGINEDLANDS\williams' and a 'Change...' button. The top right corner shows 'TARGET SERVER' as 'CorpServer17.imaginedlands.local'.

You'll also need to specify whether DNS Server should be installed along with Active Directory and provide additional details regarding whether the server should be a global catalog. Also, every domain controller has a special password that is used when you need to log on as part of Active Directory recovery and you'll need to set the server's password during the configuration.

The image shows the 'Domain Controller Options' window in Windows Server. The left sidebar has a menu with 'Deployment Configuration' and 'Domain Controller Options' selected, followed by 'DNS Options', 'Additional Options', 'Paths', 'Review Options', 'Prerequisites Check', 'Installation', and 'Results'. The main area is titled 'Specify domain controller capabilities and site information' and contains three checkboxes: 'Domain Name System (DNS) server' (checked), 'Global Catalog (GC)' (checked), and 'Read only domain controller (RODC)' (unchecked). Below this is a 'Site name:' label and a dropdown menu showing 'Chicago-First-Site'. At the bottom, there is a section 'Type the Directory Services Restore Mode (DSRM) password' with 'Password:' and 'Confirm password:' labels, each followed by a masked password field (displayed as eight dots). The top right corner shows 'TARGET SERVER' as 'CorpServer17.imaginedlands.local'.

To verify that a domain controller is installed correctly, you should do the following:

- Check the Directory Service event log for errors
- Ensure that the SYSVOL folder is accessible to clients
- Verify that name resolution is working through DNS
- Verify replication of changes to Active Directory

If the installation fails, note any errors and take appropriate corrective action before restarting the configuration procedure. Typical installation errors will relate to permissions, such as those required for preparing the forest or domain for first use of Windows Server. Here, log off and then log back on with an account that has the appropriate permissions.

## Demoting a Domain Controller

Demoting a domain controller removes Active Directory Domain Services from the server. Typically, you demote a domain controller only when other domain controllers are available in the domain to manage directory services. Sometimes, you

also may need to demote the last domain controller in a domain and in this way also remove the domain itself from the Active Directory forest.


Before you demote a domain controller, you should shift any key responsibilities to other domain controllers. This means moving the global catalog off the server and transferring any operations master roles, if necessary. You must also remove any application directory partitions that are on the server. Keep in mind that in order to remove a domain controller from a domain by demoting it, you must be logged on as a member of the Domain Admins group.

You can demote a domain controller by following these steps:


1. In Server Manager, click Manage, and then click Remove Roles And Features. This starts the Remove Roles And Features Wizard. If the wizard displays the Before You Begin page, read the Welcome message, and then click Next.
2. On the Select Destination Server page, the server pool shows servers you've added for management. Click the server you are configuring, and then click Next.



3. On the Remove Server Roles page, clear the check box for Active Directory Domain Services to specify that this is the role you want to remove.



4. A new dialog box opens. Here, you might want to clear the Remove Management Tools check box to ensure the AD DS management tools aren't uninstalled and then click Continue. Otherwise, click Remove Features. Click Next twice.



- When the Validation Results dialog box is displayed, click the Demote This Domain Controller link.



- On the Credentials page, note your current logon account. If necessary, provide alternate credentials with permissions to remove the domain controller. Click Next.

**Credentials**

TARGET SERVER  
CorpServer17.imaginedlands.local

Credentials

Warnings

New Administrator Passw...

Review Options

Demotion

Results

Supply the credentials to perform this operation

IMAGINEDLANDS\williams (Current user)

☐ Force the removal of this domain controller

Change...

7. If the Warnings page is displayed, note the warnings, select Proceed With Removal, and then click Next.

**Warnings**

TARGET SERVER  
CorpServer17.imaginedlands.local

Credentials

Warnings

New Administrator Passw...

Review Options

Demotion

Results

The domain controller currently hosts the following role(s):

- Domain Name System (DNS) Server
- Global Catalog

The roles hosted by the domain controller are required for Active Directory Domain Services functionality. If you proceed, some Active Directory Domain Services operations may be impacted.

☒ Proceed with removal

8. Enter and then confirm a new password for the server's local Administrator account. The passwords must match. Click Next.

**New Administrator Password**

TARGET SERVER  
CorpServer17.imaginedlands.local

Credentials

Warnings

New Administrator Passw...

Review Options

Demotion

Results

Password:

Confirm password:

9. When you are ready to proceed, click Demote. A restart of the server is required to complete the removal, and will be performed automatically as part of the demotion process.

**Review Options**

Demotion

Results

Remove Active Directory Domain Services from this computer.

When the process is complete, this server will be a member of the domain "imaginedlands.local".

These settings can be exported to a Windows PowerShell script to automate additional installations

View script

More about removal options

< Previous

Next >

Demote

Cancel



**CAUTION** Demoting a server gracefully transfers any roles held by the server.

However, if previous attempts to demote the domain controller have failed, you can repeat this procedure and then select the Force The Removal Of This Domain Controller check box as part of the removal process. Here, the flexible single master operations (FSMO) roles of the domain controller might be left in an invalid state until they are reassigned by an administrator. The domain data also might be left in an inconsistent state.




# Managing Users, Groups and Computers

Active Directory Users And Computers is one of the primary administration tools you use to manage Active Directory. With this utility, you can handle all user, group, and computer-related tasks, and manage organizational units.



# Using Active Directory Users and Computers

You can start Active Directory Users And Computers by selecting its related option on the Tools menu in Server Manager. You can also add Active Directory Users And Computers as a snap-in to any console that can be updated. By default, Active Directory Users And Computers works with the domain to which your computer is currently connected. You can access computer and user objects in this domain through the console tree.



**NOTE** If you can't find a domain controller, or if the domain you want to work with isn't shown, you might need to connect to a domain controller in the current domain or to a domain controller in a different domain. Other high-level tasks you might want to perform with Active Directory Users And Computers are viewing advanced options or searching for objects.

When you access a domain in Active Directory Users And Computers, the following standard set of folders is displayed:

- **Builtin** The list of built-in user accounts and groups.
- **Computers** The default container for computer accounts.
- **Domain Controllers** The default container for domain controllers.
- **ForeignSecurityPrincipals** Information about objects from a trusted external domain. Normally, these objects are created when an object from an external domain is added to a group in the current domain.
- **Managed Service Accounts** The default container for managed service accounts.
- **Microsoft Exchange Security Groups** The default container for groups used by Microsoft Exchange Server. This folder is listed only if Exchange Server is running in the environment.
- **Saved Queries** Contains saved search criteria so that you can quickly perform previously run Active Directory searches.




## • Users The default container for users.

Active Directory Users And Computers has advanced options that aren't displayed by default. To access these options and display the following additional folders, click View, and then select Advanced Features:

- **LostAndFound** Contains objects that have been orphaned. You can delete or recover them.
- **NTDS Quotas** Contains directory service quota data.
- **Program Data** Contains stored Active Directory data for Microsoft applications.
- **System** Contains built-in system settings.
- **TPM Devices** Lists devices with Trusted Platform Module (TPM) owner information stored in Active Directory.

You can also add folders for organizational units. By default, you are connected to the local domain and to the first domain controller that responds to your request. You can work with any domain in the forest provided that you have the proper access permissions. To do this, you simply connect to the domain by following these steps:


1. In the console tree, right-click Active Directory Users And Computers, and then click Change Domain.
2. The Change Domain dialog box displays the current (or default) domain. Enter a new domain name, or click Browse, select a domain in the Browse For Domain dialog box, and then click OK.
3. If you always want to use this domain when working with Active Directory Users And Computers, select the Save This Domain Setting For The Current Console check box, and then click OK. Otherwise, just click OK.



If you start Active Directory Users And Computers and no objects are available, it might be because you are not connected to a domain, or a domain controller could not be located. You need to connect to a domain controller to access user, group, and computer objects. To connect to a domain controller, follow these steps:



1. In the console tree, right-click Active Directory Users And Computers, and then click Change Domain Controller. The current domain and domain controller you're working with in the Change Directory Server dialog box are displayed.



2. The Change To list displays the available controllers in the domain. The default selection is Any Writable Domain Controller. If you select this option, you'll be connected to the domain controller that responds to your request first. Otherwise, choose a specific domain controller to which you want to connect.
3. If you always want to use this domain controller when working with Active Directory Users And Computers, select the Save This Setting For The Current Console check box, and then click OK. Otherwise, just click OK.

**NOTE** The Change Directory Server dialog box also shows you the site associated with domain controllers in addition to the domain controller type, version, and status. If the domain controller type is listed as GC, the domain controller is also hosting a global catalog.


You might also want to connect to a specific domain controller for troubleshooting. For example, if you suspect that replication isn't working properly, you might want to inspect the objects on a specific controller. After you're connected, you can look for discrepancies in recently updated objects.

## Searching the Directory


Active Directory Users And Computers has a built-in search feature you can use to find accounts, shared resources, and other directory objects. You can easily search the current domain, a specific domain, or the entire directory.

You search for directory objects by following these steps:

1. In the console tree, right-click the current domain or a specific container that you want to search, and then click Find.



2. In the Find list, choose the type of search you want. The options include the following:
  - **Users, Contacts, And Groups** Search for user and group accounts, and also contacts listed in the directory service.
  - **Computers** Search for computer accounts by type, name, and owner.
  - **Printers** Search for printers by name, model, and features.
  - **Shared Folders** Search for shared folders by name or keyword.
  - **Organizational Units** Search for organizational units by name.
  - **Custom Search** Perform an advanced search or LDAP query.
  - **Common Queries** Search quickly for account names, account descriptions, disabled accounts, nonexpiring passwords, and days since the last logon.
3. Using the In list, select the location you want to search. If you chose a container to search in step 2, such as Computers, this container is selected by default. To search all objects in the directory, click Entire Directory.
4. Enter your search parameters, and then click Find Now. Any matching entries are displayed in the search results. Double-click an object to view or modify its property settings. Right-click the object to display a shortcut menu of options for managing the object.



**NOTE** The search type determines which text boxes and tabs are available in the Find dialog box. In most cases, you'll simply want to enter the name of the object you're looking for in the Name box, but other search options are available. For example, with printers, you can search for a color printer, a printer that can print on both sides of the paper, a printer that can staple, and more.


## Using Administrative Center and Windows Powershell

Active Directory Administrative Center provides a task-orientated interface for managing Active Directory. To start this tool, select the related option from the Tools menu in Server Manager. You can use this tool to perform many common tasks, including the following ones:


- Connect to one or more domains.
- Create and manage user accounts, groups, and organizational units.
- Create and manage password settings objects.
- Create and manage authentication policy objects.
- Perform global searches of Active Directory.
- Raise forest and domain functional levels.
- Recover deleted objects from the Active Directory Recycle Bin.

Administrative Center is installed by default on Windows Server 2016 and is available on client computers when you install the Remote Server Administration Tools (RSAT).

This tool uses Windows PowerShell to perform administration tasks and relies on the Microsoft .NET Framework. Both of these features must be installed and properly configured for you to use Administrative Center.




In Active Directory Administrative Center, the local domain is opened for management by default. If you want to work with a different domain, click **Manage**, and then click **Add Navigation Nodes**. In the **Add Navigation Nodes** dialog box, select the domain you want to work with, and then click **OK**. Afterward, you can select the domain by clicking it in the left pane.



By default, you are connected to the first domain controller that responds to your request. For troubleshooting replication, you might want to connect to a specific domain controller. After you're connected, you can inspect the objects on that controller and look for discrepancies in recently updated objects. To connect to a specific domain controller, click the domain node in the left pane, and then click **Change Domain Controller**.




Displayed in the Change Domain Controller dialog box are the current domain and domain controller you're working with. Select a domain controller to use, and then click Change.



Like Active Directory Users And Computers, Administrative Center has built-in search features you can use to find directory objects. The most basic of these is the search filter, which is available when you select a directory container in the left pane.

Using the search filter, you can quickly find container-level objects within a domain or child OUs within a selected OU. When you select a domain node in the left pane, you can use the filter to quickly find top-level organizational units or built-in containers that start with the letters or words you enter as the filter. For example, you could select the domain node in the left pane and then enter **sa** in the Filter box to find top-level organizational units that begin with the letters “sa,” such as Sales . As such, a search doesn't include child OUs or subcontainers; it wouldn't include the SalesVT or SalesCA organizational units that were child OUs of Sales.





When you select a specific container, you can search within that container by using the same filtering technique. When you select the Global Search node, you can search the names of all container-level objects and also users, groups, computers, and so on for the currently selected container node.

With global searches, you can change the associated container node by clicking Scope and then selecting a node to use. Select Global Catalog Search as the node to search nonstandard objects, such as attribute schemas, display specifiers, intersite transports, or class schemas.

**REAL WORLD** Technically, the filter is based on the start string of any name part of an object. For groups, this means the group name and group Security Accounts Manager (SAM) account name are included. For users, this means the first name, last name, full name, user principal name (UPN), and group SAM account name are included.

In addition to relying on Windows PowerShell to perform administration tasks, Administrative Center makes use of the web services provided by Active Directory Web Services (ADWS). When you perform administrative tasks, those tasks are performed using Windows PowerShell and the web gateway provided by ADWS. At least one domain controller in each Active Directory domain you want to manage must have ADWS installed and have the related services running. Connections are made over TCP port 9389 by default, and firewall policies must enable an exception on this port for ADWS.



You can also work directly with Active Directory by using the Active Directory module for Windows PowerShell. The module is available and automatically imported when you are working with a computer that has the Active Directory administration tools installed.

**NOTE** If the Active Directory module isn't imported automatically, you can import it by entering **Import-Module ActiveDirectory** . After the module is imported, you can use it with the currently running instance of Windows

PowerShell. The next time you start Windows PowerShell, you need to import the module again if you want to use its features.

At the PowerShell prompt, you can list all available cmdlets by entering **getcommand** . Use Get-Help to get more information about how cmdlets are used. If you enter **get-help \*-\*** , you get a list of all cmdlets that includes a synopsis of the purpose of each cmdlet. To get Help documentation about a specific cmdlet, enter **get-help** followed by the cmdlet name. such as **get-helpadd-computer** . Several dozen Active Directory cmdlets are available, and you can get a list of the ones you'll use the most by entering **get-help \*-ad\*** at the PowerShell prompt.

## Working with Organizational Units

As discussed in “Using Organizational Units” in Chapter 4, organizational units are logical containers for directory objects that you can use to improve manageability, set Group Policy with a limited scope, and more. In this section, you'll learn how to create and manage organizational units.






## Creating Organizational Units

You usually create organizational units to mirror your organization's business or functional structure. You might also want to create units for administrative reasons, such as if you want to delegate rights to users or administrators. You can create organizational units as subgroups of a domain or as child units within an existing organizational unit.

To create an organizational unit, follow these steps:

1. In Active Directory Users And Computers or Active Directory Administrative Center, right-click the domain node or existing organizational unit folder in which you want to add an organizational unit. Click New on the shortcut menu, and then click Organizational Unit.
2. Enter the name of the organizational unit, and then click OK.
3. You can now move accounts and shared resources to the organizational unit. See the "Moving Computer Accounts" section of Chapter 11 for an example.




## Viewing and Editing Organizational Unit Properties

As with users, groups and other Active Directory objects, organizational units have properties you can set, including a description, address information and details regarding the manager of the resource. You can view and edit organizational unit properties by following these steps:

1. Open Active Directory Users And Computers or Active Directory Administrative Center.
2. Right-click the organizational unit you want to work with, and then click Properties. This displays a Properties dialog box that lets you view and edit settings.




The image shows a Windows-style dialog box titled "Technology Properties". It has three tabs: "General", "Managed By", and "COM+". The "General" tab is selected. Below the tabs is a folder icon and the text "Technology". The main area contains several input fields: "Description:" with a single-line text box, "Street:" with a multi-line text box and a vertical scrollbar, "City:" with a single-line text box, "State/province:" with a single-line text box, "Zip/Postal Code:" with a single-line text box, and "Country/region:" with a dropdown menu. The dialog box has a question mark icon and a close (X) button in the top right corner.

## Renaming and Deleting Organizational Units


As you work with organizational units, you may need to rename or delete them from time to time. To rename an organizational unit, right-click the organizational unit folder you want to work with in Active Directory Users And Computers and then select Rename. Next, type a new name for the organizational unit, and then press Enter.

Organizational Units are protected from accidental deletion by default and this setting prevents deletion without clearing the protection setting. Thus, you typically must edit an organizational unit's properties prior to actually being able to remove it. This makes removing organizational units a two-part process:


1. In Active Directory Users And Computers, ensure advanced features are displayed by clicking View and making sure the Advanced Features option is selected.



2. Right-click the organizational unit you want to delete and then select Properties.
3. On the Object tab, clear the Protect Object... checkbox and then click OK.



4. You can now delete the organizational unit by right-clicking it and then selecting Delete.
5. When prompted, confirm the action by clicking Yes.




## Moving Organizational Units

You can move organizational units to different locations within a domain at any time. In Active Directory Users And Computers, simply select the organizational unit and then drag it to the desired location.

Using Active Directory Users And Computers, you can also follow these steps to move organizational units:

1. Right-click the organizational unit folder you want to move, and then click Move.
2. In the Move dialog box, expand the domain, select the container to which you want to move the organizational unit, and then click OK.



## Chapter 6. Maintaining Active Directory

Although basic management of Active Directory focuses on establishing domain controllers and managing organizational units, your work with Active Directory domains typically is much bigger, especially if you have administrator responsibilities that include ensuring directory services remain operational. As part of operations support, you may need to:

- Configure physical structure
- Maintain domain controllers
- Manage or recover deleted objects
- Troubleshoot issues

Active Directory doesn't require a great deal of support once you've set up domain controllers but it does require proper configuration of physical structures—sites and



subnets—for communications. It also requires periodic monitoring and maintenance to ensure proper operations.






# Working with Sites

Sites are groups of IP subnets that are connected by reliable, high-speed links. Most of the time, all subnets on the same local network are part of the same site. Networks with multiple sites are connected via site links. Site links are logical, transitive connections between two or more sites. Each site link has a replication schedule, a replication interval, a link cost, and a replication transport.

When you install the first domain controller in a domain, the installation process creates a default site and a default site link. The default site is named Default-First-Site-Name, and the default site link is called DEFAULTIPSITELINK. You can rename the default site and site link as necessary. You must create subsequent sites and site links manually using Active Directory Sites And Services or Windows PowerShell.



Configuring a site is a multipart process that includes:

1. Creating the site
2. Creating one or more subnets and associating them with the site
3. Associating a domain controller with the site
4. Linking the site to other sites by using site links and, if necessary, creating site link bridges

I discuss these tasks in the sections that follow.


## Creating Sites

Any administrator who is a member of Domain Admins or Enterprise Admins can create sites. You can create a site by following these steps:

1. In Active Directory Sites And Services, right-click the Sites container in the console root, and then click New Site.



2. In the New Object - Site dialog box, enter a name for the site, such as **ChicagoFirst-Site** . Site names cannot contain spaces or any special characters other than a dash.



3. Click the site link that you will use to connect this site to other sites. If the site link you want to use doesn't exist, select the default site link and change the site link settings later.
4. Click OK. If a prompt is displayed detailing the steps you must complete to finish site configuration, click OK again.

To complete site configuration, complete the remaining configuration tasks:

- Add subnets to the site.
- Add domain controllers to the site.
- Add site links to the site.


**TIP** You can rename a site at any time. In Active Directory Sites And Services, right-click the site, and then select Rename. Type the new name for the site, and then press Enter.

## Creating Subnets

Each site you define must have associated subnets that detail the network segments that belong to the site. Any computer with an IP address on a network segment associated with a site is considered to be located in the site. Although a single site can have multiple subnets associated with it, a subnet can be associated with only one site.

To create a subnet and associate it with a site, follow these steps:


1. In Active Directory Sites And Services, right-click the Subnets container in the console tree, and then click New Subnet. This displays the New Object - Subnet dialog box.



The 'New Object - Subnet' dialog box is shown. It has three main sections. The first section, 'Prefix:', contains a text box with '192.168.35.0/24'. The second section, 'Prefix name in Active Directory Domain Services:', also contains a text box with '192.168.35.0/24'. The third section, 'Select a site object for this prefix.', features a list box with three items: 'Miami-First-Site' (selected), 'NY-First-Site', and 'NY-Second-Site'. At the bottom are three buttons: 'OK', 'Cancel', and 'Help'.

2. In the Prefix box, enter the IPv4 or IPv6 network address prefix by using the network prefix notation. In network prefix notation, you enter the network ID and then a forward slash, and then you specify which bits are used for the network ID. For example, if the network ID is 192.168.35.0 and the first 24 bits identify the network ID, you enter **192.168.35.0/24** as the network prefix notation.
3. Select the site with which the subnet should be associated, and then click OK.

You can change the site association for a subnet at any time. In Active Directory Sites And Services, double-click the subnet in the Subnets folder, and then on the General tab, use the Site list to change the site association.



The '192.168.35.0/24 Properties' dialog box is shown. It has five tabs: 'General' (selected), 'Location', 'Object', 'Security', and 'Attribute Editor'. The 'General' tab contains a description icon, a 'Description:' text box, a 'Site:' dropdown menu with 'Miami-First-Site' selected, and a 'Prefix:' text box with '192.168.35.0/24'.



# Associating Domain Controllers with Sites

Every site should have at least one domain controller associated with it. By adding a second domain controller to a site, you provide fault tolerance and redundancy. If at least one domain controller in the site is also a global catalog server, you can ensure that directory searches and authentication traffic are isolated to the site.

You can add domain controllers to sites automatically or manually. When you associate subnets with a site, any new domain controllers you install are placed in the site automatically if the domain controller’s IP address is within the valid range of IP addresses for the subnet. Existing domain controllers are not automatically associated with sites, however. You must manually associate any existing domain controllers with a new site by moving the domain controller object into the site.

Before you can move a domain controller from one site to another, you must determine in which site the domain controller is currently located. A quick way to do that is to enter the following command at a command prompt:

```
dsquery server -s DomainControllerName | dsget server -site
```


Here, DomainControllerName is the fully qualified domain name of the domain controller, such as:

```
dsquery server -s server241.tvpress.com | dsget server -site
```

The output of this command is the name of the site in which the designated domain controller is located.


To move a domain controller from one site to another site, follow these steps:

- 1. In Active Directory Sites And Services, any domain controllers associated with a site are listed in the site’s Servers node. Select the site that the domain controller is currently associated with.





2. Right-click the domain controller, and then click Move. In the Move Server dialog box, click the site that should contain the server, and then click OK.




**NOTE** Don't move a domain controller to a site if the domain controller is not on a subnet associated with the site. If you change subnet and site associations, you need to move domain controllers in the affected subnets to the appropriate site containers.

# Configuring Site Links

Because site links are used over wide area network links, bandwidth availability and usage are important considerations when configuring site links. By default, site links are scheduled to replicate data 24 hours a day, 7 days a week at an interval of at least 180 minutes. If you know a link has bandwidth limitations, you might need to alter the schedule to allow user traffic to have priority during peak usage times.

When you have multiple links between sites, you need to consider the relative priority of each link. You assign priority based on the availability and reliability of the connection. The default link cost is set to 100. If there are multiple possible routes to a site, the route with the lowest site link cost is used first. Therefore, the most reliable paths with the most bandwidth between sites should be configured in most cases to have the lowest site link cost.



The screenshot shows the 'Active Directory Sites and Services' console window. The left pane displays a tree view with 'Sites' expanded, showing 'Inter-Site Transports' and 'Subnets'. The right pane displays a table of site links.

| Name                | Type      | Description | Cost | Replication Interval |
|---------------------|-----------|-------------|------|----------------------|
| DCMiamiLink         | Site Link |             | 100  | 180                  |
| ChicagoLALink       | Site Link |             | 100  | 180                  |
| ChicagoNewYorkLink  | Site Link |             | 100  | 180                  |
| LAPortlandLink      | Site Link |             | 100  | 180                  |
| PortlandSeattleLink | Site Link |             | 100  | 180                  |
| SeattleChicagoLink  | Site Link |             | 100  | 180                  |


You can configure site links to use either remote procedure call (RPC) over IP or Simple Mail Transfer Protocol (SMTP) as the transport protocol. With IP as the transport, domain controllers establish an RPC over IP connection with a single replication partner at a time and replicate Active Directory changes synchronously. Because RPC over IP is synchronous, both replication partners must be available at the time the connection is established. You should use RPC over IP when there are reliable, dedicated connections between sites.

With SMTP as the transport, domain controllers convert all replication traffic to email messages that are sent between the sites asynchronously. Because SMTP replication is asynchronous, both replication partners do not have to be available at the time the connection is established, and replication transactions can be stored until a destination server is available. You should use SMTP when links are unreliable or not always available.

**NOTE** If you plan to use SMTP, you must set up a certificate authority (CA). Certificates from the CA are used to digitally sign and encrypt the SMTP messages sent between the sites. With IP, CAs are not required by default.


You can create a site link between two or more sites by following these steps:

1. In Active Directory Sites And Services, expand the Sites container, and then expand the Inter-Site Transports container.
2. Right-click the container for the transport protocol you want to use (either IP or SMTP), and then click New Site Link.
3. In the New Object - Site Link dialog box, enter a name for the site link, such as **ChicagoSeattleLink**. Site link names cannot contain spaces or special characters other than a dash.
4. In the Sites Not In This Site Link list, click the first site that should be included in the link, and then click Add to add the site to the Sites In This Site Link list. Repeat this process for each site you want to add to the link. You must include at least two sites. Click OK.



You can change the sites associated with a site link at any time by following these steps:

1. In Active Directory Sites And Services, right-click the site link in the details pane, and then click Properties.
  2. In the Properties dialog box, the General tab is selected by default. In the Sites Not In This Site Link list, click the first site that should be included in the link, and then click Add to add the site to the Sites In This Site Link list. Repeat this process for each site you want to add to the link.
  3. In the Sites In This Site Link list, click the first site that should not be included in the link, and then click Remove to add the site to the Sites Not In This Site Link list. Repeat this process for each site you want to remove from the link.
- Click  
OK.



When you finish creating the site link, you should configure the link's properties. This allows you to specify the link cost, replication schedule, and replication interval. To configure site link properties, follow these steps:

1. In Active Directory Sites And Services, right-click the site link in the details pane, and then click Properties.
2. In the Properties dialog box, the General tab is selected by default. In the Cost box, set the relative cost of the link. The default cost is 100.

Cost:

Replicate every  minutes

3. In the Replicate Every box, set the replication interval. The default interval is 180 minutes.
4. The default replication schedule is 24 hours a day, 7 days a week. To set a different schedule, click Change Schedule, set the replication schedule in the Schedule For dialog box, and then click OK.

Schedule for ChicagoSeattleLink

12 • 2 • 4 • 6 • 8 • 10 • 12 • 2 • 4 • 6 • 8 • 10 • 12

All

|           | 12 | 2 | 4 | 6 | 8 | 10 | 12 | 2 | 4 | 6 | 8 | 10 | 12 |
|-----------|----|---|---|---|---|----|----|---|---|---|---|----|----|
| Sunday    | ■  | ■ | ■ | ■ | ■ | ■  | ■  | ■ | ■ | ■ | ■ | ■  | ■  |
| Monday    | ■  | ■ | ■ | ■ | ■ | ■  | ■  | ■ | ■ | ■ | ■ | ■  | ■  |
| Tuesday   | ■  | ■ | ■ | ■ | ■ | ■  | ■  | ■ | ■ | ■ | ■ | ■  | ■  |
| Wednesday | ■  | ■ | ■ | ■ | ■ | ■  | ■  | ■ | ■ | ■ | ■ | ■  | ■  |
| Thursday  | ■  | ■ | ■ | ■ | ■ | ■  | ■  | ■ | ■ | ■ | ■ | ■  | ■  |
| Friday    | ■  | ■ | ■ | ■ | ■ | ■  | ■  | ■ | ■ | ■ | ■ | ■  | ■  |
| Saturday  | ■  | ■ | ■ | ■ | ■ | ■  | ■  | ■ | ■ | ■ | ■ | ■  | ■  |

○ Replication Not Available  
● Replication Available

OK  
Cancel

Sunday through Saturday from 12:00 AM to 12:00 AM

## Configuring Site Link Bridges

All site links are transitive by default. This means that when more than two sites are linked for replication and use the same transport, site links are bridged automatically, allowing links to be transitive between sites. Because of transitivity, any two domain controllers can make a connection across any consecutive series of links. For example, a domain controller in site A could connect to a domain controller in site C through site B.

The link path that domain controllers choose for connections across sites is largely determined by the site link bridge cost. The site link bridge cost is the sum of all the links included in the bridge; generally, the path with the lowest total site link bridge cost is used.


Knowing the costs of links and link bridges, you can calculate the effects of a network link failure and determine the paths that will be used when a connection is down. For example, a domain controller in site A would normally connect to a domain controller in site C through site B. However, if the connection to site B is down, the two domain controllers would choose an alternate path automatically if one is available, such as going through site D and site E to establish a connection.

Intersite replication topology is optimized for a maximum of three hops by default. In large-site configurations, this can have unintended consequences, such as the same replication traffic going over the same link several times. In this case, you should disable automatic site link bridging and manually configure site link bridges. Otherwise, you typically do not want to disable automatic site link bridging.

Within an Active Directory forest, you can enable or disable site link transitivity on a per-transport protocol basis. This means all site links that use a particular transport either use site link transitivity or don't use it. You can configure transitivity for a transport protocol by following these steps:

1. In Active Directory Sites And Services, expand the Sites container, and then expand the Inter-Site Transports container.
2. Right-click the container for the transport protocol you want to work with (either IP or SMTP), and then click Properties.
3. To enable site link transitivity, select Bridge All Site Links, and then click OK. When site link transitivity is enabled, any site link bridges you've created for a particular transport protocol are ignored.


4. To disable site link transitivity, clear the Bridge All Site Links check box, and then click OK. When site link transitivity is disabled, you must configure site link bridges for the affected protocol.



After you've disabled transitive links, you can manually create a site link bridge between two or more sites by following these steps:

1. In Active Directory Sites And Services, expand the Sites container, and then expand the Inter-Site Transports container.
2. Right-click the container for the transport protocol you want to work with (either IP or SMTP), and then click New Site Link Bridge.
3. In the New Object—Site Link Bridge dialog box, enter a name for the site link bridge. Bridge names cannot contain spaces or special characters other than a dash.
4. In the Site Links Not In This Site Link Bridge list, select a site link that should be included in the bridge, and then click Add to add the site link to the Site Links In This Site Link Bridge list. Repeat this process for each site link you want to add to the bridge. A bridge must include at least two site links. Click OK.





You can change the site links associated with a site link bridge at any time by following these steps:

1. In Active Directory Sites And Services, select the container for the transport protocol you want to work with, right-click the site link bridge you want to work with, and then click Properties.
2. In the Properties dialog box, the General tab is selected by default. In the Site Links Not In This Site Link Bridge list, click the first site link that should be included in the bridge, and then click Add to add the site link to the Site Links In This Site Link Bridge list. Repeat this process for each site link you want to add to the bridge.
3. In the Site Links In This Site Link Bridge list, click the first site link that should not be included in the bridge, and then click Remove to add the site link to the Site Links Not In This Site Link Bridge list. Repeat this process for each site link you want to remove from the bridge. Click OK.

## Maintaining Domain Controllers

Like driving a car, working with domain controllers is reasonably straight forward, as long as you keep your eye on the road. It's when you look away or get careless that





things go wrong. To avoid potholes, you need to monitor and maintain domain controllers to ensure operations and keep things running smoothly.

As part of your routine work with domain controllers, you may need to:

- [Work with domainwide and forestwide roles.](#)
- [Configure functional levels and global catalogs](#)
- [Configure group membership caching.](#)


These tasks and other related tasks are discussed in the sections that follow.



# Configuring the Functional Level


You can raise the domain level functionality by following these steps:

1. Open Active Directory Domains And Trusts. In the console tree, right-click the domain you want to work with, and then click Raise Domain Functional Level. The current domain name and functional level are displayed in the Raise Domain Functional Level dialog box.
2. To change the domain functionality, select the new domain functional level from the list provided, and then click Raise.
3. Click OK. The new domain functional level is replicated to each domain controller in the domain. This operation can take some time in a large organization.




You can raise the forest level functionality by following these steps:

1. Open Active Directory Domains And Trusts. In the console tree, right-click the Active Directory Domains And Trusts node, and then click Raise Forest Functional Level. The current forest name and functional level are displayed in the Raise Forest Functional Level dialog box.
2. To change the forest functionality, select the new forest functional level by using the list provided, and then click Raise.
3. Click OK. The new forest functional level is replicated to each domain controller in each domain in the forest. This operation can take some time in a large organization.



Another way to raise domain or forest functional level is to use Active Directory Administrative Center. This tool is available as an option on the Tools menu in Server Manager. Follow these steps to raise the domain functional level:

1. In Active Directory Administrative Center, the local domain is opened for management by default. If you want to work with a different domain, click Manage, and then click Add Navigation Nodes. In the Add Navigation Nodes dialog box, select the domain you want to work with, and then click OK.



2. Next, select the domain you want to work with by clicking it in the left pane. In the Tasks pane, click Raise The Domain Functional Level. The current domain name and functional level are displayed in the Raise Domain Functional Level dialog box.
3. To change the domain functionality, select the new domain functional level by using the list provided.

- 
4. Click OK. The new domain functional level is replicated to each domain controller in the domain. This operation can take some time in a large organization.


Follow these steps to raise the forest functional level:

1. In Active Directory Administrative Center, select the domain you want to work with by clicking it in the left pane. In the Tasks pane, click Raise The Forest Functional Level. The current forest name and functional level are displayed in the Raise Forest Functional Level dialog box.
2. To change the forest functionality, select the new forest functional level by using the list provided, and then click Raise.
3. Click OK. The new forest functional level is replicated to each domain controller in each domain in the forest. This operation can take some time in a large organization.

If you want to lower the domain or forest functional level, you'll need to use Windows PowerShell. You use `Get-ADDomain` and `Get-ADForest` to get information about domains and forest, including the fully qualified name of a domain or forest. You use `Set-ADDomainMode` to configure the domain functional level and the `SetADForestMode` to configure the forest functional level.

## Working with Domainwide Roles

Operations master roles are discussed in the “Understanding Operations Master Roles” section in Chapter 4. The fastest way to determine the current FSMO for all roles is to enter **netdom query fsmo** at a command prompt.



```
Select Command Prompt
Microsoft Windows
(c) 2016 Microsoft Corporation. All rights reserved.

C:\Users\williams.IMAGINEDLANDS>netdom query fsmo
Schema master           CorpServer14.imaginedlands.local
Domain naming master    CorpServer30.imaginedlands.local
PDC                     CorpServer28.imaginedlands.local
RID pool manager        CorpServer28.imaginedlands.local
Infrastructure master    CorpServer43.imaginedlands.local
The command completed successfully.

C:\Users\williams.IMAGINEDLANDS>
```

The domain-level roles are relative ID (RID) master, primary domain controller (PDC) emulator master, and infrastructure master. The forest-level roles are domain naming master and schema master.

You can view the current operations master roles for the domainwide roles by following these steps:

1. In Active Directory Users And Computers, right-click Active Directory Users And Computers in the console tree. On the shortcut menu, point to All Tasks, and then click Operations Masters. This opens the Operations Masters dialog box.
2. The Operations Masters dialog box has three tabs. The RID tab shows the location of the current RID master. The PDC tab shows the location of the current PDC emulator master. The Infrastructure tab shows the location of the current infrastructure master.




You can transfer the current operations master roles by following these steps:

1. Open Active Directory Users And Computers. In the console tree, right-click Active Directory Users And Computers, and then click Change Domain Controller.

2. In the Change Directory Server dialog box, click This Domain Controller Or AD LDS Instance, select the domain controller to which you want to transfer an operations master role, and then click OK.



3. In the console tree, right-click Active Directory Users And Computers. On the shortcut menu, point to All Tasks, and then click Operations Masters.
4. In the Operations Masters dialog box, click the RID, PDC, or Infrastructure tab as appropriate for the type of role you want to transfer.
5. Click Change to transfer the role to the previously selected domain controller, and then click OK.




## Working with Forestwide Roles

You can use Active Directory Domains And Trusts to view or change the location of the domain naming master in the domain forest. In Active Directory Domains And Trusts, the root level of the control tree shows the currently selected domain.


To transfer the domain naming master role, follow these steps:

1. Start Active Directory Domains And Trusts. In the console tree, right-click Active Directory Domains And Trusts and then click Change Active Directory Domain Controller.

2. In the Change Directory Server dialog box, select the This Domain Controller Or AD LDS Instance option, select the domain controller to which you want to transfer the domain naming master role, and then click OK.



3. In the console tree, right-click Active Directory Domains And Trusts, and then click Operations Master. This opens the Operations Master dialog box.
4. The Domain Naming Operations Master box displays the current domain naming master. Click Change to transfer this role to the previously selected domain controller.
5. Click Close.




You use Active Directory Schemato view or change the schema master's location.


Enter **regsvr32 schmmgmt.dll** at an administrator command prompt to register Active Directory Schema. You can then transfer the schema master role by following these steps:

1. Type **MMC.exe** in the Search box and then press Enter. Add the Active Directory Schema snap-in to the console. To do this, select Add/Remove Snap-in on the File menu. Next, in the Add Or Remove Snap-Ins dialog box, select Active Directory Schema under Available Snap-Ins, click Add and then click OK.





2. The Active Directory Schema snap-in is added to the console. In the console tree, right-click Active Directory Schema, and then click Change Active Directory Domain Controller.



3. Select Any Writable Domain Controller to let Active Directory select the new schema master, or select This Domain Controller Or AD LDS Instance and then select the new schema master.
4. Click OK. In the console tree, right-click Active Directory Schema, and then click Operations Master.
5. Click Change in the Change Schema Master dialog box. Click OK, and then click Close.





Change Schema Master

✕

The schema master manages modifications to the schema. Only one server in the enterprise performs this role.

Current schema master (online):

CorpServer28.imaginedlands.local

To transfer the schema master role to the targeted schema FSMO holder below, click Change.

CorpServer17.imaginedlands.local

Change



# Transferring Roles by Using the Command Line

Another way to transfer roles is to use Netdom to list current FSMO role holders and then Ntdsutil.exe to transfer roles. Ntdsutil is a command-line tool for managing Active Directory. Follow these steps to transfer roles at the command line:

1. Get a list of the current FSMO role holders by entering **netdom query fsmo** at a command prompt.
2. It is recommended (but not required) that you log on to the console of the server you want to assign as the new operations master. You can log on to the console locally or use Remote Desktop Connection.
3. Open a command prompt. One way to do this is by pressing the Windows key, typing **cmd.exe** , and then pressing Enter.
4. At the command prompt, enter **ntdsutil** . This starts the Directory Services Management Tool.
5. At the ntdsutil prompt, enter **roles** . This puts the utility in Operations Master Maintenance mode.
6. At the fsmo maintenance prompt, enter **connections** . At the server connections prompt, enter **connect to server** followed by the fully qualified domain name of the domain controller to which you want to assign the FSMO role, such as:

connect to server engdc01.technology.imaginedlands.com

7. After you've established a successful connection, enter **quit** to exit the server connections prompt. At the fsmo maintenance prompt, type **transfer** , type the identifier for the role to transfer, and then press Enter. The identifiers are as follows:
  - **pdc** For the PDC emulator role
  - **rid master** For the RID master role
  - **infrastructure master** For the infrastructure master role
  - **schema master** For the schema master role
  - **domain naming master** For the domain naming master role
8. Enter **quit** at the fsmo maintenance prompt, and then enter **quit** at the ntdsutil prompt.

## Configuring Global Catalogs


Global catalogs have an important role on the network. This role is discussed in the “Understanding the Directory Structure” section in Chapter 4. You configure additional



global catalogs by enabling domain controllers to host the global catalog. In addition, if you have two or more global catalogs within a site, you might want a domain controller to stop hosting the global catalog. You do this by disabling the global catalog on the domain controller.

You enable or disable a global catalog by following these steps:

- 1. In Active Directory Sites And Services, expand the site you want to work with in the console tree.
- 2. Expand the Servers folder for the site, and then select the server you want to configure to host the global catalog.



- 3. In the details pane, right-click NTDS Settings, and then click Properties.
- 4. To enable the server to host the global catalog, select the Global Catalog check box on the General tab.
- 5. To disable the global catalog, clear the Global Catalog check box on the General tab.





**CAUTION** Don't enable or disable global catalogs without proper planning and analysis of the impact on the network. In a large enterprise environment, designating a domain controller as a global catalog can cause data related to thousands of Active Directory objects to be replicated across the network.




# Configuring Universal Group Membership Caching

Universal membership caching eliminates the dependency on the availability of a global catalog server during logons. When you enable this feature on a domain controller operating at the Windows Server 2003 or higher functional level, any domain controller can resolve logon requests locally without having to go through the global catalog server. As discussed in the “Universal Group Membership Caching” section in Chapter 4, this has advantages and disadvantages.

You can enable or disable universal group membership caching by following these steps:

1. In Active Directory Sites And Services, select the site you want to work with.



2. In the details pane, right-click NTDS Site Settings, and then click Properties.
3. To enable universal group membership caching, select the Enable Universal Group Membership Caching check box on the Site Settings tab. Then, in the Refresh Cache From list, choose a site from which to cache universal group memberships. The selected site must have a working global catalog server.
4. To disable universal group membership caching, clear the Enable Universal Group Membership Caching check box on the Site Settings tab.
5. Click OK.

NTDS Site Settings Properties ? X

Site Settings Object Security Attribute Editor

 NTDS Site Settings

Description:

Inter-Site Topology Generator

Server:

Site:

Universal Group Membership Caching

☒ Enable Universal Group Membership Caching

Refresh cache from:



## Using the Active Directory Recycle Bin

When your Active Directory forest is operating in the Windows Server 2008 R2 or higher mode, you can use the Active Directory Recycle Bin. The Active Directory Recycle Bin adds an easy-to-use recovery feature for Active Directory objects. When you enable this feature, all link-valued and nonlink-valued attributes of a deleted object are preserved, allowing you to restore the object to the same state it was in before it was deleted. You can also recover objects from the recycle bin without having to initiate an authoritative restore.





## Preparing Schema for the Recycle Bin

Before you can make the Recycle Bin available, you must update Active Directory schema with the required Recycle Bin attributes. You do this by installing a domain controller running Windows Server 2008 R2 or higher in a domain in the forest. When you do this, the schema is updated, and then every object in the forest is also updated with the Recycle Bin attributes. This process is irreversible after it is started.

After you prepare Active Directory, you need to upgrade all domain controllers in your Active Directory forest to Windows Server 2008 R2 or higher and then raise the domain and forest functional levels to the Windows Server 2008 R2 level or higher. Optionally, you can update Active Directory schema in your forests and domains for Windows Server 2012 or later to enable the enhanced Recycle Bin.

After these operations, you can enable and access the Recycle Bin. When Recycle Bin has been enabled, it cannot be disabled. Now when an Active Directory object is deleted, the object is put in a state referred to as logically deleted and moved to the Deleted Objects container. Also, its distinguished name is altered. A deleted object remains in the Deleted Objects container for the period of time set in the deleted object lifetime value, which is 180 days by default.

**REAL WORLD** The `msDS-deletedObjectLifetime` attribute replaces the `tombstoneLifetime` attribute. However, when `msDS-deletedObjectLifetime` is set to `$null`, the lifetime value comes from the `tombstoneLifetime`. If the `tombstoneLifetime` is also set to `$null`, the default value is 180 days.

## Recovering Deleted Objects

If you elect not to use the Recycle Bin, you can still recover deleted objects from the Deleted Objects container by using an authoritative restore and other techniques I'll discuss in this section. The procedure has not changed from early releases of Windows Server. What has changed, however, is that the objects are restored to their previous state with all link-valued and non-link-valued attributes preserved. To perform an authoritative restore, the domain controller must be in Directory Services Restore Mode.

Rather than using an authoritative restore and taking a domain controller offline, you can recover deleted objects by using the Active Directory cmdlets for Windows PowerShell. If you updated the Active Directory schema in your forests and domains for Windows Server 2012 or later, you also can enable the enhanced Recycle Bin, which allows you to recover deleted objects by using Active Directory Administrative Center.







Keep in mind that Active Directory blocks access to an object for a short while after the object is deleted. During this time, Active Directory processes the object's link-value table to maintain referential integrity on the linked attribute's values. Active Directory then permits access to the deleted object.

## Using Windows Powershell for Basic and Advanced Recovery

The Active Directory cmdlets for Windows PowerShell allow you to recover deleted objects by using scripts or by entering commands at a Windows PowerShell prompt. You use Get-ADObject to retrieve the object or objects you want to restore, pass that object or objects to Restore-ADObject, and then Restore-ADObject restores the object or objects to the directory database.

To use the Active Directory cmdlets for recovery, you need to open an administrator Windows PowerShell prompt by right-clicking the Windows PowerShell entry on the menu, and then clicking Run As Administrator. The basic syntax for recovering an object is as follows

```
Get-ADObject -Filter { ObjectId } -IncludeDeletedObjects | Restore-ADObject
```

ObjectId is a filter value that identifies the object you want to restore. For example, you could restore a deleted user account by display name or SAM account name, as shown in these examples.

```
Get-ADObject -Filter {DisplayName -eq "Rich Tuppy"}  
-IncludeDeletedObjects | Restore-ADObject  
Get-ADObject -Filter {SamAccountName -eq "richt"} -  
IncludeDeletedObjects | Restore-ADObject
```

Note that nested objects must be recovered from the highest level of the deleted hierarchy to a live parent container. For example, if you accidentally deleted an OU and all its related accounts, you need to restore the OU before you can restore the related accounts.

The basic syntax for restoring container objects such as an OU is as follows.

```
Get-ADObject -ldapFilter:"(msDS-LastKnownRDN= ContainerID )" -  
IncludeDeletedObjects | Restore-ADObject
```

ContainerID is a filter value that identifies the container object you want to restore. For example, you could restore the Corporate Services OU, as shown in this example.

```
Get-ADObject -ldapFilter:"(msDS-LastKnownRDN=Corporate_Services)" -  
IncludeDeletedObjects | Restore-ADObject
```



If the OU contains accounts you also want to restore, you can now restore the accounts by using the technique discussed previously, or you can restore all accounts at the same time. The basic syntax requires that you establish a search base and associate the accounts with their last known parent, as shown here.

```
Get-ADObject -SearchBase "CN=Deleted Objects, ForestRootDN " -Filter {lastKnownParent -eq  
" ContainerCN,ForestRootDN "} -IncludeDeletedObjects | Restore-  
ADObject
```

ForestRootDN is the distinguished name of the forest root domain, such as DC=tvpress,DC=Com, and ContainerCN is the common name of the container, such as OU=Corporate\_Services or CN=Users. The following example restores all the accounts that were in the Corporate Services OU when it was deleted.

```
Get-ADObject -SearchBase "CN=Deleted Objects,DC=tvpress,DC=com" -Filter  
{lastKnownParent -eq "OU=Corporate_Services,DC=tvpress,DC=com"} -  
IncludeDeletedObjects | Restore-ADObject
```


## Using the Enhanced Recycle Bin for Recovery

The enhanced Recycle Bin makes recovering deleted objects as easy as pointing and clicking or tapping and holding. After you update the Active Directory schema in your forests and domains for Windows Server 2012 or later, you enable the enhanced Recycle Bin for use by following these steps:

1. In Active Directory Administrative Center, select the domain you want to work with by clicking it in the left pane. In the Tasks pane, click Enable Recycle Bin, and then click OK in the confirmation dialog box.




2. Active Directory will begin replicating the change to all domain controllers in the forest. After the change is replicated, the enhanced Recycle Bin will be available for use. Click OK when prompted. If you then click Refresh in Active Directory Administrative Center, you'll notice that a Deleted Object container is now available for domains that use the enhanced Recycle Bin.



Keep in mind that the enhanced Recycle Bin is a forestwide option. When you enable this option in one domain of a forest, Active Directory replicates the change to all domain controllers in all domains of the forest.


With the enhanced Recycle Bin enabled, you can recover deleted objects with ease. In Active Directory Administrative Center, domains using the enhanced Recycle Bin will have a Deleted Object container. In this container, a list of deleted objects is displayed. As discussed previously, deleted objects remain in this container for the deleted object lifetime value, which is 180 days by default.



Each deleted object is listed by name, when it was deleted, the last known parent, and the type. When you select a deleted object by clicking it, you can use the options in the Tasks pane to work with it. The Restore option restores the object to its original container. For example, if the object was deleted from the Users container, it is restored to this container.

The Restore To option restores the object to an alternate container within its original domain or to a different domain within the current forest. Specify the alternate container in the Restore To dialog box. For example, if the object was deleted from the Users

container in the tech.tvpress.com domain, you could restore it to the Devs OU in the eng.tvpress.com domain.



## Troubleshooting Active Directory


To ensure proper operations of Active Directory, you need to perform periodic monitoring and maintenance. In your monitoring and maintenance efforts, you'll find that some tools are instrumental to your success. In this section, I'll introduce these tools and some general maintenance and troubleshooting tasks.

## Using Adsi Edit

When you are diagnosing problems and troubleshooting, the Active Directory administration tool you should use is Active Directory Service Interfaces Editor (ADSI Edit). You can use ADSI Edit to manage the definitions of object classes and their attributes in the schema and to work with other naming contexts, including the default naming context, the Configuration naming context, and the RootDSE naming context. If you want to create custom attributes for users or groups, use ADSI Edit, which you can start by using the related option on the Tools menu in Server Manager.


You can use the ADSI Edit snap-in to connect to a naming context by following these steps:

1. Right-click the ADSI Edit node in the console tree, and then select **Connect To...**. This displays the Connection Settings dialog box.



2. In the Connection Settings dialog box, the Select A Well Known Naming Context list is enabled by default. Choose the naming context you want to work with.
3. When you click OK, you are connected to any available domain controller in your logon domain. To connect to a different domain or server, choose Select Or Type A Domain Or Server, and then choose the server or domain you want to work with along with an optional port number for the connection, such as FileServer25.tvpress.com:389. Port 389 is the default port for LDAP.

After you select a naming context, domain, and server, you are connected to and can work with the naming context. When you connect to multiple naming contexts, you have separate nodes for managing each context. For troubleshooting, you can also connect to the same naming context on different servers in the same domain. By comparing the values associated with properties on one server with those on another, you can identify a replication problem.




## Examining Intersite Topology


The intersite topology generator (ISTG) on a site is responsible for generating the intersite replication topology. When calculating the replication topology, the ISTG can use considerable processing power, especially as the size of the network grows. Because of this, you should closely monitor the ISTGs in each site to ensure that they are not overloaded.

You can determine which domain controller is the ISTG by following these steps:

1. In Active Directory Sites And Services, expand the Sites container, and then select the site for the ISTG you want to locate in the console tree.



2. In the details pane, double-click NTDS Site Settings. In the NTDS Site Settings dialog box, the current ISTG is listed in the Inter-Site Topology Generator panel.




Although replication normally occurs because of change notifications and doesn't require a schedule, some scenarios, such as domain controller outages, might mean



that some changes aren't replicated within a site. In these situations, the replication schedule is used and a domain controller that had an outage would receive the changes on the next cycle after it became available according to the replication schedule you configured.

While you are working with a site, you can review the intrasite replication interval and the periods of the day when non-standard replication is performed by clicking Change Schedule. The default schedule is for daily once an hour replication. You can modify this to replicate on specify days and times and with specific frequency, such as twice an hour or four times an hour.



Replication between sites normally is performed by bridgehead servers . A bridgehead server is a domain controller designated by the ISTG to perform intersite replication. The ISTG configures a bridgehead server for each Active Directory partition that needs to be replicated and maintains a separate replication topology for each type of partition. Although a single bridgehead server can be responsible for replicating multiple directory partitions, the replication topology for each partition is maintained separately.

Domain controllers that operate as bridgehead servers have an additional workload that increases with the number and frequency of replication changes. As you should do with the ISTG, you should periodically monitor designated bridgehead servers to ensure that they do not become overloaded. You can list the bridgehead servers in a site by entering the following command at a command prompt:

```
repadmin /bridgeheads site:SiteName
```

Here, SiteName is the name of the site, such as:





repadmin /bridgeheads site:SacramentoSite


If current bridgehead servers become overloaded, or if you have domain controllers you would prefer to be bridgehead servers, you can designate preferred bridgehead servers to use. Once you designate a preferred bridgehead server for a site, the ISTG uses the preferred bridgehead server for intersite replication. If the preferred bridgehead server goes offline or is unable to replicate for any reason, intersite replication stops until the server is again available or you change the preferred bridgehead server configuration.

When you designate preferred bridgeheads, you should always configure multiple preferred bridgehead servers in each site. The ISTG will then choose one of the servers you’ve designated as the preferred bridgehead server. If this server fails, the ISTG would then choose another server from the list of preferred bridgehead servers.

You must configure a bridgehead server for each partition that needs to be replicated. This means you must configure at least one domain controller with a replica of each directory partition as a bridgehead server. If you don’t do this, replication of the partition will fail and the ISTG will log an event in the Directory Services event log detailing the failure.


You can configure a domain controller as a preferred bridgehead server by following these steps:

- 1. In Active Directory Sites And Services, domain controllers associated with a site are listed in the site’s Servers node. Right-click the server you want to designate as a preferred bridgehead, and then click Properties.



- 2. In the Properties dialog box, select the intersite transport protocol for which the server should be a preferred bridgehead in the Transports Available For InterSite Data Transfer list and then click Add. Repeat as necessary to specify both IP and SMTP. Click OK.





When you've designated preferred bridgehead servers, you can recover from replication failure in several ways. You can remove the failed servers as preferred bridgehead servers and then specify different preferred bridgehead servers, or you can remove all servers as preferred bridgehead servers and then allow the ISTG to select the bridgehead servers that should be used. To stop a server from being a preferred bridgehead for a particular transport protocol, follow these steps:

1. In Active Directory Sites And Services, domain controllers associated with a site are listed in the site's Servers node. Right-click the server you want to stop using as a preferred bridgehead, and then click Properties.
2. Select the transport protocol in the This Server Is A Preferred Bridgehead Server For The Following Transports list, and then click Remove. Click OK.

## Diagnosing and Resolving Issues

As part of routine maintenance, you need to monitor domain controllers, global catalog servers, bridgehead servers, and site links. If you suspect problems with Active Directory, you should look at replication in most cases as the starting point for your diagnostics and troubleshooting. By configuring monitoring of Active Directory intrasite and intersite replication, you can diagnose and resolve most replication problems. Keep in mind, though, that Active Directory replication has several service dependencies, including LDAP, DNS, Kerberos version 5 authentication, and RPC.



These important services must be functioning properly to allow directory updates to be replicated. During replication, Active Directory relies on various TCP and UDP ports being open between domain controllers. By default, the ports used are as follows:

- LDAP uses TCP and UDP on port 389 for standard traffic, and TCP on port 686 for secure traffic.
- Global catalogs use TCP on port 3268. Kerberos version 5 uses TCP and UDP on port 88.
- DNS uses TCP and UDP on port 53.
- SMB over IP uses TCP and UDP on port 445.

Additionally, for replication of files in the System Volume (SYSVOL) shared folders on domain controllers, Active Directory uses either the file replication service (FRS) or the DFS Replication Service. The appropriate replication service must be running and properly configured to replicate the SYSVOL.

As mentioned earlier in the chapter, Active Directory tracks changes by using update sequence numbers (USNs). Any time a change is made to the directory, the domain controller processing the change assigns the change a USN. Each domain controller maintains its own local USNs and increments the value each time a change occurs. The domain controller also assigns the local USN to the object attribute that changed. Each object has a related attribute, `usnChanged`, which is stored with the object and identifies the highest USN that has been assigned to any of the object's attributes.

Each domain controller tracks its local USN and also the local USNs of other domain controllers. During replication, domain controllers compare the USN values received to what is stored. If the current USN value for a particular domain controller is higher than the stored value, changes associated with that domain controller need to be replicated. If the current value for a particular domain controller is the same as the stored value, changes for that domain controller do not need to be replicated.

You can monitor replication from the command line by using Repadmin. With Repadmin, most command-line parameters accept a list of the domain controllers you want to work with, called `DSA_List`. You can specify the values for `DSA_List` as follows:

- **\*** A wildcard character that includes all domain controllers in the organization
- **PartialName** A partial server name followed by the **\*** wildcard character to match the remainder of the server name

- **Site:SiteName** The name of the site for which you want to include domain controllers
- **Gc** All global catalog servers in the organization

Although Repadmin has many parameters and you can use it in many ways, you'll perform certain tasks more than others. To view detailed help for each task enter the following example at the command prompt, **repadmin /?:kcc**. Some of these tasks include:

- Forcing the Knowledge Consistency Checker (KCC) to recalculate the intrasite replication topology for a specified domain controllerby using **repadmin /kcc DSA\_List [/async]**
- Listing a summary of the replication stateby using **repadmin /replsummary [DSA\_List]**
- Listing bridgehead servers that match the DSA\_Listby using **repadmin /bridgeheads [DSA\_List] [/verbose]**
- Listing calls made but not yet answered by the specified server to other serversby using **repadmin /showoutcalls [DSA\_List]**
- Listing computers that have opened sessions with a specified domain controllerby using **repadmin /showctx [DSA\_List]**
- Listing connection objects for the specified domain controllers. Defaults to the local siteby using **repadmin /showconn [DSA\_List]**
- Listing domains trusted by a specified domainby using **repadmin /showtrust [DSA\_List]**
- Listing failed replication events that were detected by the KCCby using **repadmin /failcache [DSA\_List]**
- Listing replication partners for each directory partition on the specified domain controllerby using **repadmin /showrepl [DSA\_List]**
- Listing server certificates loaded on the specified domain controllersby using **repadmin /showcert [DSA\_List]**
- Listing tasks waiting in the replication queueby using **repadmin /queue [DSA\_List]**
- Listing the name of the ISTG for a specified siteby using **repadmin istg [DSA\_List] [/verbose]**
- Listing the time between intersite replications that are using the ISTG Keep Alive time stampby using **repadmin /latency [DSA\_List] [/verbose]**

## Seizing Roles By Using the Command Line

Occasionally, you might find yourself in a situation where you can't gracefully transfer server roles. For example, a domain controller acting as the RID master might have a drive failure that takes down the entire server. If you're unable to get the server back online, you might need to seize the RID master role and assign this role to another domain controller.

**NOTE** Seize a server role only if the domain controller managing the current role is out of service. When the original server comes back online, the server will recognize the role change and accept it.

Don't seize a role without first determining how up to date the domain controller that will take over the role is with respect to the previous role owner. Active Directory tracks replication changes by using update sequence numbers (USNs). Because replication takes time, not all domain controllers will necessarily be up to date. If you compare a domain controller's USN to that of other servers in the domain, you can determine whether the domain controller is the most up to date with respect to changes from the previous role owner. If the domain controller is up to date, you can transfer the role safely. If the domain controller isn't up to date, you can wait for replication to occur and then transfer the role to the domain controller.

Windows Server 2016 includes several tools for working with Active Directory replication, including Repadmin which I introduced in the previous section. You can display the status of the last inbound replication for a domain controller by using Repadmin /Showrepl. The syntax is:

```
repadmin /showrepl DomainControllerName NamingContext
```

Here, DomainControllerName is the fully qualified domain name of the domain controller, and NamingContext is the distinguished name of the domain in which the server is located. In this example, you examine the default partition for Server52 in the Tvpress.com domain:

```
repadmin /showrepl server52.tvpress.com DC=tvpress,dc=com
```

**NOTE** Windows PowerShell parses commands differently from the command prompt. Typically, you can enter commands at the Windows PowerShell prompt in the same way you enter them at the command prompt. In this case, though, Windows PowerShell will misinterpret DC=tvpress,dc=com as two separate parameters. To prevent this, enclose the value in quotation marks: "DC=tvpress,dc=com".

To display the highest sequence number for a specified naming context on each replication partner of a designated domain controller, enter the following at a command prompt:

```
repadmin /showutdvec DomainControllerName NamingContext
```

In this example, you display the highest sequence number for the default configuration partition on Server52 in the Tvpress.com domain:

```
repadmin /showutdvec server52.tvpress.com DC=tvpress,dc=com
```

The output shows the highest USN on replication partners for the default configuration partition:

```
Default-First-Site-Name\SERVER52 @ USN 87164 @ Time 2016-09-30 12:15:14
```

```
Default-First-Site-Name\SERVER47 @ USN 87414 @ Time 2016-09-30 12:14:56
```

If Server52 was the previous role owner, and the domain controller you are examining has an equal or larger USN for Server52, the domain controller is up to date.

However, if Server52 was the previous role owner, and the domain controller you are examining has a lower USN for Server52, the controller is not up to date and you should wait for replication to occur before seizing the role. You can also use Repadmin /Syncall to force the domain controller that is the most up to date with respect to the previous role owner to replicate with all of its replication partners.

In Windows PowerShell, you can use replication management cmdlets to view and troubleshoot Active Directory replication. Related cmdlets include the following:

- **Get-ADReplicationAttributeMetadata** Gets the replication metadata for the attributes of the distinguished name specified
- **Get-ADReplicationFailure** Gets information about replication failure for a specified server, site, domain, or forest, if applicable
- **Get-ADReplicationPartnerMetadata** Gets replication metadata for a specified server, site, domain, or forest
- **Get-ADReplicationQueueOperation** Gets pending operations in a server's replication queue
- **Get-ADReplicationUpToDatenessVectorTable** Gets the highest USN for the specified server, site, domain, or forest
- **Sync-ADObject** Replicates the specified directory object

Use Get-ADReplicationPartnerMetadata to get information about inbound replications for a server with the following syntax:

```
Get-ADReplicationPartnerMetadata -Target Object
```





```
[-Scope Server|Site|Domain|Forest] [-Partition Domain|Schema|Configuration|*]
```

Here, –Target sets the name of the server, site, domain, or forest to work with. Setting the scope is required when you are working with objects other than servers. Setting the partition is required when you want to work with partitions other than the default. In this example, you examine the default partition on CorpServer98:

```
get-adreplicationpartnermetadata -target corpserver98
```

You also can examine all partitions on the server by using the following syntax:

```
get-adreplicationpartnermetadata -target corpserver98 -partition *
```

Like Repadmin /Showutdvec, Get-ADReplicationUpToDatenessVectorTable displays the highest sequence numbers for partitions being replicated and can help you troubleshoot replication issues. Here is the basic syntax:

```
Get-ADReplicationUpToDatenessVectorTable -Target Object  
[-Scope Server|Site|Domain|Forest] [-Partition Domain|Schema|Configuration|*]
```

In this example, you display the highest sequence number for the default partition (the domain configuration partition) on CorpServer98:

```
get-adreplicationuptodatenessvectortable -target corpserver98
```

The output shows the highest USN on replication partners for the default configuration partition:

```
LastReplicationSuccess : 9/30/2016 8:45:57 PM  
Partition              : DC=tvpress,DC=com  
PartitionGuid          : c39cfdbd-e1a1-4c4c-9355-85d7ea05c10a  
Partner               : CN=NTDS Settings,CN=CORPSEVER72,CN=Servers,  
CN=Default-First-Site-Name,CN=Sites,CN=Configuration,DC=tvpress,DC=com  
PartnerInvocationId    : fb32931c-e319-473a-8069-d781f980057b  
Server                : CorpServer98.tvpress.com UsnFilter  
: 82656
```

```
LastReplicationSuccess : 9/30/2016 8:48:44 PM  
Partition              : DC=tvpress,DC=com  
PartitionGuid          : c39cfdbd-e1a1-4c4c-9355-85d7ea05c10a  
Partner               : CN=NTDS Settings,CN=CORPSEVER98,CN=Servers,  
CN=Default-First-Site-Name,CN=Sites,CN=Configuration,DC=tvpress,DC=com  
PartnerInvocationId    : d8bf2da2-b08d-4d36-bc53-1b7f62643437  
Server                : CorpServer98.tvpress.com  
UsnFilter              : 12593
```

You interpret the output much like you would interpret output from Repadmin /Showutdvec. If you suspect a problem, you can use Get-ADReplicationFailure to examine replication issues. Here is the basic syntax:



Get-ADReplicationFailure -Target Object [-Scope Server|Site|Domain|Forest]

Knowing this, you can display information about all replication failures in the Tvpress.com domain by entering the following:

```
get-adreplicationfailure -Target "tvpress.com" -Scope Domain
```

You can display information for a specific site by entering the following:

```
get-adreplicationfailure -Target "NewYork-FirstSite" -Scope Site
```

Or you can display information for a specific server by entering the following:

```
get-adreplicationfailure -Target CorpServer72
```

Follow these steps to seize a server role:

1. Enter **netdom query fsmo** at a command prompt to get a list of the current FSMO role holders.
2. Ensure that the current domain controller with the role you want to seize is permanently offline. If the server can be brought back online, don't perform this procedure unless you intend to completely reinstall this server.
3. It is recommended that you log on to the console of the server you want to assign as the new operations master. You can log on to the console locally or use Remote Desktop Connection.
4. Open a Command Prompt.
5. At the command prompt, enter **ntdsutil** . This starts the Directory Services Management Tool.
6. At the ntdsutil prompt, enter **roles** . This puts the utility in Operations Master Maintenance mode.
7. At the fsmo maintenance prompt, enter **connections** . At the server connections prompt, enter **connect to server** followed by the fully qualified domain name of the domain controller to which you want to assign the FSMO role, such as:

```
connect to server engdc01.technology.imaginedlands.com
```

8. After you've established a successful connection, enter **quit** to exit the server connections prompt. At the fsmo maintenance prompt, type **seize** , type the identifier for the role to seize, and then press Enter. The identifiers are as follows:





- **pdc** For the PDC emulator role **rid master** For the
  - RID master role **infrastructure master** For the
  - infrastructure master role **schema master** For the
  - schema master role
  - **domain naming master** For the domain naming master role
9. Enter **quit** at the fsmo maintenance prompt, and then enter **quit** at the ntdsutil prompt.







## Chapter 7. Accounts: The Essentials

In this chapter, you'll learn how to manage domain accounts. With user accounts, you can enable individual users to log on to the network and access network resources. With group accounts, you manage resources for multiple users. The permissions and privileges you assign to user and group accounts determines which actions users can perform and which computer systems and resources they can access.

Although local system accounts are discussed, they are not the primary focus. For further discussion of configuring local system accounts, see Chapters 7 and 8 in *Windows 10: Essentials for Administration* (Stanek & Associates, 2016). Keep in mind that current versions of Windows and Windows Server add a special type of local account called a Microsoft account. Microsoft accounts can be thought of as synchronized accounts. When you connect a local or domain account to a Microsoft account, the account becomes a connected local or connected domain account.

Although you might be tempted to give users wide access, you need to balance a user's need for access to job-related resources with the organization's need to protect sensitive resources and privileged information. For example, you don't want everyone in the company to have access to payroll data. Consequently, you should be sure that only those who require that information as part of their job have access to it.

### Managing Access to Network Resources

You control access to network resources with the components of the Windows Server security model. The key components you need to know about are those used for authentication and access controls.



# Navigating Authentication Protocols

In Windows and Windows Server, authentication is implemented as a two-part process consisting of an interactive logon and network authentication. When a user logs on to a computer by using a domain account, the interactive logon process authenticates the user's logon credentials, which confirms the user's identity to the local computer and grants access to Active Directory Domain Services (AD DS). Afterward, whenever the user attempts to access network resources, network authentication is used to determine whether the user has permission to do so.

Windows Server 2016 supports multiple network authentication protocols, including Kerberos V5 and NTLM. Active Directory uses Kerberos V5 as the default authentication protocol. NTLM authentication is maintained only for backward compatibility. In Group Policy, you can control how NTLM is used with the security option Network Security: LAN Manager Authentication Level. The default authentication level in most cases is Send NTLMv2 Response Only. With this authentication level, clients use NTLM 2 for authentication and session security if the server supports it. Active Directory can also use client certificates for authentication.

The authentication model used by Windows and Windows Server supports single signon, which works as follows:

1. A user logs on to the domain by using a logon name and password, or by inserting a smart card into a card reader.
2. The interactive logon process authenticates the user's access. With a local account, the credentials are authenticated locally, and the user is granted access to the local computer. With a domain account, the credentials are authenticated in Active Directory, and the user has access to local and network resources.
3. The user can authenticate to any computer in the domain through the network authentication process.

With domain accounts, the network authentication process typically is automatic (through single sign-on). With local accounts, users must provide a user name and password every time they access a network resource. When you need to integrate the on-premises organization into the cloud, you have several options.

To set up a private cloud for users, you can use Active Directory Federation Services (AD FS), which extends single sign-on to trusted resources on the Internet. Using AD FS, organizations can extend their existing Active Directory infrastructure to provide access to trusted Internet resources, which can include third parties and geographically



separated units of the same organization. After you configure federated servers, users at the organization can sign on once to the organization's network and are then automatically logged on to trusted Web applications hosted by partners on the Internet. Federated Web Single Sign-On uses federated authorization for seamless access. In addition to user identity and account information, security tokens used in federated authorization include authorization claims that detail user authorization and specific application entitlement.

When you need to integrate the on-premises organization with Microsoft Online services, you can do so using Azure AD Connect. By connecting your on-premises directories with Azure AD and Office 365, you allow users to use a common identity for accessing both on-premises and cloud resources. At the heart of Azure AD Connect are synchronization services responsible for ensuring identity information is matched between the on-premises organization and the cloud, and health services responsible for providing monitor and diagnostic resources. AD FS is an optional third component of the integration if you need to synchronization identities between multiple forests or customize sign in options.





## Using Access Controls

Active Directory is object-based. Users, computers, groups, shared resources, and many other entities are all defined as objects. Access controls are applied to these objects with security descriptors. Security descriptors do the following:

- List the users and groups that are granted access to objects
- Specify permissions the users and groups have been assigned
- Track events that should be audited for objects
- Define ownership of objects

Individual entries in the security descriptor are referred to as access control entries (ACEs). Active Directory objects can inherit ACEs from their parent objects. This means that permissions for a parent object can be applied to a child object. For example, all members of the Domain Admins group inherit permissions granted to this group.

When working with ACEs, keep the following points in mind:

- ACEs are created with inheritance enabled by default.
- Inheritance takes place immediately after the ACE is created and saved. All
- ACEs contain information specifying whether the permission is inherited or explicitly assigned to the related object.

## Using Claims-Based Access Controls

To the standard access controls, current releases of Windows Server add Kerberos armoring, compound identities, and claims-based access controls. Kerberos armoring improves domain security by allowing domain-joined clients and domain controllers to communicate over secure, encrypted channels. Compound identities incorporate not only the groups that a user is a member of, but also user claims, device claims, and resource properties.

Claims-based access controls can be configured in several ways. The most basic approach is to define conditions that limit access as part of a resource's advanced security permissions. Typically, these conditions add device claims or user claims to the access controls. User claims identify users; device claims identify devices. For example, to access the Human Resources share, you might add a device claim to ensure that the computer being used to access a resource is a member of HR Computers, and add a user claim that ensures the user is a member of the HR Managers group.





Kerberos armoring, compound identities, and claims-based access controls can also work together as part of a new authorization platform that allows dynamic access to resources by configuring central access policies. With central access policies, you define central access rules in Active Directory, and those rules are applied dynamically throughout the enterprise. Central access rules use conditional expressions that require you to determine the resource properties required for the policy, the claim types and security groups required for the policy, and the servers where the policy should be applied.

Before you can define and apply an access rule, you'll likely need to define resource properties and claim types:

- **Resource properties** create property definitions for resources. For example, you might add Department and Country properties to files so that you can dynamically control access by department and country.
- **Claim types** create claim definitions for resources. For example, you might create a user claim to add Department and Country properties to User objects so that you can dynamically control access by department and country.

After you create resource properties and claim types and determine where the policy should be applied, you can create an access rule and then add it to a central access policy. Adding the rule to a policy makes the rule available for dynamic control. You then need to apply the policy across file servers by using Group Policy.

Claims-based policy must be enabled for Default Domain Controllers policy. You do this by enabling and configuring the KDC Support For Claims, Compound Authentication And Kerberos Armoring policy in the Administrative Templates policies for Computer Configuration under System\KDC. The policy must be configured to use a specific mode. Here are the available modes:

- **Supported** Domain controllers support claims, compound identities, and Kerberos armoring. Client computers that don't support Kerberos armoring can be authenticated.
- **Always Provide Claims** This mode is the same as Supported, but domain controllers always return claims for accounts.
- **Fail Unarmored Authentication Requests** Kerberos armoring is mandatory. Client computers that don't support Kerberos armoring cannot be authenticated.

The Kerberos Client Support For Claims, Compound Authentication And Kerberos Armoring policy in the Administrative Templates policies for Computer Configuration under System\Kerberos controls whether the Kerberos client running




on Windows and Windows Server requests claims and compound authentication. The policy must be enabled for compatible Kerberos clients to request claims and compound authentication for Dynamic Access Control and Kerberos armoring.

**REAL WORLD** Claims-based policy should be enabled for all domain controllers in a domain to ensure consistent application. A domain must have at least one domain controller running Windows Server 2012 or later, and file servers must run Windows Server 2012 or later. By default, domain controllers are placed in the Domain Controllers organizational unit (OU), and Default Domain Controllers policy has the highest precedence among Group Policy Objects (GPOs) linked to the Domain Controllers OU. If your organization uses a different approach, you need to ensure the GPO with the highest precedence for the appropriate OU has claims-based policy enabled and configured properly.

## Working with Central Access Policies


Central access policies don't replace traditional access controls. Instead, they are designed to enhance existing access controls by defining very precisely the specific attributes users and devices must have to access resources. The easiest way to manage central access policy is to use Active Directory Administrative Center.




An overview of the policy creation and deployment process follows:

1. Open Active Directory Administrative Center. In the left pane, List View is selected by default. Click Tree View to display the tree view.






2. Expand Dynamic Access Control in the left pane, and then select Claim Types.



3. Use the Claim Types node to create and manage claim types. For example, rightclick the Claim Types node, click New, and then select Claim Type to start creating a new claim type.



4. Use the Resource Properties node to create and manage resource properties. For example, right-click the Resource Properties node, click New, and then select Resource Property to start creating a new resource property.

**Create Resource Property: Protected**

**General**

**Suggested Values**

A resource property describes a characteristic of a resource, such as a file or a folder. It is used to define target resources and permissions when authoring central access rules. It is also used to classify resources.

Display name: \* Protected

Value type: Single-valued Choice

Description:

**NOTE** Resource properties are also added as classification definition properties on file servers.

5. Use the Central Access Rules node to create and manage central access rules. For example, right-click the Central Access Rules node, click New, and then select Central Access Rule to start creating a new access rule.

**Create Central Access Rule: Default Access Rule**

**General**

A central access rule defines the assignment of permissions to resources that qualify the scope of the target resource. It is used to construct a central access policy, which can be then be applied on a resource, such as a folder, after it is published.

Name: \* Default Access Rule

Description:

☒ Protect from accidental deletion

6. Use the Central Access Policies node to create and manage central access policies. For example, right-click the Central Access Policies node, click New, and then select Central Access Policy to start creating a new access policy.

**Create Central Access Policy: Primary Access Policy**

**General**


A central access policy contains several central access rules that determine who has access to resources such as files and folder. The policy can be published and then applied to a resource to control the access to that resource.

Name: \* Primary Access Policy


Description:

To complete the deployment, you need to edit the highest precedence GPO linked to the OU where you put file servers and then enable central access policies. To do this, follow these steps:


1. In Group Policy Management, open the GPO for editing. Browse Computer Configuration policies to Windows Settings\Security Settings\File System.



2. Right-click Central Access Policy, and then click Manage Central Access Policies. This opens the Central Access Policies Configuration dialog box.



3. In the Central Access Policies Configuration dialog box, available policies are listed in the left pane, and currently applied policies are listed in the right pane. To apply a policy, click it in the left pane and then click Add. To remove a policy, click it in the right pane and then click Remove. Click OK.



After the Group Policy changes take effect on your servers, the dynamic controls are available. You can speed the refresh along by entering **gpupdate /force** at an administrator command prompt.

Servers that you want to apply dynamic controls to must have, at a minimum, the File

And Storage Services role with the File Server, Storage Services, and File Server Resource Manager role services. You need the File Server Resource Manager role service and the related tools to apply classification property definitions to folders.



After you enable central access policy, and any time you update your classification property definitions, you need to wait for Global Resource Properties from Active Directory to refresh on your file servers. You can speed this along by opening Windows PowerShell and entering **update-fsrmclassificationpropertydefinition** . Do this on each file server where you want to configure central access policies.

To complete the deployment of central access policies, you need to edit the properties of each folder where you want a central access policy to apply and do the following:

1. Add the appropriate classification definitions on the folder's Classification tab. On the Classification tab, each resource property you created will be listed. Select each property in turn and then set its value as appropriate.
2. Enable the appropriate policy by using advanced security settings for the folder. On the Security tab, click Advanced, and then click the Central Policy tab. Any currently selected or applied policy is listed along with a description that allows you to review the rules of that policy. When you click Change, you can use the selection list provided to select a policy to apply, or you can choose No Central Access Policy to stop using policy. Click OK.

Repeat this process for each top-level or other folder where you want to limit access. Files and folders within the selected folder will inherit the access rule automatically unless you specify otherwise. As an example, if you create an access rule called "HR Managers in the US" and define Department and Country resource definitions, you could edit the HR folder's properties, select the Classification tab, and use the options available to set Department to HR and Country to US. Then you could apply the "HR Managers in the US" policy by using the advanced security settings for the folder.



# Navigating User and Group Accounts

Windows Server provides user accounts and group accounts (of which users can be a member). User accounts are designed for individuals. Group accounts are designed to make the administration of multiple users easier. Although you can log on with user accounts, you can't log on with a group account. Group accounts are usually referred to simply as groups .

**REAL WORLD** Windows Server supports the InetOrgPerson object. Essentially, this object is the same as a user object, and you can use it as such. However, the real purpose for the InetOrgPerson object is to allow for compatibility and transition from third-party X.500 and Lightweight Directory Access Protocol (LDAP) directory services that use this object to represent users. If you are migrating from a third-party directory service and end up with many InetOrgPerson objects, don't worry. These objects are security principals just like user accounts. You can set passwords for InetOrgPerson objects and change the object class if you want to. When you change the object class, the InetOrgPerson object is converted to a user object, and from then on it is listed as the User type in Active Directory Users And Computers.



# Working with User Accounts

Two types of user accounts are defined in Windows Server:

- User accounts defined in Active Directory are called domain user accounts . Through single sign-on, domain user accounts can access resources throughout the domain. You create domain user accounts in Active Directory Users And Computers.
- User accounts defined on a local computer are called local user accounts . Local user accounts have access to the local computer only, and they must authenticate themselves before they can access network resources. You create local user accounts with the Local Users And Groups utility.

**NOTE** In a domain, only member servers and workstations have local user and group accounts. On the initial domain controller for a domain, these accounts are moved from the local Security Account Manager (SAM) database to Active Directory and then become domain accounts.

All user accounts are identified with a logon name. In Windows Server, this logon name has two parts:

- **User name** The text label for the account
- **User domain or workgroup** The workgroup or domain where the user account exists

For the user wrstanek, whose account is created in the tvpress.com domain, the standard logon name is wrstanek@tvpress.com. The legacy logon name is TVPRESS\wrstanek.

When working with Active Directory, you might also need to specify the fully qualified domain name (FQDN) for a user. The FQDN for a user is the combination of the Domain Name System (DNS) domain name, the container or organizational unit that contains the user, and the user name. For example, for the user tvpress.com\users\wrstanek, tvpress.com is the DNS domain name, users is the container or organizational unit location, and wrstanek is the user name.

User accounts can also have passwords and public certificates associated with them. Passwords are authentication strings for an account. Public certificates combine a public and private key to identify a user. You log on with a password interactively. You log on with a public certificate by using a smart card and a smart card reader.

Although Windows Server displays user names to describe privileges and permissions, the key identifiers for accounts are security identifiers (SIDs). SIDs are unique identifiers that are generated when you create accounts. Each account's SID



consists of the domain's security ID prefix and a unique relative ID (RID), which is allocated by the relative ID master.

Windows Server uses these identifiers to track accounts independently from user names. SIDs serve many purposes. The two most important purposes are to allow you to change user names easily and to allow you to delete accounts without worrying that someone might gain access to resources simply by re-creating an account with the same name.

When you change a user name, you tell Windows Server to map a particular SID to a new name. When you delete an account, you tell Windows Server that a particular SID is no longer valid. Afterward, even if you create an account with the same user name, the new account won't have the same privileges and permissions as the previous one.

That's because the new account will have a new SID.





# Working with Group Accounts

In addition to user accounts, Windows Server provides groups. Generally speaking, you use groups to grant permissions to similar types of users and to simplify account administration. If a user is a member of a group that can access a resource, that particular user can access the same resource. Thus, you can give a user access to various work-related resources just by making the user a member of the correct group. Note that although you can log on to a computer with a user account, you can't log on to a computer with a group account.

Because different Active Directory domains might have groups with the same name, groups are often referred to by domain \ groupname , such as TVPRESS\gmarketing for the gmarketing group in the tvpress domain. When you work with Active Directory, you might also need to specify the FQDN for a group. The FQDN for a group is the concatenation of the DNS domain name, the container or organizational unit location, and the group name. For the group tvpress.com\users\gmarketing, tvpress.com is the DNS domain name, users is the container or organizational unit location, and gmarketing is the group name.

**REAL WORLD** Employees in a marketing department probably need access to all marketing-related resources. Instead of granting access to these resources to each individual employee, you could make the users members of a marketing group. That way, they automatically obtain the group's privileges. Later, if a user moves to a different department, you simply remove the user from the group, thus revoking all access permissions. Compared to having to revoke access for each individual resource, this technique is pretty easy, so you'll want to use groups whenever possible.

Windows Server supports three types of groups:

- **Local** Groups that are defined on a local computer. Local groups are used on the local computer only. You create local groups with the Local Users And Groups utility.
- **Security** Groups that can have security descriptors associated with them. You define security groups in domains by using Active Directory Users And Computers.
- **Distribution** Groups that are used as email distribution lists. They can't have security descriptors associated with them. You define distribution groups in domains by using Active Directory Users And Computers.

**NOTE** Most general discussions about groups focus on local groups and security groups rather than on distribution groups. Distribution groups are only for email distribution and are not for assigning or managing access.





In Active Directory, groups can have different scopes—domain local, built-in local, global, and universal. That is, the groups are valid in different areas, as described here:

- **Domain local** Groups primarily used to assign access permissions to resources within a single domain. Domain local groups can include members from any domain in the forest and from trusted domains in other forests. Typically, global and universal groups are members of domain local groups.
- **Built-in local** Groups with a special group scope that have domain local permissions and, for simplicity, are often included in the term domain local groups. The difference between built-in local groups and other groups is that you can't create or delete built-in local groups. You can only modify built-in local groups. References to domain local groups apply to built-in local groups unless otherwise noted.
- **Global** Groups that are used primarily to define sets of users or computers in the same domain that share a similar role, function, or job. Members of global groups can include only accounts and groups from the domain in which they're defined.
- **Universal** Groups that are used primarily to define sets of users or computers that should have wide permissions throughout a domain or forest. Members of universal groups include accounts, global groups, and other universal groups from any domain in the domain tree or forest.

**BEST PRACTICES** Universal groups are very useful in large enterprises where you have multiple domains. If you plan properly, you can use universal groups to simplify system administration. You shouldn't change the members of universal groups frequently. Each time you change the members of a universal group, you need to replicate those changes to all the global catalogs in the domain tree or forest. To reduce changes, assign other groups rather than user accounts to the universal group.

When you work with groups, the group's scope restricts what you can and cannot do:

- Domain local groups can have as members accounts, global groups, and universal groups from any domain; domain local groups from the same domain only. These groups can be put into other domain local groups and assigned permissions only in the same domain, and can be converted to universal scope provided that it doesn't have as its member another group having domain local scope.
  - Global groups can have as members accounts and global groups from the same domain only. These groups can be put into other groups and assigned permissions in any domain, and can be converted to universal scope provided that it's not a member of any other group having global scope.
- 

- Universal groups can have as members accounts from any domain, in addition to global and universal groups from any domain. These groups can be put into other groups and assigned permissions in any domain, and can't be converted to any other group scope.

As with user accounts, Windows Server tracks group accounts with unique SIDs. This means that you can't delete a group account, re-create it, and then expect all the permissions and privileges to remain the same. The new group will have a new SID, and all the permissions and privileges of the old group are lost.

Windows Server creates a security token for each user logon. The security token specifies the user account ID and the SIDs of all the security groups to which the user belongs. The token's size grows as the user is added to additional security groups, which has the following consequences:

- The security token must be passed to the user logon process before logon can be completed. As the number of security group memberships grows, the logon process takes longer.
- To determine access permissions, the security token is sent to every computer that the user accesses. Therefore, the size of the security token has a direct impact on the network traffic load.

**NOTE** Distribution group memberships aren't distributed with security tokens, so distribution group memberships don't affect the token size.

Domain local, global, and universal groups provide many options for configuring groups in the enterprise. Although these group scopes are designed to simplify administration, poor planning can make them your worst administration nightmare. Ideally, you use group scopes to help you create group hierarchies that are similar to your organization's structure and the responsibilities of particular groups of users. The best uses for domain local, global, and universal groups are as follows:

- **Domain local** Groups with domain local scope have the smallest extent of all groups. Use groups with domain local scope to help you manage access to resources such as printers and shared folders.
- **Global** Use groups with global scope to help you manage user and computer accounts in a particular domain. Then you can grant access permissions to a resource by making the group with global scope a member of the group with domain local scope.
- **Universal** Groups with universal scope have the largest extent. Use groups with universal scope to consolidate groups that span domains. Typically, you do this by



adding global groups as members. Then, when you change membership of the global groups, the changes aren't replicated to all global catalogs because the membership of the universal group didn't change.

**TIP** If your organization doesn't have two or more domains, you don't really need to use universal groups. Instead, build your group structure with domain local and global groups. Then, if you ever bring another domain into your domain tree or forest, you can easily extend the group hierarchy to accommodate the integration.

To put this in perspective, consider the following scenario. Suppose that you have branch offices in Seattle, Chicago, and New York. Each office has its own domain, which is part of the same domain tree or forest. These domains are called Seattle, Chicago, and NY, respectively. You want to make it easy for any administrator (from any office) to manage network resources, so you create a group structure that is very similar at each location. Although the company has marketing, IT, and engineering departments, let's focus on the structure of the marketing department. At each office, members of the marketing department need access to a shared printer called MarketingPrinter and a shared data folder called MarketingData. You also want users to be able to share and print documents. For example, Bob in Seattle should be able to print documents so that Ralph in New York can pick them up on his local printer, and Bob should also be able to access the quarterly report in the shared folder at the New York office.

To configure the groups for the marketing departments at the three offices, you'd follow these steps:

1. Start by creating global groups for each marketing group. In the Seattle domain, create a group named **GMarketing** and add the members of the Seattle marketing department to it. In the Chicago domain, create a group named **GMarketing** and add the members of the Chicago marketing department to it. In the NY domain, create a group named **GMarketing** and add the members of the New York marketing department to it.
2. In each location, create domain local groups that grant access to the shared printers and shared folders. Name the printer group **LocalMarketingPrinter**, and name the shared folder group **LocalMarketingData**. The Seattle, Chicago, and NY domains should each have their own local groups.
3. Create a group with universal scope in the domain at any branch office. Name the group **UMarketing**. Add Seattle\GMarketing, Chicago\GMarketing, and NY\GMarketing to this group.



4. Add UMarketing to the LocalMarketingPrinter and LocalMarketingData groups at each office. Marketing users should now be able to share data and printers.

## Granting Privileges, Rights and Permissions

When you set up a user account, you can grant the user specific capabilities. You generally assign these capabilities by making the user a member of one or more groups, thus giving the user the capabilities of these groups. You withdraw capabilities by removing group membership.

In Windows Server, you can assign the following types of user rights as capabilities to an account:

- **Privileges** Grants permissions to perform specific administrative tasks. An example of a privilege is the ability to shut down the system. You assign privileges through group policies, which can be applied to individual computers, organizational units, and domains. Although you can assign privileges to both users and groups, you'll usually want to assign privileges to groups. In this way, users are automatically assigned the appropriate privileges when they become members of a group. Assigning privileges to groups also makes it easier to manage user accounts.
- **Logon rights** Grants logon permissions. You can assign logon rights to both user and group accounts. An example of a logon right is the ability to log on locally. As with privileges, you assign logon rights through group policies, and you'll usually want to assign logon rights to groups rather than to individual users.
- **Built-in capabilities** Assigned to groups and includes the group's automatic capabilities. Built-in capabilities are predefined and unchangeable, but they can be delegated to users with permission to manage objects, organizational units, or other containers. An example of a built-in capability is the ability to create, delete, and manage user accounts. This capability is assigned to administrators and account operators. Thus, if a user is a member of the Administrators group, the user can create, delete, and manage user accounts.
- **Access permissions** Defines the operations that can be performed on network resources. You can assign access permissions to users, computers, and groups. An example of an access permission is the ability to create a file in a directory.

As an administrator, you deal with account capabilities every day. To help track built-in capabilities, refer to the following sections. Keep in mind that although you can't change a group's built-in capabilities, you can change a group's default rights. For




example, an administrator could revoke network access to a computer by removing a group's right to access the computer from the network.

The difference between privileges and logon rights is an important one:

- A privilege is a user right assignment that grants permissions to perform a specific administrative task. You assign privileges through group policies, which can be applied to individual computers, organizational units, and domains. Although you can assign privileges to both users and groups, you'll usually want to assign privileges to groups. In this way, users are automatically assigned the appropriate privileges when they become members of a group. Assigning privileges to groups also makes it easier to manage user accounts.
- A logon right is a user rights assignment that grants logon permissions. You can assign logon rights to both user and group accounts. As with privileges, you assign logon rights through group policies, and you'll usually want to assign logon rights to groups rather than to individual users.

The built-in capabilities that are assigned to groups in Active Directory depend on a computer's configuration. Using the Local Group Policy Editor, you can view the capabilities that have been assigned to each group by expanding Windows Settings\Security Settings under Computer Configuration\Local Policies and then selecting the User Rights Assignment node.



Note that any action that's available to the Everyone group is available to all groups, including the Guests group. This means that although the Guests group doesn't have explicit permission to access the computer from the network, a member of the Guests group can still access the system because the Everyone group has this right.





# Using the Default User Accounts

When you install Windows Server 2016, the operating system installs default users and groups. These accounts are designed to provide the basic setup necessary to grow your network. Three types of default accounts are provided:

- **Built-in** User and group accounts installed with the operating system, applications, and services
- **Predefined** User and group accounts installed with the operating system
- **Implicit** Special groups, also known as special identities , created implicitly when accessing network resources

**NOTE** Although you can modify default users and groups, you can't delete default users and groups created by the operating system because you wouldn't be able to re-create them. The SIDs of the old and new accounts wouldn't match, and the permissions and privileges of these accounts would be lost.







## Working with Built-in User Accounts

Built-in user accounts have special purposes in Windows Server. All Windows Server systems have several built-in user accounts, including the following ones:

- **LocalSystem** A pseudoaccount for running system processes and handling systemlevel tasks. This account is part of the Administrators group on the server and has all user rights on the server. If you configure applications or services to use this account, the related processes have full access to the server system. Many services run under the LocalSystem account. In some cases, these services also have the privilege to interact with the desktop. Services that need alternative privileges or logon rights run under the LocalService or NetworkService account.
- **LocalService** A pseudoaccount with limited privileges. This account grants access to the local system only. The account is part of the Users group on the server and has the same rights as the NetworkService account, except that it is limited to the local computer. Configure applications or services to use this account when related processes don't need to access other servers.
- **NetworkService** A pseudoaccount for running services that need additional privileges and logon rights on a local system and the network. This account is part of the Users group on the server and provides fewer permissions and privileges than the LocalSystem account (but more than the LocalService account). Specifically, processes running under this account can interact throughout a network by using the credentials of the computer account.

When you install add-ons or other applications on a server, other default accounts might be installed.

## Working with Predefined User Accounts

Several predefined user accounts are installed with Windows Server, including Administrator and Guest. With member servers, predefined accounts are local to the individual system they're installed on. Predefined accounts have counterparts in Active Directory. These accounts have domainwide access and are completely separate from the local accounts on individual systems.

Administrator is a predefined account that provides complete access to files, directories, services, and other facilities. In Active Directory, the Administrator account has domainwide access and privileges. Otherwise, the Administrator account generally has access only to the local system. Although files and directories can be protected from the Administrator account temporarily, the Administrator account can take control of these resources at any time by changing the access permissions. By





default, the Administrator account is enabled for use, but you can disable or rename it to enhance security.

**SECURITY ALERT** To prevent unauthorized access to the system or domain, be sure to give the Administrator account an especially secure password. Also, because this is a known Windows account, you might want to rename the account as an extra security precaution. If you rename the original Administrator account, you might also want to create a dummy Administrator account. This dummy account should have no permissions, rights, or privileges, and you should disable it.

You usually won't need to change the basic settings for the Administrator account. However, you might need to change its advanced settings, such as membership in particular groups. By default, the Administrator account for a domain is a member of these groups: Administrators, Domain Admins, Domain Users, Enterprise Admins, Group Policy Creator Owners, and Schema Admins. You'll find more information about these groups in the next section.

**REAL WORLD** In a domain environment, you use the local Administrator account primarily to manage the system when you first install it. This allows you to set up the system without getting locked out. You probably won't use the account after the system has been installed. Instead, you should make your administrators members of the Administrators group. This ensures that you can revoke administrator privileges without having to change the passwords for all the Administrator accounts.

For a system that's part of a workgroup where each individual computer is managed separately, you typically rely on this account any time you need to perform your system administration duties. Here, you probably don't want to set up individual accounts for each person who has administrative access to a system. Instead, use a separate administrator account on each computer.

Another predefined user is the Guest account, which is designed for users who need one-time or occasional access. Although guests have limited system privileges, you should be very careful about using this account. Whenever you use this account, you open the system to potential security problems. The risk is so great that the account is initially disabled when you install Windows Server.

The Guest account is a member of the Domain Guests and Guests groups by default. Note that the Guest account—like all other named accounts—is also a member of the implicit group Everyone. The Everyone group typically has access to files and folders by default. The Everyone group also has a default set of user rights.







**SECURITY ALERT** If you decide to enable the Guest account, be sure to restrict its use and change the password regularly. As with the Administrator account, you might want to rename the account as an added security precaution.





## Using Other Default Accounts

Built-in groups are installed with all Windows Server systems. Use built-in and predefined groups to grant a user the group's privileges and permissions. You do this by making the user a member of the group. For example, you give a user administrative access to the system by making a user a member of the local Administrators group. You give a user administrative access to the domain by making a user a member of the Domain Admins group in Active Directory.

Implicit groups are assigned implicitly during logon and are based on how a user accesses a network resource. For example, if a user accessed a resource through an interactive logon, the user is automatically a member of the implicit group named Interactive. Although you can't view the membership of special identities, you can grant membership in implicit groups to users, groups, and computers.

To reflect the modified role, implicit groups are also referred to as special identities . A special identity is a group whose membership can be set implicitly, such as during logon, or explicitly through security access permissions. As is the case with other default groups, the availability of a specific implicit group depends on the current configuration. Implicit groups are discussed later in this chapter.





# Using the Default Group Accounts

The default group accounts are designed to be versatile. By assigning users to the correct groups, you can make managing your Windows Server 2016 workgroup or domain a lot easier. Unfortunately, with so many groups, understanding the purpose of each isn't easy. To help, let's take a closer look at groups used by administrators and groups that are implicitly created.





## Working with the Admin Groups

An administrator is someone who has wide access to network resources.

Administrators can create accounts, modify user rights, install printers, manage shared resources, and more. The main administrator groups are Administrators, Domain Admins, and Enterprise Admins.

**TIP** The Administrator account and the global groups Domain Admins and Enterprise Admins are members of the Administrators group. The Administrator account is used to access the local computer. Domain Admins membership allows other administrators to access the system from elsewhere in the domain. Enterprise Admins membership allows other administrators to access the system from other domains in the current domain tree or forest. To prevent enterprisewide access to a domain, you can remove Enterprise Admins from this group.

Administrators is a local group that provides full administrative access to an individual computer or a single domain, depending on its location. Because this account has complete access, you should be very careful about adding users to this group. To make someone an administrator for a local computer or domain, all you need to do is make that person a member of this group. Only members of the Administrators group can modify this account.

Domain Admins is a global group designed to help you manage resources in a domain. Members of this group have full control of a domain. This group has administrative control over all computers in a domain because it's a member of the Administrators group by default on all domain controllers, all domain workstations, and all domain member servers at the time they join the domain. To make someone an administrator for a domain, make that person a member of this group.

**TIP** The Administrator account is a member of Domain Admins by default. This means that if a user logs on to a computer as the administrator and the computer is a member of the domain, the user will have complete access to all resources in the domain.

Enterprise Admins is a global group designed to help you manage resources in a forest. Members of this group have full control of all domains in a forest. This group has administrative control over all domain controllers in the enterprise because the group is a member of the Administrators group by default on all domain controllers in a forest. To make someone an administrator for the enterprise, make that person a member of this group.

**TIP** The Administrator account is a member of Enterprise Admins by default.





This means that if someone logs on to a computer as the administrator and the computer is a member of the domain, the user will have complete access to the domain tree or forest.

Schema Admins is a universal group designed to help you manage schema in Active Directory. Members of this group can work with and manage schema in the domain. Before someone can edit schema, they need to be a member of this group.



# Working with Implicit Groups

Windows Server defines a set of special identities you can use to assign permissions in certain situations. You usually assign permissions implicitly to special identities. However, you can assign permissions to special identities directly when you modify Active Directory objects. The special identities include the following:

- **Anonymous Logon** Any user accessing the system through anonymous logon has the Anonymous Logon identity. This identity allows anonymous access to resources such as a webpage published on a corporate intranet server.
- **Authenticated Users** Any user accessing the system through a logon process has the Authenticated Users identity. This identity allows access to shared resources within the domain, such as files in a shared folder that should be accessible to all the workers in the organization.
- **Batch** Any user or process accessing the system as a batch job (or through the batch queue) has the Batch identity. This identity allows batch jobs to run scheduled tasks, such as a nightly cleanup job that deletes temporary files.
- **Creator Group** Windows Server uses this special identity group to automatically grant access permissions to users who are members of the same group or groups as the creator of a file or a directory.
- **Creator Owner** The person who created the file or the directory is a member of this special identity group. Windows Server uses this identity to automatically grant access permissions to the creator of a file or directory.
- **Dial-Up** Any user accessing the system through a dial-up connection has the DialUp identity. This identity distinguishes dial-up users from other types of authenticated users.
- **Enterprise Domain Controllers** Domain controllers with enterprisewide roles and responsibilities have the Enterprise Domain Controllers identity. This identity allows them to perform certain tasks in the enterprise by using transitive trusts.
- **Everyone** All interactive, network, dial-up, and authenticated users are members of the Everyone group. This special identity group gives wide access to a system resource.
- **Interactive** Any user logged on to the local system has the Interactive identity. This identity allows only local users to access a resource.
- **Network** Any user accessing the system through a network has the Network identity. This identity allows only remote users to access a resource.
- **Proxy** Users and computers accessing resources through a proxy have the Proxy identity. This identity is used when proxies are implemented on the network. •



**Remote Desktop Services User** Any user accessing the system through Remote Desktop Services has the Remote Desktop Services User identity. This identity allows Remote Desktop Services users to access Remote Desktop Services applications and to perform other necessary tasks with Remote Desktop Services. •

**Restricted** Users and computers with restricted capabilities have the Restricted identity.

- **Self** The Self identity refers to the object itself and allows the object to modify itself.
- **Service** Any service accessing the system has the Service identity. This identity grants access to processes being run by Windows Server services.
- **System** The Windows Server operating system itself has the System identity. This identity is used when the operating system needs to perform a system-level function.







## Chapter 8. Managing Account Policies

User and group accounts allow Windows Server 2016 to track and manage information about users, including permissions and privileges. To create user accounts, you primarily use:

- [Active Directory Users And Computers](#), which is designed to administer accounts throughout an Active Directory domain
- [Local Users And Groups](#), which is designed to administer accounts on a local computer

The most important aspects of account creation are account setup and account organization. Without the appropriate guidelines and policies, you might quickly find that you need to rework all your user accounts. Before you create accounts, determine the policies you'll use for setup and organization.





# Understanding Account Naming Policies

The naming scheme for accounts is something you should consider as part of account management and general administration. User accounts have display names and logon names. The display name (or full name) is the name displayed to users and the name referenced in user sessions. The logon name is the name used to log on to the domain.

For domain accounts, the display name is normally the concatenation of the user's first name, middle initial, and last name, but you can set it to any string value. The display names must follow these rules:

- Local display names must be unique on an individual computer.
- Display names must be unique throughout a domain.
- Display names must be no more than 64 characters.
- Display names can contain alphanumeric characters and special characters.

Logon names must follow these rules:

- Local logon names must be unique on an individual computer, and global logon names must be unique throughout a domain.
- Logon names can contain as many as 256 characters. However, it isn't practical to use logon names that have more than 64 characters.
- A legacy logon name is given to all accounts. By default, this logon name is set to the first 20 characters of the Windows logon name. This logon name must be unique throughout a domain.
- Users logging on to the domain can use their standard logon names or their legacy logon names.
- Logon names can't contain certain characters. The following characters are invalid: `/ \ [ ] ; | = , + * ? < > .`
- Logon names can contain all other special characters, including spaces, periods, dashes, and underscores. Generally, however, it is not a good idea to use spaces in account names.

Windows Server stores user names in the case that you enter. However, user names aren't case sensitive. For example, you can access the Administrator account with the user name Administrator, administrator, or ADMINISTRATOR. Thus, user names are case aware but not case sensitive.

Most small organizations tend to assign logon names that use the user's first or last name. But you can have more than one person with the same name in an organization of any size. Rather than having to rework your logon naming scheme when you run





into problems, select a good naming scheme now and make sure that other administrators use it. You should use a consistent procedure for naming accounts—one that allows your user base to grow, limits the possibility of name conflicts, and ensures that your accounts have secure names that aren't easily exploited. If you follow these guidelines, the types of naming schemes you might want to use include the following:

- User's first name and last initial
- User's first initial and last name
- User's first initial, middle initial, and last name
- User's first initial, middle initial, and first five characters of the last name
- User's first name and last name





# Using Password and Account Policies

Domain accounts use passwords or private keys from certificates to authenticate access to network resources. This section focuses on passwords.





## Using Secure Passwords

A password is a case-sensitive string that can contain more than 127 characters with Active Directory and up to 14 characters with the legacy Security Manager. Valid characters for passwords are letters, numbers, and symbols. When you set a password for an account, Windows Server stores the password in an encrypted format in the account database.

Simply having a password isn't enough to ensure proper security. The key to preventing unauthorized access to network resources is to use secure passwords. The difference between an average password and a secure password is that secure passwords are difficult to guess and crack. You make passwords difficult to guess and crack by using combinations of all the available character types—including lowercase letters, uppercase letters, numbers, and symbols. For example, instead of using happydays for a password, you would use haPPy2Days&, Ha\*\*y!day5, or even h\*99Y%d\*ys.

You also might want to use password phrases. With a password phrase, the password contains multiple words and punctuation, like a sentence. For example, you might use the password phrase This problem is 99 times ten! A password phrase that includes punctuation and numbers meets all complexity requirements and is incredibly difficult to crack.

Unfortunately, no matter how secure you make a user's password initially, the user will eventually choose his own password. Therefore, you should set account policies that define a secure password for your systems. Account policies are a subset of the policies configurable in Group Policy.

## Setting Account Policies

As I mentioned in earlier chapters, you can apply group policies at various levels within the network structure. You manage local group policies in the manner discussed in the “Managing Local Group Policies” section in Chapter 14, “Using Group Policy.” You manage global group policies as explained in “Managing Enterprise Policies,” also in Chapter 14.


Account policies should be configured in the highest precedence GPO that is linked to a domain. By default, the highest precedence GPO linked to a domain is the Default Domain Policy GPO. After you access the Default Domain Policy GPO or other appropriate GPO, you can set account policies by following these steps:

1. In the Group Policy Management Editor, open the Account Policies node by expanding Windows Settings\Security Settings under Computer Configuration.



The console tree shows the name of the computer or domain you are configuring. Be sure that this is the appropriate network resource to configure.

**NOTE** Domain policies have precedence over local policies. The GPO with a link order of 1 in the domain always has the highest precedence.



2. You can now manage account policies through the Password Policy, Account Lockout Policy, and Kerberos Policy nodes. To configure a policy, double-click its entry, or right-click it and then click Properties. This opens a Properties dialog box for the policy.

All policies are either defined or not defined. That is, they are either configured for use or not configured for use. A policy that isn't defined in the current container could be inherited from another container.

**NOTE** Kerberos policies aren't used with local computers. Kerberos policies are available only with group policies that affect domains. For stand-alone servers, you can change the local policy settings. However, you cannot change the local policy settings for domain controllers or member servers.

3. Select or clear the Define This Policy Setting check box to specify whether a policy is defined.

**TIP** Policies can have additional options for configuration. Often these options are buttons labeled Enabled and Disabled. Clicking Enabled turns on the policy restriction. Clicking Disabled turns off the policy restriction. Some policies are negations, which means that by enabling them you are actually negating the item. For example, Disable Log On As A Service is the negation of the item Log On As A Service.

Specific procedures for working with account policies are discussed in the "Configuring Password Policies," "Configuring Account Lockout Policies," and "Configuring Kerberos Policies" sections later in this chapter.



# Configuring Password Policies

Password policies, listed here, control security for passwords:

- [Enforce Password History](#)
- [Maximum Password Age](#)
- [Minimum Password Age](#)
- [Minimum Password Length](#)
- [Password Must Meet Complexity Requirements](#)
- [Store Passwords Using Reversible Encryption](#)

The uses of these policies are discussed in the following sections.



## Enforce Password History

Enforce Password History sets how frequently old passwords can be reused. With this policy, you can discourage users from alternating between several common passwords. Windows Server can store up to 24 passwords for each user in the password history.



To disable this feature, set the value of the password history to 0. To enable this feature, set the value of the password history by using the Passwords Remembered box. Windows Server then tracks old passwords by using a password history that's unique for each user, and users aren't allowed to reuse any of the stored passwords.

**NOTE** To prevent users from bypassing settings for Enforce Password History, don't allow them to change passwords immediately. This stops users from changing their passwords several times to get back to an old password. You can set the time required to keep a password with the Minimum Password Age policy as discussed later in the chapter.

## Maximum Password Age

Maximum Password Age determines how long users can keep a password before they have to change it. The aim is to force users to change their passwords periodically. When you use this feature, set a value that makes sense for your network. Generally, you use a shorter period when security is very important and a longer period when security is less important.





You can set the maximum password age to any value from 0 through 999. A value of 0 specifies that passwords don't expire. Although you might be tempted to set no expiration date, users should change passwords regularly to ensure the network's security. Where security is a concern, good values are 30, 60, or 90 days. Where security is less important, good values are 120, 150, or 180 days.

**NOTE** Windows Server notifies users when the password expiration date is approaching. Any time the expiration date is fewer than 30 days away, a warning is displayed when the user logs on stating that he or she has to change the password within a specific number of days.

## Minimum Password Age

Minimum Password Age determines how long users must keep a password before they can change it. You can use this box to prevent users from bypassing the password system by entering a new password and then changing it right back to the old one.




If the minimum password age is set to 0, users can change their passwords immediately. To prevent this, set a specific minimum age. Reasonable settings are from three to seven days. In this way, you make sure that users are less inclined to switch back to an old password but are able to change their passwords in a reasonable amount of time if they want to. Keep in mind that a minimum password age could

prevent a user from changing a compromised password. If a user can't change the password, an administrator has to make the change.

## Minimum Password Length

Minimum Password Length sets the minimum number of characters for a password. If you haven't changed the default setting, you should do so immediately. The default in some cases is to allow empty passwords (passwords with zero characters), which is definitely not a good idea.




For security reasons, you'll generally want passwords of at least eight characters, because long passwords are usually harder to crack than short ones. If you want greater security, set the minimum password length to 14 characters.

## Password Must Meet Complexity Requirements

Beyond the basic password and account policies, Windows Server includes facilities for creating additional password controls. These facilities enforce the use of secure passwords that follow these guidelines:


- Passwords must have at least six characters.
- Passwords can't contain the user name, such as steve, or parts of the user's full name, such as steve.
- Passwords must use at least three of the four available character types: lowercase letters, uppercase letters, numbers, and symbols.

To enforce these rules, enable Passwords Must Meet Complexity Requirements.



## Store Passwords Using Reversible Encryption

Passwords in the password database are encrypted. This encryption can't normally be reversed. The only time you would want to change this setting is when your organization uses applications that need to read the password. If this is the case, enable Store Password Using Reversible Encryption For All Users.



With this policy enabled, passwords might as well be stored as plain text—it presents the same security risks. With this in mind, a much better technique is to enable the option on a per-user basis and then only as required to meet the user's actual needs.



# Configuring Account Lockout Policies

Account lockout policies, listed here, control how and when accounts are locked out of the domain or the local system:

- [Account Lockout Threshold](#)
- [Account Lockout Duration](#)
- [Reset Account Lockout Counter After](#)


These policies are discussed in the sections that follow.



## Account Lockout Threshold

Account Lockout Threshold sets the number of logon attempts that are allowed before an account is locked out. If you decide to use lockout controls, you should use a value that balances the need to prevent account cracking with the needs of users who are having difficulty accessing their accounts.


The main reason users might not be able to access their accounts properly the first time is that they forgot their passwords. If this is the case, they might need several attempts to log on properly. Workgroup users could also have problems accessing a remote system if their current passwords don't match the passwords that the remote system expects. For example, the remote system might record several bad logon attempts before a user receives a prompt to enter the correct password because Windows Server has attempted to automatically log on to the remote system. In a domain environment, this normally doesn't happen because of the single sign-on feature.



You can set the lockout threshold to any value from 0 to 999. The lockout threshold is set to 0 by default, which means that accounts won't be locked out because of invalid logon attempts. Any other value sets a specific lockout threshold. Keep in mind that the higher the lockout value, the higher the risk that a hacker might be able to break into your system. A reasonable range of values for this threshold is from 7 to 15. This is high enough to rule out user error and low enough to deter hackers.

## Account Lockout Duration

If someone violates the lockout controls, Account Lockout Duration sets the length of time that the account is locked. You can set the lockout duration to a specific length of time using a value between 1 and 99,999 minutes or to an indefinite length of time by setting the lockout duration to 0.




The best security policy is to lock the account indefinitely. When you do, only an administrator can unlock the account. This prevents hackers from trying to access the system again and forces users who are locked out to seek help from an administrator, which is usually a good idea. By talking to the user, you can determine what the user is doing wrong and help the user avoid further problems.

**TIP** When an account is locked out, open the Properties dialog box for the account in Active Directory Users And Computers. Click the Account tab, and then select the Unlock Account check box.

## Reset Account Lockout Counter After

Every time a logon attempt fails, Windows Server raises the value of a threshold that tracks the number of bad logon attempts. To maintain a balance between potential lockouts from valid security concerns and lockouts that could occur from simple human error, another policy determines how long to maintain information regarding bad logon attempts. This policy is called Reset Account Lockout Counter After, and you use it to reset the bad logon attempts counter to 0 after a certain waiting period. The way the policy works is simple: If the waiting period for Reset Account Lockout Counter After has elapsed since the last bad logon attempt, the bad logon attempts counter is reset to 0. The bad logon attempts counter is also reset when a user logs on successfully.





If the Reset Account Lockout Counter After policy is enabled, you can set it to any value from 1 to 99,999 minutes. As with Account Lockout Threshold, you need to select a value that balances security needs against user access needs. A good value is from one to two hours. This waiting period should be long enough to force hackers to wait longer than they want to before trying to access the account again.

If the Reset Account Lockout Counter After policy isn't set or is disabled, the bad logon attempts counter is reset only when a user successfully logs on.

**NOTE** Bad logon attempts against a password-protected screen saver at a workstation don't increase the lockout threshold. Similarly, if you press Ctrl+Alt+Delete to lock a server or workstation, bad logon attempts against the Unlock dialog box don't count.

## Configuring Kerberos Policies

Kerberos V5 is the primary authentication mechanism used in an Active Directory domain. The Kerberos protocol uses tickets to verify the identification of users and network services. Tickets contain encrypted data that confirms identity for the purposes of authentication and authorization.

You can control ticket duration, renewal, and enforcement with the following policies:

- [Enforce User Logon Restrictions](#)
- [Maximum Lifetime For Service Ticket](#)
- [Maximum Lifetime For User Ticket](#)
- [Maximum Lifetime For User Ticket Renewal](#)
- [Maximum Tolerance For Computer Clock Synchronization](#)

These policies are discussed in the sections that follow.

**SECURITY ALERT** Only administrators with an intimate understanding of Kerberos security should change these policies. If you change these policies to inefficient settings, you might cause serious problems on the network. The default Kerberos policy settings usually work just fine.



# Enforce User Logon Restrictions

Enforce User Logon Restrictions ensures that any restrictions placed on a user account are enforced. For example, if the user's logon hours are restricted, this policy enforces the restriction. By default, the policy is enabled and you should disable it only in rare circumstances.





# Maximum Lifetime

Maximum Lifetime For Service Ticket and Maximum Lifetime For User Ticket set the maximum duration for which a service or user ticket is valid. By default, service tickets have a maximum duration of 600 minutes, and user tickets have a maximum duration of 10 hours.



You can change the duration of tickets. For service tickets, the valid range is from 0 to 99,999 minutes. For user tickets, the valid range is from 0 to 99,999 hours. A value of 0 effectively turns off expiration. Any other value sets a specific ticket lifetime.



A user ticket that expires can be renewed, provided that the renewal takes place within the time set for Maximum Lifetime For User Ticket Renewal. By default, the maximum renewal period is seven days. You can change the renewal period to any




value from 0 to 99,999 days. A value of 0 effectively turns off the maximum renewal period, and any other value sets a specific renewal period.



# Maximum Tolerance

Maximum Tolerance For Computer Clock Synchronization is one of the few Kerberos policies you might need to change. By default, computers in the domain must be synchronized within five minutes of one another. If they aren't, authentication fails.



If you have remote users who log on to the domain without synchronizing their clocks to the network time server, you might need to adjust this value. You can set any value from 0 to 99,999. A value of 0 indicates that there's no tolerance for a time difference, which means the remote user's system must be precisely time-synchronized or authentication will fail.



# Configuring Authentication Policies

Authentication policies are an advanced configuration option for large enterprises. You use authentication policies to do the following:

- Restrict the authentication methods available to user, service, and computer accounts.
- Specify access control conditions that restrict the devices that can request Kerberos tickets for user, service and computer accounts.




# Getting Started with Authentication Policies

Authentication policies and related options are available only when an Active Directory domain is operating at the Windows Server 2012 R2 or higher functional level. Authentication policies are enforced during initial Kerberos authentication and whenever a user, computer or service requests access to a server services. You use authentication policy silos to identify the users, computers, managed services accounts, and group managed service accounts to which you want to apply authentication settings; to fine-tune which accounts are to be protected; and to define explicitly permitted accounts.

Several object classes in the Active Directory schema were implemented to support authentication policies and silos, including:

- **Authentication Policies container** Used to store Authentication Policy objects in Active Directory.
- **Authentication Policy object** Used to configure Kerberos Ticket Granting Ticker properties and specify authentication conditions based on account type.
- **Authentication Policy Silos container** Used to store Authentication Policy Silo objects in Active Directory Domain Services.
- **Authentication Policy Silo object** Defines the authentication policy or policies to be applied to accounts that inherit this object.

**NOTE** A default Authentication Policies container is created under the System container in a domain and it stores the Authentication Policy objects (APOs). You add APOs to the Authentication Policies Container to define the various sets of authentication settings you want to use.



You use authentication policies in conjunction with the Protected Users security group.



Any accounts that should have additional protections should be added to the Protected Users group, restricted using an Authentication Policy object and then optionally identified in an authentication policy silo to fine-tune the way controls are applied.

Membership in the Protected Users group ensures that only Kerberos can be used for authentication. Otherwise, if NTLM is used for authentication, access control conditions that you define in policies will not be properly applied.

Both types of policies can be configured in auditing only or enforce mode. With auditing only mode, policy restrictions are audited rather than enforced. With enforced mode, policy restrictions are strictly enforced.



# Implementing Authentication Policy

In Active Directory Administrative Center, you create and manage authentication policies using the options on the Authentication\Authentication Policies node. To create an authentication policy, select the Authentication\Authentication Policies node in the left pane. Next, under Tasks, select New and then select Authentication Policy. In the Create Authentication Policy dialog box, use the options provided to define the new policy and specify access control conditions for each type of account to which the policy will apply.

The screenshot shows the 'Create Authentication Policy' dialog box. The left pane has a tree view with 'General' selected. The main area is titled 'General' and contains the following fields and options:

- Display name:** A text box with a red asterisk indicating it is required.
- Description:** A text box.
- Protect from accidental deletion:** A checked checkbox.
- Policy Enforcement:** Two radio buttons: 'Only audit policy restrictions' (unselected) and 'Enforce policy restrictions' (selected). Below the radio buttons is a note: 'Note: Audit policy applied through a silo will'.

Below the 'General' section is an 'Accounts' section with a table header showing 'Name' and 'Account Type'. An 'Add...' button is located at the bottom right of the 'Accounts' section.

In Active Directory Administrative Center, you create and manage authentication silos by using the options on the Authentication\Authentication Policy Silos node. To create an authentication policy silo, select the Authentication\Authentication Policies node in the left pane. Next, under Tasks, select New and then select Authentication Policy Silo. Use the options in the Create Authentication Policy Silo dialog box to create the authentication policy silo.

The screenshot shows the 'Create Authentication Policy Silo' dialog box. The left pane has a tree view with 'General' and 'Policy' selected. The main area is titled 'General' and contains the following fields and options:

- Display name:** A text box with a red asterisk indicating it is required.
- Description:** A text box.
- Protect from accidental deletion:** A checked checkbox.
- Policy Enforcement:** Two radio buttons: 'Only audit silo policies' (selected) and 'Enforce silo policies' (unselected).

Below the 'General' section is a 'Permitted Accounts' section with a table header showing 'Name', 'Account Type', and 'Assigned'. An 'Add...' button is located at the bottom right of the 'Permitted Accounts' section.

Authentication Policy Silos allow you to apply authentication policies in several different ways. When you create the silo object, you can specify a single authentication policy to apply to all member accounts or use a separate authentication policy for each

type of security principal. As part of creating the authentication policy silo, you also can specify explicitly permitted accounts.

Authentication Policy

☐ Use a single policy for all principals that belong to this authentication policy silo.

The authentication policy that applies to all accounts in this silo:

☒ Use a separate authentication policy for each type of principal.

\*

User account policy:

Open

\*

Managed Service Account policy:

Open

\*

Computer account policy:

Open



# Configuring User Rights Policies

User accounts have built-in capabilities and user rights. Although you can't change built-in capabilities for accounts, you can manage user rights for accounts. Typically, you apply user rights to users by making them members of the appropriate group or groups. You can also apply rights directly, and you do this by managing the user rights for the user's account.


**SECURITY ALERT** Any user who's a member of a group that's assigned a certain right also has that right. For example, if the Backup Operators group has the right, and jsmith is a member of this group, jsmith also has this right. Keep in mind that changes you make to user rights can have a far-reaching effect.

Because of this, only experienced administrators should make changes to the user rights policy.

You assign user rights through the Local Policies node of Group Policy. As the name implies, local policies pertain to a local computer. However, you can configure local policies and then import them into Active Directory. You can also configure these local policies as part of an existing GPO for a site, a domain, or an organizational unit. When you do this, the local policies apply to computer accounts in the site, domain, or organizational unit.

To administer user rights policies, follow these steps:


1. Open the GPO you want to work with, and then open the Local Policies node by working your way down the console tree. To do so, expand Windows Settings\Security Settings\Local Policies under Computer Configuration.
2. Select User Rights Assignment to manage user rights. To configure a user rights assignment, double-click a user right, or right-click it and then click Properties. This opens a Properties dialog box.
3. Configure the user rights. To configure local user rights, follow steps 1–3 in the “Configuring Local User Rights.” section later in this chapter. To configure global user rights, follow steps 1–6 in the following section.



# Configuring Global User Rights

For a site, a domain, or an organizational unit, you configure individual user rights by following these steps:

1. Open the Properties dialog box for the user right. If the policy isn't defined, select Define These Policy Settings.
2. To apply the right to a user or group, click Add User Or Group. Then, in the Add User Or Group dialog box, click Browse. This opens the Select Users, Computers, Service Accounts, Or Groups dialog box.
3. Enter the name of the user or group you want to use in the box provided, and then click Check Names. By default, the search is configured to find built-in security principals and user accounts. To add groups to the search, click Object Types, select Groups in the list box, and then click OK.
4. After you select the account names or groups to add, click OK. The Add User Or Group dialog box should now show the selected accounts. Click OK again.
5. The Properties dialog box is updated to reflect your selections. If you made a mistake, select a name and remove it by clicking Remove.
6. When you have finished granting the right to users and groups, click OK.



Select Users, Computers, Service Accounts, or Groups

Select this object type:

Users, Service Accounts, Groups, or Built-in security principals

Object Types...

From this location:

imaginedlands.local

Locations...

Enter the object names to select (examples):

Engineering

Check Names

Advanced...

OK

Cancel



# Configuring Local User Rights

For local computers, apply user rights by following these steps:

1. Open the Properties dialog box for the user right. Remember that site, domain, and organizational unit policies have precedence over local policies.
2. The Properties dialog box shows current users and groups that have been given a user right. To remove the user right, select a user or group and then click Remove.
3. You can apply the user right to additional users and groups by clicking Add User Or Group. This opens the Select Users, Computers, Service Accounts, Or Groups dialog box. You can now add users and groups.







## Chapter 9. Creating Accounts

As part of routine administration, you'll often need to create accounts for users, groups and computers. You create user accounts for each user who requires access to network resources and group accounts to manage privileges for multiple users. You use computer accounts to control access to the network and its resources. After you create user accounts, you may need to provide contact information. After you create groups, you may need to manage membership. These and other tasks are discussed in this chapter.

For more information on managing and maintaining accounts, see Chapter 11, "Managing Computers, Users and Groups."





# Adding User Accounts

You can create domain accounts with either Active Directory Users And Computers or

Active Directory Administrative Center and local accounts with Local Users And Groups. When you are working with Active Directory, keep in mind that you typically create users and groups in an organizational unit or in the Users container rather than the domain container itself. When you are working with the administrative center, be sure to use the Tree View so that the step-by-step instructions required are the same as provided.



# Creating Domain Users

Whether you prefer to create users in Active Directory Users And Computers or Active Directory Administrative Center, the options available are similar. To get started, simply right-click the container in which you want to place the user account, click New, and then click User.

The screenshot shows the 'Create User' dialog box for 'William R. Stanek'. The 'Account' tab is selected. The 'Full name' field is highlighted with a red asterisk, indicating it is required. The 'User UPN logon' field is set to 'wstanek' and the 'User SamAccountName' field is set to 'imaginedlands\wstanek'. The 'Password' field is empty. The 'Account expires' section has 'Never' selected. The 'Password options' section has 'Other password options' selected. The 'Encryption options' section has 'Store password using reversible encryption' selected. The 'Other options' section has 'Account is sensitive and cannot be delegated' and 'Do not require Kerberos pre-authentication' selected.

You'll then need to configure the user display name and logon name. Enter the user's first name, middle initials, and last name in the boxes provided. These boxes are used to create the full name, which is the user's display name.

Make changes to the Full Name field as necessary. For example, you might want to enter the name in LastName FirstName MiddleInitial format or in FirstName MiddleInitial LastName format. The full name must be unique in the domain and must have 64 or fewer characters.

Next, enter the user's logon name and use the drop-down list to select the domain to associate the account with. This sets the fully qualified logon name.

The first 20 characters of the logon name are used to set the SAM account name (which is also referred to as the pre-Windows 2000 account name). The logon name and the SAM account name must be unique in the domain. If necessary, change the account name.

Finally, you'll need to manage the account password. The related options include:

- **Password** The password for the account. This password should follow the conventions of your password policy.



- **Confirm Password** A box to ensure that you assign the account password correctly. Simply reenter the password to confirm it.
- **User Must Change Password At Next Logon** If selected, the user must change the password upon logon.
- **User Cannot Change Password** If selected, the user can't change the password.
- **Password Never Expires** If selected, the password for this account never expires. This setting overrides the domain account policy. Generally, it's not a good idea to set a password so that it doesn't expire—this defeats the purpose of having passwords in the first place.
- **Account Is Disabled** If selected, the account is disabled and can't be used. Use this check box to temporarily prevent anyone from using an account.

If you have problems creating the account, you receive a warning, and you need to use the Back button to reenter information in the user name and password pages as necessary.

# Creating Local Users

You create local user accounts with Local Users And Groups. You can open this utility and create an account by following these steps:

1. In Server Manager, click Tools and then click Computer Management.  
Alternatively, you can press Windows+X and then click Computer Management.
2. Right-click the Computer Management entry in the console tree, and then click Connect To Another Computer. You can now choose the system whose local accounts you want to manage. Domain controllers don't have local users and groups.
3. Under System Tools, choose Local Users And Groups.
4. Right-click Users, and then click New User. This opens the New User dialog box. You use each of the boxes in the dialog box as follows:
  - **User Name** The logon name for the user account. This name should follow the conventions for the local user name policy.
  - **Full Name** The full name of the user, such as William R. Stanek.
  - **Description** A description of the user. Typically, you enter the user's job title, such as Webmaster. You could also enter the user's job title and department.
  - **Password** The password for the account. This password should follow the conventions of your password policy.
  - **Confirm Password** A second entry to ensure that you assign the account password correctly. Simply reenter the password to confirm it.
  - **User Must Change Password At Next Logon** If selected, the user must change the password upon logon.
  - **User Cannot Change Password** If selected, the user can't change the password.
  - **Password Never Expires** If selected, the password for this account never expires.  
This setting overrides the local account policy.
  - **Account Is Disabled** If selected, the account is disabled and can't be used. Use this check box to temporarily prevent anyone from using an account.
5. Click Create when you have finished configuring the new account.

## Adding Groups

You use group accounts to manage privileges for multiple users. You can create global group in either Active Directory Users And Computers or Active Directory Administrative Center and local groups in Local Users And Groups.

As you set out to create group accounts, remember that you create group accounts for similar types of users. The types of groups you might want to create include the following:

- **Groups for departments within the organization** Generally, users who work in the same department need access to similar resources. You'll often create groups that are organized by department, such as Business Development, Sales, Marketing, and Engineering.
- **Groups for users of specific applications** Users often need access to an application and resources related to the application. If you create application-specific groups, you can be sure that users have proper access to the necessary resources and application files.
- **Groups for roles within the organization** You can also organize groups by user roles within the organization. For example, executives probably need access to different resources than supervisors and general users do. By creating groups based on roles within the organization, you can ensure that proper access is given to the users who need it.

## Creating Global Groups

As with user accounts, the options for creating groups are similar whether you are working with Active Directory Users And Computers or Active Directory Administrative Center. To get started, simply right-click the container in which you want to place the user account, click New, and then click Group.

The screenshot shows the 'Create Group: AllMarketing' dialog box. The 'Group' section is active, showing the group name 'AllMarketing' and the group type 'Global'. The 'Managed By' section is also visible, showing the manager's name and contact information. The 'Protect from accidental deletion' checkbox is checked.

**Create Group: AllMarketing**

**Group**

Group name: AllMarketing

Group (SamAccountName): AllMarketing

Group type: ☒ Security ☐ Distribution

Group scope: ☐ Domain local ☒ Global ☐ Universal

☒ Protect from accidental deletion

E-mail:

Create in: CN=Users,DC=imaginedlands,DC=local

Description:

Notes:

**Managed By**

Managed by:

☐ Manager can update membership list

Phone numbers:

Main:

Mobile:

Fax:

Office:

Address:

City: State/Province: Zip/Postal code:

Country/Region:

More Information

OK Cancel

**NOTE** Optionally, when you are working with the administrative center, select **Protect From Accidental Deletion** to mark the account as protected in Active Directory. Protected accounts can be deleted only if you remove the Protect flag



prior to attempting to delete the account. To market accounts as protected when you are using Active Directory Users And Computers, you must displayed advanced settings, and then edit the group properties after creating the group.

You'll then need to configure the group name. Group names aren't case sensitive and can be up to 64 characters.

The first 20 characters of the group name are used to set the SAM account name for the group (which is also referred to as the pre-Windows 2000 group name). The group name and the SAM account name must be unique in the domain. If necessary, change the account name.

Next, set the group scope as either Domain Local, Global, or Universal and then select a group type as either Security or Distribution. When you create a distribution group and have integrated Active Directory with Exchange Server, a default email address based on the group name and domain is set for the group automatically. With the administrative center, you can specify a different email address when you create the group. With Active Directory Users And Computers, you must edit the group properties after creating the group to change the default email address.

After you create the account, you can add members and set additional properties, as discussed later in this chapter.





## Creating Local Groups

You create local groups with Local Users And Groups. You can access this utility and create a group by following these steps:

1. In Server Manager, click Tools and then click Computer Management. Rightclick the Computer Management entry in the console tree, and then click Connect To Another Computer. You can now choose the system whose local accounts you want to manage. Domain controllers don't have local users and groups.
2. Under System Tools, choose Local Users And Groups.
3. Right-click Groups, and then click New Group. This opens the New Group dialog box.
4. After you enter a name and description of the group, click the Add button to add names to the group. This opens the Select Users dialog box.
5. In the Select Users dialog box, enter the name of a user you want to use in the Name box, and then click Check Names. If matches are found, select the account you want to use, and then click OK. If no matches are found, update the name you entered and try searching again. Repeat this step as necessary, and then click OK.
6. The New Group dialog box is updated to reflect your selections. If you made a mistake, select a name and click Remove.
7. Click Create when you've finished adding or removing group members.

## Adding Computer Accounts

Computer accounts are stored in Active Directory as objects. You use them to control access to the network and its resources. You can add computer accounts to any standard container displayed in Active Directory Users And Computers. The best folders to use are Computers, Domain Controllers, and any organizational units you've created.

The easiest way to create a computer account is to log on to the computer you want to configure and then join a domain, as described in the "Joining a Computer to a Domain or Workgroup" section in Chapter 11. When you do this, the necessary computer account is created automatically and placed in the Computers folder. You can also create a computer account in either Active Directory Users And Computers or Active Directory Administrative Center before you try to install the computer.




# Creating Standard and Managed Computer Accounts

You can create two types of computer accounts: standard computer accounts and managed computer accounts. Managed computer accounts are available when you've installed Windows Deployment Services in your domain.

Using Active Directory Users And Computers, you can create a standard computer account by following these steps:

1. In the Active Directory Users And Computers console tree, right-click the container in which you want to place the computer account, click New, and then click Computer.



The screenshot shows the 'New Object - Computer' dialog box. At the top, it says 'Create in: imaginedlands.local/Computers'. Below this, there are two text boxes: 'Computer name:' with the value 'Creatives29' and 'Computer name (pre-Windows 2000):' with the value 'CREATIVES29'. A line of text states: 'The following user or group can join this computer to a domain.' Below this, there is a text box labeled 'User or group:' containing 'Default: Domain Admins' and a 'Change...' button. At the bottom, there is a checkbox labeled 'Assign this computer account as a pre-Windows 2000 computer:' which is currently unchecked. At the very bottom are three buttons: 'OK', 'Cancel', and 'Help'.

2. Enter a name for the computer in the Computer Name text box. The name you provide also sets the NetBIOS name, which is used primarily for legacy operations.
3. By default, only members of Domain Admins can join this computer to the domain. To allow a different user or group to join the computer to the domain, click Change, and then select a user or group account in the Select User Or Group dialog box.

**NOTE** You can select any existing user or group account. This allows you to delegate the authority to join this computer account to the domain.

4. If this account will be used with applications written for legacy operating systems, select Assign This Computer Account As A Pre–Windows 2000 Computer.
5. If Windows Deployment Services are not installed, click OK to create the computer account. Otherwise, click Next twice, and then click Finish.

When you are working with Windows Deployment Services, managed computer accounts are used to prestage computer accounts so that a computer can be automatically installed. Using Active Directory Users And Computers, you can create a managed computer account by following these steps:

1. Complete steps 1–4 in the previous procedure. Click Next to display the Managed page.
2. Select the This Is A Managed Computer check box, and then enter the computer's globally unique identifier/universally unique identifier (GUID/UUID). Click Next.
3. On the Host Server page, you have the option to specify which host server to use or to allow any available host server to be used for remote installation. To select a host server, select The Following Remote Installation Server. In the Find dialog box, click Find Now to display a list of all remote installation servers in the organization. Click the host server you want to use, and then click OK to close the Find dialog box.
4. Click Next, and then click Finish.

**REAL WORLD** You can find the GUID/UUID in the system BIOS or displayed on the computer case. If Windows PowerShell is installed, you can collect the GUID/UUID by using the Win32\_ComputerSystemProduct class of the Windows Management Instrumentation (WMI) interface. The following example returns the universally unique identifier (UUID) of the computer you are logged on to:

```
get-wmiobject -class win32_computersystemproduct | fl uuid
```

Here, you return the UUID of a remote computer:


```
get-wmiobject -class win32_computersystemproduct -computername  
engpc24 | format-list pscomputername, uuid
```

After you create a standard or manager computer account in Active Directory Users And Computers, you might want to mark the account as protected. Protected accounts can be deleted only if you remove the Protect flag prior to attempting to delete the account.



To mark a computer account as protected, follow these steps:

- 1. In Active Directory Users And Computers, ensure Advanced Features is selected on the View menu.
- 2. Double-click the computer account to open its Properties dialog box.
- 3. On the Object tab, select Protect Object From Accidental Deletion, and then click OK.






# Creating Computer Accounts in Administrative Center

Using Active Directory Administrative Center, you can create a standard computer account, add the account as a member of specific groups, and set properties about the manager of the computer. To do this, follow these steps:

1. Open Active Directory Administrative Center. In the left pane, List View is selected by default. Click Tree View to display the tree view.
2. In the Active Directory Administrative Center console tree, right-click the container in which you want to place the computer account, click New, and then click Computer. This opens the Create Computer dialog.



3. Enter a name for the computer in the Computer Name text box. The name you provide also sets the NetBIOS name, which is used primarily for legacy operations.
4. By default, only members of Domain Admins can join this computer to the domain. To allow a different user or group to join the computer to the domain, click Change, and then select a user or group account in the Select User Or Group dialog box.

**NOTE** You can select any existing user or group account. This allows you to delegate the authority to join this computer account to the domain.

5. If the computer account will be used with applications written for legacy operating systems, select Assign This Computer Account As A Pre-Windows 2000 Computer.
6. Optionally, select Protect From Accidental Deletion to mark the account as protected in Active Directory. Protected accounts can be deleted only if you remove the Protect flag prior to attempting to delete the account.
7. Optionally, assign a security principal as the manager of the computer by clicking Edit under Managed By, and then selecting a user or group to designate as the manager in the Select User Or Group dialog box. Who you assign as a computer's manager depends on corporate policy and can include



the primary user of the computer, a branch manager at a particular office, or a support contact.

8. The computer account will be added to the appropriate default computer group automatically. Typically, this is Domain Computers. You can add the computer account to other groups by clicking Add under Member Of and then using the Select Groups dialog box to specify groups that have accounts to which the computer account should belong.
  9. Click OK to create the computer account.
- 



# Managing User Contact Information

When you create user accounts, those accounts can have detailed contact information associated with them, including:

- [Job title](#)
- [Business address](#)
- [Department](#)
- [Telephone number](#)
- [Manager and direct reports](#)

The contact information is then available for anyone in the enterprise to use as criteria to search for users and to create address book entries. When you are setting contact information, keep in mind you generally only want to provide business-related contact details. Thus, while you might want to provide the business address where a user works or business phone number, you typically would not provide a user's home address or home phone number. If you have any questions about what information can be provided, refer to company policy or discuss the matter with your human resources and legal departments.



## Setting Contact Information

You can set a user's contact information in Active Directory Users And Computers by following these steps:

1. Double-click the user name in Active Directory Users And Computers. This opens the account's Properties dialog box.
2. Click the General tab. As needed, set general contact information, including:
  - **First Name, Initials, Last Name** The user's full name.
  - **Display Name** The user's display name as shown in logon sessions and in Active Directory.
  - **Description** A description of the user.
  - **Office** The user's office location.
  - **Telephone Number** The user's primary business telephone number. If the user has other business telephone numbers you want to track, click Other, and then enter additional phone numbers in the Phone Number (Others) dialog box.
  - **E-Mail** The user's business email address.
  - **Web Page** The URL of the user's web page, which can be on the Internet or on the company intranet. If the user has other webpages you want to track, click Other, and then enter additional webpage addresses in the Web Page Address (Others) dialog box.

The image displays two side-by-side screenshots of the 'Tom Stevens Properties' dialog box, specifically the 'General' tab. The left screenshot shows the 'General' tab with fields for First name (Tom), Last name (Stevens), Display name (Tom Stevens), Office (Engineering), Telephone number (206-555-1212), E-mail (tomstevens@imaginedlands.com), and Web page. The right screenshot shows the 'General' tab with fields for Job Title (Development Manager), Department (Engineering), Company (Imagined Lands LLC), Manager Name (William Stanek), and Direct reports.

**TIP** You must fill in the E-Mail and Web Page boxes if you want to use the Send Mail and Open Home Page features of Active Directory Users And Computers.

3. On the Address tab, as needed, set the user's business location or business address in the boxes provided.
4. A user's primary business phone number should be entered on the General tab. If a user has other phone numbers, such as for FAX or mobile phone, click the Telephones tab and enter the phone numbers that should be used. Generally, you should not provide personal or home numbers without prior authorization.

**NOTE** You can configure other numbers for each type of telephone number. Click the associated Other buttons, and then enter additional phone numbers in the dialog box provided.

5. As appropriate, use the text boxes provided on the Organization tab to enter the user's job title, department, and company. To specify the user's manager, click Change, and then use the Select User Or Contact dialog box to specify the user's manager. When you specify a manager, the user shows up as a direct report in the manager's account.
6. Click OK to apply the changes.

You also can set contact information using Active Directory Administrative Center. Double-click the user name. In the account's Properties dialog box, click Organization to display the Organization panel. This panel provides a one-stop location for setting general, address, telephone, and organization details.

The screenshot shows the 'Organization' tab for user 'Tom Stevens'. The left sidebar contains a navigation menu with options: Account, Organization (selected), Member Of, Password Settings, Profile, Policy, Silo, and Extensions. The main content area is divided into several sections:

- Display name:** Tom Stevens
- Office:** Engineering
- E-mail:** tomstevens@imaginedlands.com
- Web page:** (empty text box)
- Other web pages...** (button)
- Phone numbers:**
  - Main:** 206-555-1212
  - Home:** (empty text box)
  - Mobile:** (empty text box)
  - Fax:** (empty text box)
  - Pager:** (empty text box)
  - IP Phone:** (empty text box)
  - Other phone numbers...** (button)
- Descriptions:** (empty text box)
- Job title:** Development Manager
- Department:** Engineering
- Company:** Imagined Lands LLC
- Manager:** William Stanek (with 'Edit...' and 'Clear' buttons)
- Direct reports:** (empty list with 'Add...' and 'Remove' buttons)
- Address:**
  - Street:** (empty text box)
  - City:** (empty text box)
  - State/Province:** (empty dropdown)
  - Zip/Postal code:** (empty text box)
  - Country/Region:** (empty dropdown)

The Web Page box sets the URL of the user's home page, which can be on the Internet or on the company intranet. If the user has other webpages you want to track, click Other Web Pages and then enter additional webpage addresses in the Web Page Address (Others) dialog box.

Under Phone Numbers, enter the primary telephone numbers that should be used to contact the user, such as main, home, mobile, fax, pager, and IP phone. You can configure other numbers for each type of telephone number. Click Other Phone



Numbers, and then enter additional phone numbers in the dialog box provided.

If the user has a manager set, this is listed in the Manager box. If no manager is set or you want to change the manager, click the Edit button next to Manager to specify the user's manager in the Select User Or Contact dialog box. When you specify a manager, the user shows up as a direct report in the manager's account.

If the user has direct reports, they are listed under Direct Reports. You can add and remove direct reports by using the related Add and Remove buttons. To add a direct report, click Add, specify the direct report in the Select User Or Contact dialog box, and then click OK. To remove a direct report, click the name in the list, and then click Remove.



# Searching for Users and Groups

Active Directory makes it easy for you to find users and groups in the directory, which you can do by following these steps:

1. In Active Directory Users And Computers, right-click the starting container, and then click Find. The domain or OU you right-click sets the base context for the search.
2. In the Find Users, Contacts, And Groups dialog box, the In list shows the previously selected domain or container. If you want to search the entire directory instead, select Entire Directory, or click Browse to select a domain or container to search.
3. On the Users, Contacts, And Groups tab, enter the name of the user, contact, or group you want to search for.
4. Click Find Now to begin the search. If matches are found, the search results are displayed. Otherwise, enter new search parameters and search again.
5. To manage an account, right-click its entry. If you right-click an account entry and then select Properties, you can open the account's Properties dialog box.

Find Users, Contacts, and Groups

File Edit View

Find: Users, Contacts, and Groups In: Engineering Browse...

Users, Contacts, and Groups Advanced

Name: Mark

Description:

Find Now Stop Clear All

Search results:

| Name           | Type | Description |
|----------------|------|-------------|
| Mark Jefferson | User |             |
| Mark Thomas    | User |             |
| Mark Bradley   | User |             |

You can also search for users and groups by using the filter and global search features of

Active Directory Administrative Center. For more information, see the “Using



Administrative Center and Windows PowerShell” section in Chapter 5, “Managing Active Directory.”







# Managing Group Membership

To configure group membership, you use Active Directory Users And Computers or Active Directory Administrative Center. When working with groups, keep the following points in mind:

- All new domain users are members of the group Domain Users, and their primary group is specified as Domain Users.
- All new domain workstations and member servers are members of Domain Computers, and their primary group is Domain Computers.
- All new domain controllers are members of Domain Controllers, and their primary group is Domain Controllers.

You manage group membership several ways:


- Manage individual memberships
- Manage multiple memberships
- Set primary group membership for individual users and computers






## Managing Membership in Groups

You can quickly add a user or a group to one or more groups by right-clicking the account and then selecting Add To Group. This opens the Select Groups dialog box. You can now choose groups that the currently selected account should be a member of.



You can manage group membership for any type of account by following these steps:

1. Double-click the user, computer, or group entry in Active Directory Users And Computers or Active Directory Administrative Center. This opens the account's Properties dialog box.
2. On the Member Of tab or panel, groups that the user is currently a member of are listed. Click Add to make the account a member of an additional group. This opens the Select Groups dialog box. You can now choose groups that the currently selected account should be a member of.
3. To remove the account from a group, select a group, and then click Remove.
4. Click OK.



## Adding or Removing Multiple Accounts


If you're working exclusively with user accounts, you can add one or more users to groups by following these steps:



1. Select the user accounts you want to work with in Active Directory Users And Computers or Active Directory Administrative Center.

**TIP** To select multiple users individually, hold down the Ctrl key and then click each user account you want to select. To select a sequence of accounts, hold down the Shift key, select the first user account, and then select the last user account.

2. Right-click one of the selections, and then click Add To A Group or Add To Group as appropriate. This opens the Select Groups dialog box. You can now choose groups that the currently selected accounts should be members of.
3. Click OK.



Another way to add or remove multiple accounts is to use a group's Properties dialog box. To do this, follow these steps:

1. Double-click the group entry in Active Directory Users And Computers or Active Directory Administrative Center. This opens the group's Properties dialog box.
2. On the Members tab or panel, current members of the group are listed alphabetically. To add accounts to the group, click Add. This opens the Select Users, Contacts, Computers, Service Accounts, Or Groups dialog box. You can now choose users, computers, service accounts, and groups that should be members of the currently selected group.
3. To remove members from a group, select an account, and then click Remove.
4. Click OK.

Domain Admins

TASKSSECTIONSD

Group

Managed By

Member Of

Members

Password Settings

Extensions

Members

Filter

Name

Active Directory Domain Services Fol...

Administrator

imaginedlands-Users-Administrator

Ed Green

imaginedlands-Users-Ed Green

Tom S

imaginedlands-Users-Tom S

Add..

Remove


# Setting the Primary Group for Users and Computers

Users who access Windows Server from Macintosh or POSIX-compliant applications use primary groups. When a Macintosh user creates files or directories on a system running Windows Server, the primary group is assigned to these files or directories.

**NOTE** All user and computer accounts must have a primary group regardless of whether the accounts access Windows Server systems through Macintosh. This group must be a group with global or universal scope, such as the global group Domain Users or the global group Domain Computers.

To set the primary group, follow these steps:

1. Double-click the user or computer entry in Active Directory Users And Computers or Active Directory Administrative Center. This opens the account's Properties dialog box.
2. On the Member Of tab or panel, select a group with global or universal scope in the Member Of list.
3. Click Set Primary Group.



All users must be a member of at least one primary group. You can't revoke membership in a primary group without first assigning the user to another primary group. To do this, follow these steps:

1. Select a different group with global or universal scope in the Member Of list, and then click Set Primary Group.
2. In the Member Of list, click the former primary group, and then click Remove. The group membership is now revoked.



# Chapter 10. Working with Managed Accounts

In addition to traditional types of accounts discussed in previous chapters, Windows and Windows Server support three additional types of accounts:

- [Managed service accounts](#)
- [Managed virtual accounts](#)
- [Group managed service accounts](#)

These special types of accounts have specific usage scenarios and are the subject of this chapter.





## Using Managed Accounts

Microsoft Exchange Server, Microsoft Internet Information Services, Microsoft SQL Server, and other types of applications often use service accounts. On a local computer, you can configure the application to run as a built-in user account, such as Local Service, Network Service, or Local System. Although these service accounts are easy to configure and use, they usually are shared among multiple applications and services and cannot be managed on a domain level. If you configure the application to use a domain account, you can isolate the privileges for the application, but then you must manually manage the account password and any service principal names (SPNs) required for Kerberos authentication. To resolve these and other issues, Windows and Windows Server support managed service accounts, managed virtual accounts, and group managed service accounts.

Managed service accounts are a special type of domain user account for managed services. These accounts reduce service outages and other issues by having Windows manage the account password and related SPNs automatically.

Managed virtual accounts are a special type of local computer account for managed services. These accounts provide the ability to access the network with a computer identity in a domain environment. Because the computer identity is used, no password management is required.

All current releases of Windows and Windows Server support an extension of the original implementation of service and virtual accounts called group managed service accounts. Group managed service accounts provide the same functionality as standard managed service accounts but extend that functionality over multiple servers. As an example, when a client computer connects to a service hosted by a server farm, mutual authentication cannot succeed unless all the instances of the services use the same principal. By using a group managed service account, you allow each server in the farm to use the same service principal, which is managed by Windows itself rather than individually by the administrator.

Group managed service accounts are, in fact, the default type of service account for all current releases of Windows and Windows Server. Because of this, managed service accounts can span multiple computers by default. This means you can add the account to more than one computer at a time as necessary to support clustered nodes, network load-balancing server farms, and so on. If you want to restrict a managed service account to a single computer, set the `-RestrictToSingleComputer` option when creating the account. Don't forget that a single computer can also have multiple managed service accounts.



In Active Directory schema, managed service accounts are represented by `msDSManagedServiceAccounts`. This object class inherits its attributes from the `Computer` object class, but the objects also are users. Managed service accounts use the same password-update mechanism as regular computer accounts. This means the password for the account is updated whenever the computer updates its password, which by default occurs every 30 days. Managed service accounts can automatically maintain their Kerberos SPN and support delegation.

**TIP** Some applications, such as SQL Server and IIS, make use of Kerberos extensively and know how to register themselves with SPNs. If an application supports writing its own SPNs, managed service accounts will work for automatic SPN management.

**NOTE** By default, all managed service accounts are created in the Managed Service Accounts container in Active Directory. This container is visible in Active Directory Users And Computers when you display advanced features.

Like computer accounts, managed service accounts do not use either domain or finegrained password policies. Instead, they use a randomly generated 240-byte (120character) password. Managed service accounts cannot perform interactive logons or be locked out like user accounts can be. You can add managed service accounts to groups by using Active Directory Users And Computers or `Add-ADGroupMember`.

## Implementing Managed Accounts

With managed service accounts, you create an actual account, which is stored by default in the Managed Service Accounts container in Active Directory. Next, you associate the account with a computer in Active Directory and then install the managed service account on a local server to add it to the account as a local user. Finally, you configure the local service to use the account. Put another way, you must do the following:

1. Create the managed service account.
2. Associate the account with a computer in Active Directory.
3. Install the managed service account on the computer that was associated.
4. Configure the local service to use the account.

You can use Windows PowerShell cmdlets to install, uninstall, and reset passwords for managed service accounts. After a managed service account has been installed, you can configure a service or application to use the account and no longer have to specify or change passwords because the account password is maintained by the







computer. You can also configure the SPN on the service account without requiring domain administrator privileges.



# Creating Managed Accounts

You create a managed service account by using `New-ADServiceAccount`. The basic syntax is as follows:

```
New-ADServiceAccount -DisplayName DisplayName -SamAccountName SAMName  
-Name Name [-RestrictToSingleComputer]
```

`DisplayName` is the display name for the account, `SAMName` is the pre-Windows 2000 name of the account, and `Name` is the pre-Windows 2000 name of the account, such as:

```
New-ADServiceAccount -DisplayName "SQL Agent Account" -  
SamAccountName sqlagent -Name "SQL Agent"
```

The account will be created as a group account by default. It will have a randomly generated 240-byte (120-character) password and be created in the Managed Service Accounts container. The account is also enabled by default, but you can create the account in a disabled state by adding `-Enabled $false`. If you need to pass in credentials to create the account, use the `-Credential` parameter, as shown in this example:

```
$cred = Get-Credential  
New-ADServiceAccount -DisplayName "IIS App Pool 1"  
-SamAccountName pool1 -Name "IIS Pool 1" -Credential $cred
```

Although the account is listed in Active Directory Users And Computers, you shouldn't use this management tool to work with the account. Instead, you should use the following Windows PowerShell cmdlets:

- [Get-ADServiceAccount](#), to get information about one or more managed service accounts
- [Set-ADServiceAccount](#), to set properties on an existing managed service account
- [Remove-ADServiceAccount](#), to remove a managed service account from Active Directory

## Selecting Target Computers

After you create a managed service account in Active Directory, you associate it with a target computer in Active Directory by using `Add-ADComputerServiceAccount`. You use `Remove-ADComputerServiceAccount` to remove a computer association from Active Directory.

The basic syntax for `Add-ADComputerServiceAccount` is as follows:

```
Add-ADComputerServiceAccount [-Identity] ComputerName [-  
ServiceAccount] MSAName
```



ComputerName is the name of the target computer, and MSAName is the name of the managed service account, such as:

```
Add-ADComputerServiceAccount IISServer84 WebServicesAccount
```

If you need to pass in credentials to create the account, use the `–Credential` parameter, as shown in this example:

```
$cred = Get-Credential
```

```
Add-ADComputerServiceAccount IISServer32 FarmFourServicesAccount
```





## Installing Managed Accounts

You can install the account on a local computer by using `Install-ADServiceAccount`. The basic syntax is this:

```
Install-ADServiceAccount [-Identity] ServiceAccountId
```

`ServiceAccountId` is the display name or SAM account name of the service account, such as:

```
Install-ADServiceAccount sqlagent
```

If you need to pass in credentials to create the account, use the `–Credential` parameter. Use `Uninstall-ADServiceAccount` to uninstall an account.



# Configuring Services to Use Managed Accounts

You can configure a service to run with the managed service account by following these steps:

1. In Server Manager, click Tools, and then click Computer Management.
2. As necessary, connect to the computer you want to manage. In the left pane, rightclick the Computer Management node, and then click Connect To Another Computer. Enter the host name, fully qualified domain name, or IP address of the remote server, and then click OK.
3. In the left pane, expand the Services And Applications node, and then select the Services node.
4. Right-click the name of the service you want to work with, and then click Properties.
5. On the Log On tab, select This Account, and then enter the name of the managed service account in the format AccountName@DomainName ; or click Browse to search for the account.

6. Confirm that the password box is blank, and then click OK.
7. Select the name of the service, and then click Start to start the service, or click Restart to restart the service as appropriate. Confirm that the newly configured account name appears in the Log On As column for the service.

**NOTE** A dollar sign (\$) appears at the end of the account name in the Services snap-in console. When you use the Services console to configure the logon as an account, the Service Logon Right logon right is automatically assigned to the account. If you use a different tool, the account has to be explicitly granted this right.



## Maintaining Managed Accounts

After you create a managed service account, there isn't much you need to do maintain the account. That said, you may find that you need to move a managed account from one computer to another or you may need to remove a managed account that is no longer needed.





## Moving Managed Accounts

To move a managed service account from a source computer to a new destination computer, you need to do the following:

1. On the source computer, configure any services that are using the managed account to use a different account, and then run `Uninstall-ADServiceAccount`.
2. On the new destination computer, run `Install-ADServiceAccount`, and then use the Services snap-in console to configure the service to run with the managed service account.

To migrate a service from a user account to a managed service account, you need to do the following:

1. Create a new managed service account in Active Directory by using `NewADServiceAccount`.
2. Install the managed service account on the appropriate computer by using `InstallADServiceAccount`, and then use the Services snap-in console to configure the service to run with the managed service account.
3. You also might need to configure the access control lists on the service resources for the service management account.





## Removing Managed Accounts

If a managed service account is no longer being used on a computer, you might want to uninstall the account. Before you do this, however, you should check the Services snapin to ensure that the account isn't being used. To uninstall a managed service account from a local computer, use `Uninstall-ADServiceAccount`. The basic syntax is shown here:

```
Uninstall-ADServiceAccount -Identity ServiceAccountId
```

`ServiceAccountId` is the display name or SAM account name of the service account, such as:

```
Uninstall-ADServiceAccount -Identity sqlagent
```

If you need to pass in credentials to uninstall the account, use the `-Credential` parameter.

Managed service account passwords are reset on a regular basis based on the password reset requirements of the domain, but you can reset the password manually if needed. To

reset the password for a managed service account, use `ResetADServiceAccountPassword`. The basic syntax is as follows:

```
Reset-ADServiceAccountPassword -Identity ServiceAccountId
```

`ServiceAccountId` is the display name or SAM account name of the service account, such as:

```
Reset-ADServiceAccountPassword -Identity sqlagent
```

If you need to pass in credentials to reset the password, use the `-Credential` parameter. You can modify the default password change interval for managed service accounts by using the domain policy `Domain Member: Maximum Machine Account Password Age` under `Local Policies\Security Options`. Group Policy settings under `Account Policies>Password Policy` are not used to modify managed service account passwordreset intervals, nor can the `NLTEST /SC_CHANGE_PWD` command be used to reset managed service account passwords.

## Using Virtual Accounts

Virtual accounts require very little management. They cannot be created or deleted, and they do not require any password management. Instead, they exist automatically and are represented by the machine identity of the local computer.





With virtual accounts, you configure a local service to access the network with a computer identity in a domain environment. Because the computer identity is used, no account needs to be created and no password management is required.

You can configure a service to run with a virtual account by following these steps:

1. In Server Manager, click Tools, and then click Computer Management.
2. As necessary, connect to the computer you want to manage. In the left pane, rightclick the Computer Management node, and then click Connect To Another Computer. Enter the host name, fully qualified domain name, or IP address of the remote server, and then click OK.
3. In the left pane, expand the Services And Applications node, and then select the Services node.
4. Right-click the name of the service you want to work with, and then click Properties.
5. On the Log On tab, select This Account, and then enter the name of the service account in the format `SERVICE \ ComputerName` .



6. Confirm that the password box is blank, and then click OK.
7. Select the name of the service, and then click Start to start the service, or click Restart to restart the service. Confirm that the newly configured account name appears in the Log On As column for the service.

**NOTE** A dollar sign (\$) appears at the end of the account name in the Services snap-in console. When you use the Services snap-in console to configure the logon as an account, the Service Logon Right logon right is automatically assigned to the account. If you use a different tool, the account has to be explicitly granted this right.



# Chapter 11. Managing Computers, Users and Groups


Account management is an important part of routine administration. Whether you are working in engineering, IT, support or elsewhere in technology, you may need to update, remove or otherwise work with accounts. This chapter provides guidelines and tips to make the related tasks easier.




# Maintaining Computer Accounts

Using Active Directory Users And Computers or Active Directory Administrative Center, you can work with properties of computer accounts by following these steps:

1. In the console tree, expand the domain node.
2. Select the container or organizational unit in which the computer account is located.
3. Double-click the account. This displays a Properties dialog box that allows you to view and edit settings.



In Active Directory Users And Computers, advanced tabs and settings are available only when the Advanced Features option is selected on the View menu. In Active Directory Administrative Center, most advanced options are available via tabs on the Extensions panel.





If you need to manage a computer, you can use Computer Management. Whether you're working with Active Directory Users And Computers or Active Directory Administrative Center, you can open Computer Management and connect to a specific computer by right-clicking the computer entry and then selecting Manage from the shortcut menu. This starts Computer Management and automatically connects to the selected computer.



# Resetting Locked Computer Accounts

Computer accounts have passwords, just like user accounts. Unlike user accounts, however, computer account passwords are managed and maintained automatically. To perform this automated management, computers in the domain store a computer account password, which is changed every 30 days by default, and a secure channel password for establishing secure communications with domain controllers. The secure channel password is also updated by default every 30 days, and both passwords must be synchronized. If the secure channel password and the computer account password get out of sync, the computer won't be allowed to log on to the domain, and a domain authentication error message will be logged for the Netlogon service with an event ID of 3210 or 5722.

If this happens, you need to reset the computer account password. One way to do this is to right-click the computer account in Active Directory Users And Computers and then select Reset Account. You then need to remove the computer from the domain (by making the computer a member of a workgroup or another domain) and then rejoin the computer to the domain.

**REAL WORLD** Several other ways to reset the computer account password in the same way are available. In Active Directory Administrative Center, you right-click the computer account, and then select Reset Account. At the command prompt, you can use

`Dsmod Computer -Reset` to reset a computer account password. In Windows PowerShell, you can use `Reset-ComputerMachinePassword` (or the `Set-ADAccountPassword` with the `-Reset` option) to reset a computer account password. The following command runs `Reset-ComputerMachinePassword` on a remote computer:

```
Invoke-Command -ComputerName EngPC84 -ScriptBlock {Reset-ComputerMachinePassword}
```

All of these options might require the additional steps of removing the computer from the domain (by making the computer a member of a workgroup or another domain) and then rejoining the computer to the domain. The additional steps might be required because the password must be synchronized between the local computer and the domain.

Several tools allow you to reset a computer's password and sync the changes in the domain. On computers where Windows PowerShell is installed, you can use `TestComputerSecureChannel` to test the secure connection between a local computer

and a domain. You'll want to log on locally to the computer, open a Windows PowerShell prompt, and then enter the command:

```
test-computersecurechannel
```

You can use the `-Server` option to test the communications channel with a specific domain controller. If the command returns `False`, there is a communications problem, and you can use the `-Repair` option to reset the account password on the local computer and write this change to the related Computer object on a domain controller in the domain. The password change is then replicated to other domain controllers.

Another tool for resetting a computer's password and syncing changes is the Netdom command-line utility. You can use Netdom Verify to test the secure connection between a local computer and a domain. You can use Netdom Resetpwd to reset the account password on the local computer and write this change to the related Computer object on a domain controller in the domain, which in turn ensures the password change is replicated to other domain controllers.

For a member server, you can reset the computer account password by following these steps:

1. Log on locally to the computer. At a command prompt, enter **netdom resetpwd /s:ServerName /ud:domain\UserName/pd:\*** , where **ServerName** is the name of the domain controller to use to set the password, **domain\UserName** specifies an administrator account with the authority to change the password, and **\*** indicates that Netdom should prompt you for the account password before continuing.
2. Enter your password when prompted. Netdom changes the computer account password locally and on the domain controller. The domain controller distributes the password change to other domain controllers in the domain.
3. Restart the computer.

For domain controllers, you must perform additional steps. After you log on locally, you must stop the Kerberos Key Distribution Center service and set its startup type to Manual. After you restart the computer and verify that the password has been successfully reset, you can restart the Kerberos Key Distribution Center service and set its startup type back to Automatic.

## Moving Computer Accounts

Computer accounts are normally placed in the Computers or Domain Controllers containers, or in customized organizational unit containers. You can move an account to a different container by selecting the computer account in Active Directory Users



And Computers and then dragging the account to the new location. You can't click and drag accounts in Active Directory Administrative Center.

Using either tool, you can also use the following technique to move computer accounts:

1. In the console tree, select the container in which the computer account is located.
2. Right-click the computer account you want to move, and then click Move.
3. In the Move dialog box, use the options provided to select the container to which you want to move the computer. Browse to subcontainers or child OUs as necessary, and then click OK.



## Disabling and Deleting Computer Accounts

If you no longer need a computer account, you can delete it permanently from Active Directory. You can also temporarily disable the account and later enable it to be used again.

To delete, disable, or enable computer accounts, follow these steps:

1. Open Active Directory Users And Computers or Active Directory Administrative Center. In the console tree, select the container in which the computer account is located.
2. Right-click the computer account, and then do one of the following:
  - Click **Delete** to delete the account permanently. Click **Yes** to confirm the deletion.
  - Click **Disable Account** to temporarily disable the account, and click **Yes** to confirm the action. A red circle with an X indicates that the account is disabled.
  - Click **Enable Account** to enable the account so that it can be used again.

If the account is protected, you need to clear the Protect flag before you can delete the account. Double-click the account to open its Properties dialog box. Clear the Protect Object From Accidental Deletion check box, and then click OK. In Properties dialog boxes for Active Directory Users And Computers, this check box is on the Object tab. In Active Directory Administrative Center, this check box is on the Computer panel.

**TIP** If an account is currently in use, you might not be able to disable it. Try shutting down the computer or disconnecting the computer session in the Sessions folder of Computer Management.

## Joining a Computer to a Domain or Workgroup

A computer joined to a domain or workgroup can log on and access the network. Before you get started, make sure that networking components are properly installed on the computer. These should have been installed during the setup of the operating system. TCP/IP settings must be correct and permit communications between the computer you're configuring and a controller in the domain. If Dynamic Host Configuration Protocol (DHCP), Windows Internet Name Service (WINS), and DNS are properly installed on the network, workstations don't need to be assigned a static IP address or have a special configuration. The only requirements are a computer name and a domain name, which you can specify when joining the computer to the domain.

**REAL WORLD** Windows Server 2016 automatically grants the Add Workstations To The Domain user right to the implicit group Authenticated Users. This means that any user who logs on to the domain as a User and is authenticated can add





workstations to the domain without needing administration privileges. However, as a security precaution, the number of workstations any such user can add to the domain is limited to 10. If an authenticated user exceeds this limit, an error message is displayed.

Although you can use the Ldp.exe tool from the Windows Server 2016 Support Tools to override the default limit on the number of computers an authenticated user can join to a domain (as set by the ms-DS-MachineAccountQuota attribute), this isn't a good security practice. A better technique, and a more appropriate technique where security is a concern, is to create the necessary computer account in a specific OU beforehand or to grant the user the advanced security privilege Create Account Objects for the Computers container. You also might want to grant certain users the Delete Account Objects privilege for the Computers container so that designated users can remove computer accounts from the domain.

During installation of the operating system, a network connection was probably configured for the computer, or you might have previously joined the computer to a domain or a workgroup. If so, you can join the computer to a new domain or workgroup. For joining a computer to a domain, see the "Configuring Server Naming" section in Chapter 3, "Configuring Server Settings."

If the name change is unsuccessful, a message is displayed informing you that the change was unsuccessful or that the account credentials already exist. This problem can occur when you're changing the name of a computer that's already connected to a domain and when the computer has active sessions in that domain. Exit applications that might be connected to the domain, such as File Explorer accessing a shared folder over the network. Then repeat the process for changing the computer's name.

If you have other problems joining a domain, be sure that the computer you're configuring has the proper networking configuration. The computer must have Networking Services installed, and TCP/IP properties must have the correct DNS server settings.

All authenticated users have the Add Workstations To The Domain user right by default and can create up to 10 computer accounts when joining computers to a domain. Users who have the Create Account Objects privilege for the Computers container aren't restricted in this way; they can create an unlimited number of computer accounts in the domain. However, computer accounts created by authenticated users have Domain Admins as the account owner, whereas computer accounts created by users with the



Create Account Objects privilege have the creator as the owner. If you grant the Create Account Objects privilege, you also might want to grant the Delete Account Objects privilege so that certain users can remove computer accounts from the domain.

You grant the Create Account Objects, the Delete Account Objects privilege, or both privileges for the Computers container by following these steps:

1. Open Active Directory Users And Computers or Active Directory Administrative Center. In Active Directory Users And Computers, ensure Advanced Features is enabled on the View menu.
2. Right-click the Computers container, and then click Properties.
3. On the Security tab, click Advanced. In the Advanced Security Settings For Computers dialog box, click Add to open the Permission Entry For Computers dialog box.
4. Click Select A Principal. In the Select User, Computer, Service Account, Or Group dialog box, enter the name of the user or group to whom you want to grant privileges, and then click OK. Click OK again.

## Using Offline Domain Join

Computers running editions of Windows designed for business support offline domain join, as does Windows Server. The related utility, Djoin.exe, is included with these editions of Windows. Any member of Domain Admins can perform offline domain joins (as can anyone who is granted the appropriate user rights).

The basic steps for performing an offline domain join operation follow:

1. Create the computer account in Active Directory, and then force replication of the shared secrets of the computer that is to join the domain.
2. Write the relevant state information that the computer needs to join the domain to a text file, and then make the state information available to the computer.
3. When the computer starts, Windows reads the provisioning data, and the computer is joined to the domain.

**NOTE** Client computers must be connected to the corporate network to join a domain or receive domain settings. With the remote domain join feature, Windows Server makes it possible for computers to join a domain and receive domain settings remotely over the Internet.

You run Djoin.exe at an administrator command prompt to provision the computer account metadata. The computer account metadata is written to a .txt file. After



provisioning the computer, you can run Djoin.exe again to request the computer account metadata and insert it into the Windows directory of the destination computer. Alternatively, you can save the computer account metadata in an Unattend.xml file and then specify the Unattend.xml file during an unattended operating system installation.

You can use a .txt file for provisioning by following these steps:

1. Using an account that is allowed to join computers to the domain, log on to a computer that is a member of the domain.
2. Use Djoin.exe to create a text file that contains the computer account metadata. To do this, at an administrator command prompt, enter **djoin /provision /domain DomainName /machine MachineName /savefile FileName** , where DomainName is the name of the domain to join, MachineName is the computer name, and FileName is name of the .txt file where the metadata should be saved, such as:

```
djoin /provision /domain tvpress /machine HrComputer15 /savefile  
Hrcomputer15.txt
```

**TIP** By default, computer accounts are created in the Computers container. If you want to use a different container, you can add the /Machineou parameter and then specify the container to use. If the computer account object is already created, you can still generate the required metadata by adding the /reuse parameter. If your domain controller is not yet running Windows Server 2008 R2 or Windows Server 2016, add the /downlevel command.

3. On the new computer, use Djoin.exe to import the .txt file. At an administrator command prompt, enter **djoin /requestODJ /loadfile FileName/windowspath %SystemRoot% /localos** , where FileName is the name of the metadata file, such as:

```
djoin /requestODJ /loadfile HrComputer15.txt /windowspath  
%SystemRoot% /localos
```

4. Ensure that the new computer is connected to the network, and then restart it. During startup, the computer will be joined to the domain.

You can use an Unattend.xml file for provisioning by creating a section in the Unattend.xml file and then adding the contents of the metadata .txt file to the – AccountData element, as shown in this example:

```
<Component>  
<Component name=Microsoft-Windows-UnattendedJoin>  
  <Identification>  
    <Provisioning>
```

```
<AccountData> Insert metadata here! </AccountData>
</Provisioning>
</Identification>
</Component>
```

After you create the Unattend.xml file, start the new computer in safe mode or start the computer in the Windows Preinstallation Environment (Windows PE), and then run the Setup command with an answer file, as shown in the following example:

```
setup /unattend: FullPathToAnswerFile
```

Here, FullPathToAnswerFile is the full file path to the Unattend.xml file.

## Configuring Paths and Profiles

User accounts can have profiles, logon scripts, and home directories associated with them. To configure these optional settings, double-click a display name in Active Directory Users And Computers, and then click the Profile tab. On the Profile tab, you can provide the following settings:

- **Profile Path** The path to the user's profile. Profiles provide the environment settings for users. Each time a user logs on to a computer, that user's profile is used to determine desktop and Control Panel settings, the availability of menu options and applications, and more.
- **Logon Script** The path to the user's logon script. Logon scripts are batch files that run whenever a user logs on. You use logon scripts to set commands that should be executed each time a user logs on.
- **Home Folder** The directory the user should use for storing files. Here, you assign a specific directory for the user's files as a local path on the user's system or on a connected network drive. If the directory is available to the network, the user can access the directory from any computer on the network, which is a definite advantage.

Mark Thomas Properties

Published Certificates Member Of Password Replication Dial-in Object

Remote Desktop Services Profile COM+ Attribute Editor

Security Environment Sessions Remote control

General Address Account Profile Telephones Organization

User profile

Profile path: \\services18\userprofiles\markt

Logon script: \\fileserver23\userlogin\cs-supp.vbs

Home folder

☐ Local path:

☒ Connect: Z: To: \\custsupp03\markt

In Active Directory Administrative Center, you configure a user's environment settings by using the options on the Profile panel. To configure these settings, double-click a display name in Active Directory Administrative Center, and then click Profile to display the Profile panel.

Profile

Profile path: \\services18\userprofiles\markt Log on script: \\fileserver23\userlogin\cs-supp.vbs

Home folder:


☐ Local path:

☒ Connect: Z: To: \\custsupp03\markt


## Using Environment Variables in Paths

System environment variables are often helpful when you're setting up the user's environment, especially when you work with logon scripts. You use environment variables to specify path information that can be dynamically assigned. You use the following environment variables the most:

- **%SystemRoot%** The base directory for the operating system, such as C:\Windows. Use it with the Profile tab of the user's Properties dialog box and logon scripts.
- **%UserName%** The user account name, such as wrstanek. Use it with the Profile tab of the user's Properties dialog box and logon scripts.
- **%HomeDrive%** The drive letter of the user's home directory followed by a colon, such as C:. Use it with logon scripts.
- **%HomePath%** The full path to the user's home directory on the respective home drive, such as \Users\Mkg\TinaH. Use it with logon scripts.
- **%Processor\_Architecture%** The processor architecture of the user's computer, such as x86. Use it with logon scripts.



Note that by using the %UserName% variable, you allow the system to determine the full path information on a user-by-user basis. If you use this technique, you can use the same path information for multiple users, and all the users will have unique settings.



## Using Logon Scripts

Logon scripts set commands that should be executed each time a user logs on. You can use logon scripts to set the system time, network drive paths, network printers, and more. Although you can use logon scripts to execute one-time commands, you shouldn't use them to set environment variables. Any environment settings used by scripts aren't maintained for subsequent user processes. Also, you shouldn't use logon scripts to specify applications that should run at startup. You should set startup applications by placing the appropriate shortcuts in the user's Startup folder.

Normally, logon scripts contain Windows commands. However, logon scripts can be any of the following:

- Windows PowerShell scripts with a .ps1 or other valid extension
- Windows Script Host files with .vbs, .js, or another valid script extension
- Batch files with the .bat extension
- Command files with the .cmd extension
- Executable programs with the .exe extension

One user or many users can use a single logon script. As the administrator, you control which users use which scripts. As the name implies, logon scripts are accessed when users log on to their accounts. You can specify a logon script by following these steps:

1. Open the user's Properties dialog box in Active Directory Users And Computers, and then click the Profile tab.
2. Enter the path to the logon script in the Logon Script box. Be sure to set the full path to the logon script, such as **\\CSServices08\UserLogin\Eng.vbs**.

Creating logon scripts is easier than you might think, especially when you use the Windows command language. Just about any command you can enter into a command prompt can be set to run in a logon script. The most common tasks you'll want logon scripts to handle are to set the default printers and network paths for users.

With Microsoft Visual Basic Scripting Edition (VBScript), you need to initialize the variables and objects you plan to use and then call the appropriate methods of the Network object to add the connections. Consider the following example:

Option Explicit



```
Dim wNetwork, printerPath
Set wNetwork = WScript.CreateObject("WScript.Network")
```

```
printerPath = "\\gandolf\techmain"
wNetwork.AddWindowsPrinterConnection printerPath
wNetwork.SetDefaultPrinter printerPath
wNetwork.MapNetworkDrive "G:", "\\frodo\corpfiles"
```

```
Set wNetwork = vbEmpty
Set printerPath = vbEmpty
```

Here, you use the `AddWindowsPrinterConnection` method to add a connection to the TechMain printer on Gandolf, and you use the `SetDefaultPrinter` method to set the printer as the default for the user. You then use the `MapNetworkDrive` method to define a network drive on drive G.





# Assigning Home Directories

Windows Server 2016 lets you assign a home directory for each user account. Users can store and retrieve their personal files in this directory. Many applications use the home directory as the default for File Open and File Save As operations, which helps users find their resources easily. The command prompt also uses the home directory as the initial current directory.

Home directories can be located on a user's local hard drive or on a shared network drive. On a local drive, the directory is accessible only from a single workstation, however, shared network drives can be accessed from any computer on the network, which makes for a more versatile user environment.

If you don't assign a home directory, Windows Server 2016 uses the default local home directory. To specify a local home directory, follow these steps:

1. Open the user's Properties dialog box in Active Directory Users And Computers, and then click the Profile tab.
2. Select Local Path in the Home Folder section, and then enter the path to the home directory in the associated box, such as **C:\Home\%UserName%**.

To specify a network home directory, follow these steps:

1. Open the user's Properties dialog box in Active Directory Users And Computers, and then click the Profile tab.
2. In the Home Folder section, select the Connect option, and then select a drive letter for the home directory. For consistency, you should use the same drive letter for all users. Also, be sure to select a drive letter that won't conflict with any currently configured physical or mapped drives. To avoid problems, you might want to use Z as the drive letter.
3. Enter the complete path to the home directory by using Universal Naming Convention (UNC) notation, such as **\\Support12\UserDirs\%UserName%**. You include the server name in the drive path to ensure that the user can access the directory from any computer on the network.

**NOTE** You don't need to create the user's home directory ahead of time. Active Directory Users And Computers automatically creates the directory for you. If there's a problem creating the directory, Active Directory Users And Computers will instruct you to create it manually. Additionally, although users can share home directories, it's not a good idea. You'll usually want to provide each user with a unique home directory.



# Controlling Logon and Access

Windows Server 2016 gives you many ways to control user accounts and their access to the network. You can define logon hours, permitted workstations for logon, and more.



# Restricting Logon Times

Windows Server 2016 lets you control when users can log on to the network. You do this by setting their valid logon hours. You can use logon hour restrictions to tighten security and prevent system cracking or malicious conduct after normal business hours.

During valid logon hours, users can work as they typically do. They can log on to the network and access network resources. During restricted logon hours, users can't work. They can't log on to the network or make connections to network resources. If users are logged on when their logon time expires, what happens depends on the account policy you've set for them? Generally, one of two things happens to the user:

- **Forcibly disconnected** You can set a policy that tells Windows Server to forcibly disconnect users when their logon hours expire. If this policy is set, remote users are disconnected from all network resources and logged off the system when their hours expire.
- **Not disconnected** Users aren't disconnected from the network when their logon hours expire. Instead, Windows Server doesn't allow them to make any new network connections.


**NOTE** When you set logon hours, you'll save yourself a lot of work in the long run if you give users a moderately restricted time window. For example, rather than explicit hours matching a user's work schedule, you might want to allow a few hours on either side of the normal work hours. This lets early birds onto the system and allows night owls to keep working until they finish for the day. To configure the logon hours, follow these steps:

1. Open the user's Properties dialog box. In Active Directory Users And Computers, click the Account tab, and then click Logon Hours. In Active Directory Administrative Center, click Log On Hours on the Account panel.

The screenshot shows the 'Account' tab of a user's properties dialog box. On the left, there are fields for 'First name' (Mark), 'Middle initials', 'Last name' (Jefferson), 'Full name' (Mark Jefferson), 'User UPN logon' (markj@imaginedlands.io), and 'User SamAccountName L...' (imaginedlands\markj). There is also a checkbox for 'Protect from accidental deletion'. On the right, there are sections for 'Account expires' (set to 'Never'), 'Password options' (including 'User must change password at next log on', 'Other password options', 'Microsoft Passport or smart card is required for interactive log...', 'Password never expires', and 'User cannot change password'), 'Encryption options', and 'Other options'. At the bottom left, the 'Log on hours...' button is highlighted with a dashed rectangular box.



2. You can now set the valid and invalid logon hours by using the Log On Hours dialog box. In this dialog box, you can turn on or off each hour of the day or night.
- Hours that are allowed are filled in with a dark bar—you can think of these hours as being turned on.
  - Hours that are disallowed are blank—you can think of these hours as being turned off.
3. To change the setting for an hour, click it and then select either Logon Permitted or Logon Denied. Keep the following in mind:
- **All** Allows you to select all the time periods and then either permit or deny logon.
  - **Days of the Week buttons** Allow you to select all the hours in a particular day and then either permit or deny logon for that day of the week.
  - **Hourly Buttons** Allow you to select a particular hour for all the days of the week and then either permit or deny logon for that hour.




To forcibly disconnect users when their logon hours expire, follow these steps:

1. Access the highest precedence GPO that is linked to a domain. By default, this is the Default Domain Policy GPO. For details on working with GPOs, see “Navigating Enterprise Policies” section in Chapter 14.
2. After you access the Default Domain Policy GPO or other appropriate GPO open the Security Options node by working your way down through the console tree. Expand Windows Settings\Security Settings under Computer Configuration. In Security Settings, expand Local Policies, and then select Security Options.



3. Double-click Network Security: Force Logoff When Logon Hours Expire. This opens a Properties dialog box for the policy.
4. Select the Define This Policy Setting check box, and then click Enabled. This turns on the policy restriction and enforces the logon hours. Click OK.




# Restricting Workstation Logons

By default, you can use any valid user account to log on locally to a workstation. Thus, unless you restrict workstation use, anyone who obtains a user name and password can


use them to log on to any workstation in the domain. If you want to restrict the workstations a user can log on to, you can define a permitted workstation list.

For domain users, you specify permitted logon workstations by following these steps:

1. Open the user's Properties dialog box. In Active Directory Users And Computers, click the Account tab, and then click Log On To. In Active Directory Administrative Center, in the Account panel, click Log On To.

The screenshot shows the 'Account' dialog box for a user named Mark Jefferson. The 'Log on to...' button is highlighted with a dashed box. The dialog includes fields for first name, middle initials, last name, full name, user UPN logon, and user SamAccountName. It also has sections for account expiration, password options, encryption options, and other options.

2. For the This User Can Log On To option, select The Following Computers.

The screenshot shows the 'Log on to' dialog box. The 'The following computers' radio button is selected. A list of computers (tspc-b83, tspc-d42, aspc-c17) is shown in the 'Add >>' button's dropdown menu. The dialog also includes an 'OK' button and a 'Cancel' button.

3. Enter the name of a permitted workstation, and then click Add. Repeat this procedure to specify additional workstations.
4. If you make a mistake, select the invalid entry and then click Remove.

# Configuring VPN Privileges and Remote Access

You configure remote access privileges for accounts on the Dial-In tab. I know, I know, no one has dial-in these days, but that's still where you go to manage settings for virtual private networks (VPNs) . Basically, you either explicitly allow or deny remote access, or you specify that remote access is controlled Network Policy Server (NPS) Network Policy (which is the default and preferred method of controlling remote access) . In any event, before users can remotely access the network, you need to follow these steps:

1. In Server Manager, add the following roles: Network Policy And Access Services and Remote Access.



2. Configure remote access by using Routing And Remote Access. In Computer Management, expand Services And Applications, and then select Routing And Remote Access. Configure Routing And Remote Access as appropriate.

**REAL WORLD** Binaries needed to install roles and features are referred to as payloads. With Windows Server 2016, you remove the payloads for roles and features

that you are uninstalling by using the `-Remove` parameter of the `UninstallWindowsFeature` cmdlet. Restore a removed payload by using the `InstallWindowsFeature` cmdlet. By default, payloads are restored via Windows Update. Use the `-Source` parameter to restore a payload from a Windows Imaging (WIM) format mount point. In the following example, you restore the NPS and Routing And Remote Access Service (RRAS) binaries via Windows Update:

```
install-windowsfeature -name npas-policy-server  
-includemanagementtools
```

```
install-windowsfeature -name remoteaccess -  
includeallsubfeature -includemanagementtools
```





## Managing Security Options

The Account tab/panel of the user's Properties dialog box has the following options, which are designed to help you maintain a secure network environment and control how user accounts are used:

- **User Must Change Password At Next Logon** Forces the user to change his password when the user logs on next.
- **User Cannot Change Password** Doesn't allow the user to change the account password.
- **Password Never Expires** Ensures that the account password never expires, which overrides the normal password expiration period.

**CAUTION** Selecting this option creates a security risk on the network. Although you might want to use Password Never Expires with some accounts, you usually shouldn't.

- **Store Password Using Reversible Encryption** Saves the password as encrypted clear text.
  - **Account Is Disabled** Disables the account, which prevents the user from accessing the network and logging on (Active Directory Users And Computers only).
  - **Smart Card Is Required For Interactive Logon** Requires the user to log on to a workstation by using a smart card. The user can't log on to the workstation by entering a logon name and password at the keyboard.
  - **Account Is Sensitive And Cannot Be Delegated** Specifies that the user's account credentials cannot be delegated by using Kerberos. Use this for sensitive accounts that should be carefully controlled.
  - **Use Kerberos DES Encryption Types For This Account** Specifies that the user account will use Data Encryption Standard (DES) encryption.
  - **This Account Supports Kerberos AES 128 Bit Encryption** Specifies that the account supports Advanced Encryption Standard (AES) 128-bit encryption.
  - **This Account Supports Kerberos AES 256 Bit Encryption** Specifies that the account supports AES 256-bit encryption.
  - **Do Not Require Kerberos Preauthentication** Specifies that the user account doesn't need Kerberos preauthentication to access network resources.  
Preauthentication is part of the Kerberos 5 security procedure. The option to log on without it allows authentication from clients who are using a previous, or nonstandard, implementation of Kerberos.
- 





While working with the Account tab/panel set or clear options as appropriate.

## Working with User Profiles

Every user account has a profile. Profiles control startup features for the user's session, the types of programs and applications that are available, the desktop settings, and a lot more. Problems with a profile can sometimes prevent a user from logging on. For example, if the display size in the profile isn't available on the system being used, the user might not be able to log on properly. In fact, the user might get nothing but a blank screen. You could restart the computer, go into VGA mode, and then reset the display manually. However, solutions for profile problems aren't always this easy, and you might need to update the profile itself.



## Using Roaming and Mandatory Profiles

Each computer that a user logs on to has a copy of the user's profile. Because this profile is stored on the computer's hard disk, users who access several computers have a profile on each computer. Another computer on the network can't access a locally stored profile—called a local profile—and, as you might expect, this has some drawbacks. For example, if a user logs on to three different workstations, the user could have three very different profiles—one on each system. As a result, the user might get confused about what network resources are available on a given system.

Rather than a local profile, most users prefer to have roaming profiles which allow them to use the same profile no matter which computer they're using within the domain. Roaming profiles are server-based and can be stored on any server running Windows. When a user with a roaming profile logs on and the profile isn't present, the profile is downloaded, which creates a local copy. When the user logs off, changes to the profile are updated both in the local copy and on the originating server.

When your organization uses the Encrypting File System (EFS) to make file access more secure, the use of roaming profiles becomes extremely important for users who log on to multiple computers. This is because encryption certificates are stored in user profiles, and the encryption certificate is needed to access and work with the user's encrypted files. If a user has encrypted files and doesn't have a roaming profile, that user won't be able to work with these encrypted files on another computer—unless she uses credential roaming with Digital ID Management Service (DIMS).

As an administrator, you can control user profiles administratively or let users manage their own profiles. One reason to control profiles yourself is to make sure that all users have a common network configuration, which can reduce the number of environment-related problems.

Profiles controlled by administrators are called mandatory profiles. Users who have a mandatory profile can make only transitory changes to their environment. Any changes that users make to the local environment aren't saved, and the next time they log on, they're back to the original profile. The idea is that if users can't permanently modify the network environment, they can't make changes that cause problems. A key drawback to mandatory profiles is that the user can log on only if the profile is accessible. If, for some reason, the server that stores the profile is inaccessible and a cached profile isn't accessible, the user normally won't be able to log on. If the server is inaccessible, but a cached profile is accessible, the user receives a warning message and is logged on to the local system using the system's cached profile.



**NOTE** When a user has a mandatory profile, a temporary profile (meaning they are logged on as a guest), or a system profile, Windows blocks deployment of app packages by default. To permit app deployment when using one of these special profiles, you can enable Allow Deployment Operation In Special Profiles in the Administrative Templates policies for Computer Configuration under the Windows Components\App Package Deployment.






## Restricting Roaming Profiles

Typically, users can access their roaming profiles no matter which computer they're using within the domain. Windows allows you to modify this behavior by specifying from which computers a user can access roaming profiles and redirected folders. You do this by designating certain computers as primary computers and then configuring domain policy to restrict the downloading of profiles, redirected folders, or both to primary computers.


A primary computer is a computer that has been specifically designated as permitted for use with redirected data. You designate a primary computer by editing the advanced properties of a user or group in Active Directory and then setting the msDS-PrimaryComputer property to the name of the permitted computers. You then turn on the primary computer restriction for roaming profiles by enabling the Download Roaming Profiles On Primary Computers Only policy found in the Administrative Templates policies for Computer Configuration under System\User Profiles. You also can turn on the primary computer restriction for redirected folders by enabling the Redirect Folders On Primary Computers Only policy found in the Administrative Templates policies for Computer Configuration under System\Folder Redirection.

The goal of these policies is to protect personal and corporate data when users log on to computers other than the ones they use regularly for business. Data security is improved by not downloading and caching this data on computers a user doesn't usually use. To set the msDS-PrimaryComputer of a user or group, follow these steps:

1. In Active Directory Administrative Center, open the Properties dialog box for the user or group and then click Extensions. Or in Active Directory Users And Computers, ensure Advanced Features is selected on the View menu, and then open the Properties dialog box for the user or group.
2. On the Attribute Editor tab, scroll through the list of attributes. Click msDSPrimaryComputer and then click Edit.




3. In the Multi-Valued String Editor dialog box, enter the name of the first primary computer and then click Add. Repeat this process until you've added all primary computers. Click OK twice.




## Creating User Profiles

User profiles are maintained either in a default directory or in the location set by the Profile Path box in the user's Properties dialog box. For Windows 7 and later, the default location for profiles is %SystemDrive%\Users\%UserName%. A key part of the profile is the Ntuser.dat file in this location, such as C:\Users\wrstane\Ntuser.dat. To view this protected operating system file while working with File Explorer, you'll need to display hidden and protected operating system files by following these steps:

1. On the View pane in File Explorer, select Hidden Items and then click Options.



2. On the View tab in the Folder Options dialog box, clear Hide Protected Operating System Files and then click OK.



If you want users to have roaming profiles, you can assign profile paths in Active Directory Users And Computers or Active Directory Administrative Center. Otherwise, remove the profile path assignment when you want users have local profiles instead.

When users log on to multiple computers and use EFS, they need a roaming profile to ensure that the certificates necessary to read and work with encrypted files are available on computers other than their primary work computers. If you want a user to have a roaming profile, you must set a server-based location for the profile. One way to do this is to follow these steps:

1. Create a shared directory on a Windows server, and make sure that the group Everyone has at least Change and Read access.
2. In Active Directory Users And Computers or Active Directory Administrative Center, open the user's Properties dialog box, and then access the Profile tab/panel. Enter the path to the shared directory in the Profile Path box. The path should have the form \\server name\profile folder name\user name. An example is \\FileServer71\UserProfiles\TinaH, where FileServer71 is the server name, UserProfiles is the shared directory, and TinaH is the user name. You also can use the %UserName% environment variable as discussed previously.
3. The roaming profile will then be stored in the Ntuser.dat file in the designated directory, such as \\FileServer71\UserProfiles\TinaH\Ntuser.dat. Generally, you don't need to create the profile directory. The directory is created automatically the next time the user logs on, and NTFS permissions are set so that only the user has access.

**NOTE** When you are working with the administrator tools, you can select multiple user accounts for simultaneous editing. One way to do this is by holding down the Shift key or the Ctrl key when clicking the user names. Then when you right-click one of the selected users and then click Properties, you can edit properties for all the selected users. With multiple accounts, be sure to use

%UserName% in the profile path, such as \\FileServer71\UserProfiles\%UserName%.

4. As an optional step, you can create a profile for the user or copy an existing profile to the user's profile folder. If you don't create an actual profile for the user, the next time the user logs on, he will get a copy of the default local profile. Any changes the user makes to this profile are saved at log off. The next time the user logs on, he has a personal profile.

Mandatory profiles are also stored on Windows servers. If you want a user to have a mandatory profile, you define the profile as follows:


1. Create a shared directory on a server running Windows Server, and make sure that the group Everyone has at least Change and Read access.
2. In Active Directory Users And Computers or Active Directory Administrative Center, open the user's Properties dialog box, and then access the Profile tab/panel. Enter the path to the shared directory in the Profile Path box.
3. Log on as the user or have the user log on. This creates the user's profile in the specified profile path.
4. Have the user log off. Create the mandatory profile by renaming the Ntuser.dat file as **Ntuser.man**. The next time the user logs on, she will have a mandatory profile.

Ntuser.dat contains the registry settings for the user. When you change the extension for the file to Ntuser.man, you tell Windows Server the profile is a mandatory one.

## Managing Local Profiles

To manage local profiles, you need to log on to the user's computer. Then you can use the System utility in Control Panel to manage local profiles. To view current profile information, open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter. Next, in the System Properties dialog box, under User Profiles, click Settings.





The User Profiles dialog box displays information about the profiles stored on the local system. You can use this information to help you manage profiles. The dialog box lists the following information:

- **Name** The local profile's name, which generally includes the name of the originating domain or computer and the user account name. For example, the name `RPMEDIA\wrstane` tells you that the original profile is from the domain `rpmmedia` and the user account is `wrstane`.

**NOTE** If you delete an account but don't delete the associated profile, an entry that says `Account Deleted` or `Account Unknown` might be displayed. Don't worry—the profile is still available for copying if you need it, or you can delete the profile here.


- **Size** The profile's size. Generally, the larger the profile, the more the user has customized the environment.
- **Type** The profile type, which is either local or roaming.
- **Status** The profile's current status, such as whether it's from a local cache.
- **Modified** The date that the profile was last modified.



Sometimes you might want to create the profile manually. You do this by logging on to the user account, setting up the environment, and then logging off. As you might guess, creating accounts in this manner is time consuming. A better way to handle account creation is to create a base user account, set up the account environment, and then use this account as the basis of other accounts.

If you have a base user account or a user account you want to use in a similar manner, you can copy an existing profile to the new user account by following these steps:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter. (Or, in Control Panel, click System And Security, click System, and then click Advanced System Settings.)
2. In the System Properties dialog box, under User Profiles, click Settings.
3. Select the profile you want to copy from the Profiles Stored On This Computer list.
4. Copy the profile to the new user's account by clicking Copy To. In the Copy Profile To box, enter the path to the new user's profile directory. For example, if you were creating the profile for TinaH, you'd enter **\\FileServer71\UserProfiles\TinaH**.



5. Now you need to give the user permission to access the profile. In the Permitted To Use area, click Change, and then use the Select User Or Group dialog box to grant access to the new user account.

6. Click OK to close the Copy To dialog box. Windows then copies the profile to the new location.

When you work with workgroups where each computer is managed separately, you often have to copy a user's local profile from one computer to another. Copying a profile allows users to start with similar environment settings when they use different computers. Of course, in a Windows Server domain, you can use a roaming profile to create a single profile that can be accessed from anywhere within the domain. The problem is that sometimes you might need to copy an existing local profile to replace a user's roaming profile (when the roaming profile becomes corrupt), or you might need to copy an existing local profile to create a roaming profile in another domain.

You can copy a profile to a new location by following these steps:

1. Log on to the user's computer. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter. (Or, in Control Panel, click System And Security, click System, and then click Advanced System Settings.)
2. In the System Properties dialog box, under User Profiles, click Settings.
3. In the Profiles Stored On This Computer list, select the profile you want to copy.
4. Copy the profile to the new location by clicking Copy To, and then enter the path to the new profile directory in the Copy Profile To box. For example, if you're creating the profile for janew, you could enter **\\FileServer42\UserProfiles\Janew** .
5. To give the user permission to access the profile, click the Change button in the Permitted To Use area, and then grant access to the appropriate user account in the Select User Or Group dialog box.
6. When you have finished, click OK to close the Copy To dialog box. Windows then copies the profile to the new location.

## Changing the Profile Type

With roaming profiles, the System utility lets you change the profile type on the user's computer. To do this, select the profile and then click Change Type. The options in this dialog box allow you to do the following:

- **Change a roaming profile to a local profile** If you want the user to always work with the local profile on this computer, specify that the profile is for local use. All changes to the profile are then made locally, and the original roaming profile is left untouched.



- **Change a local profile (defined originally as a roaming profile) to a roaming profile** The user will use the original roaming profile for the next logon. Windows Server then treats the profile like any other roaming profile, which means that any changes to the local profile are copied to the roaming profile.

NOTE If these options aren't available, the user's original profile is defined locally.



## Recovering from Profile Problems

Windows Server uses local profiles for all users who don't have roaming profiles. Generally, local profiles are also used if the local profile has a more recent modification date than the user's roaming profile. Therefore, sometimes you might need to delete a user's local profile to ensure the roaming profile is used instead. For example, if a user's local profile becomes corrupt, you could delete the local profile on the affected computer and then the next time the user logs on to that computer the roaming profile will be downloaded to the computer and used instead. Keep in mind that when you delete a local profile that isn't stored anywhere else on the domain, you can't recover the user's original environment settings.

To delete a user's local profile, follow these steps:

1. Log on to the user's computer by using an account with administrator privileges. Open the System Properties dialog box with the Advanced tab selected by typing `SystemPropertiesAdvanced` in the Search box, and then pressing Enter. (Or, in Control Panel, click System And Security, click System, and then click Advanced System Settings.)
2. In the System Properties dialog box, under User Profiles, click Settings.
3. Select the profile you want to delete, and then click Delete. When asked to confirm that you want to delete the profile, click Yes.

**NOTE** You can't delete a profile that's in use. If the user is logged on to the local system (the computer you're deleting the profile from), the user needs to log off before you can delete the profile. In some instances, Windows Server marks profiles as in use when they aren't. This typically results from an environment change for the user that wasn't properly applied. To correct this, you might need to restart the computer.


The next time the user logs on, Windows Server does one of two things. If a roaming profile has been assigned and is accessible, the operating system retrieves it. Otherwise, the operating system gives the user the default local profile for that system. In the event that no available profile is working correctly for the user, you can remove the local profile as discussed and then use File Explorer to remove the roaming profile from its storage location on Windows Server.

# Maintaining Users and Groups

Active Directory Administrative Center and Active Directory Users And Computers are the tools to use when you want to update a domain user or group account. If you want to update a local user or group account, use Local Users And Groups.

When you work with Active Directory, you'll often want to get a list of accounts and then do something with those accounts. For example, you might want to list all the user accounts in the organization and then disable the accounts of users who have left the company. One way to perform this task is to follow these steps:

1. In Active Directory Users And Computers, right-click the domain name and then click Find.
2. In the Find list, select Custom Search. This updates the Find dialog box to display a Custom Search tab.
3. Using the In list, select the domain or other area you want to search. To search the enterprise, select Entire Directory.
4. Click Field to display a menu. Select User, and then select Logon Name (Pre–Windows 2000).



**TIP** Be sure to select Logon Name (Pre–Windows 2000). Don't use Logon Name. User accounts aren't required to have a logon name, but they are required to have a pre–Windows 2000 logon name.

5. In the Condition list, select Present, and then click Add.
6. Click Find Now. Active Directory Users And Computers gathers a list of all users in the designated area.

- 
7. You can now work with the accounts one by one or several at a time. One way to select multiple resources not in sequence is to hold down the Ctrl key and then click each object you want to select. One way to select a series of resources at one time is to hold down the Shift key, click the first object, and then click the last object.
  8. Right-click a user account, and then select an action on the shortcut menu that's displayed, such as **Disable Account**.

**TIP** The actions you can perform on multiple accounts include **Add To Group** (used to add the selected accounts to a designated group), **Enable Account**, **Disable Account**, **Delete**, **Move**, and **Send Mail**. By choosing **Properties**, you can edit the properties of multiple accounts.

Use this same procedure to get a list of computers, groups, or other Active Directory resources. With computers, do a custom search, click **Field**, choose **Computer**, and then select **Computer Name (Pre–Windows 2000)**. With groups, do a custom search, click **Field**, choose **Group**, and then select **Group Name (Pre–Windows 2000)**.

The sections that follow examine other techniques you can use to update (rename, copy, delete, and enable) accounts and change and reset passwords. You'll also learn how to troubleshoot account logon problems.

# Copying Domain User Accounts

Creating domain user accounts from scratch can be tedious. Instead of starting anew each time, you might want to use an existing account as a starting point. This option currently isn't in Active Directory Administrative Center. To do this in Active Directory Users And Computers, follow these steps:

1. Right-click the account you want to copy, and then click Copy. This opens the Copy Object—User dialog box.
2. Create the account as you would any other domain user account, and then update the properties of the account as appropriate.

Copy Object - User

Create in: imaginedlands.local/Users

First name:  Initials:

Last name:

Full name:

User logon name:  @imaginedlands.local

User logon name (pre-Windows 2000):

< Back Next > Cancel

As you might expect, when you create a copy of an account, Active Directory Users And Computers doesn't retain all the information from the existing account and instead tries to copy only the information you need while discarding the information that you need to update. The following properties are retained:

- City, state, postal code, and country values set on the Address tab
- Department and company set on the Organization tab
- Account options set by using the Account Options boxes on the Account tab
- Logon hours and permitted logon workstations
- Account expiration date
- Group account memberships
- Profile settings



## Dial-in privileges

If you used environment variables to specify the profile settings in the original account, the environment variables are also used for the copy of the account. For example, if the original account used the %UserName% variable, the copy of the account will also use this variable.





# Importing and Exporting Accounts

Windows Server 2016 includes the Comma-Separated Value Directory Exchange (CSVDE) command-line utility for importing and exporting Active Directory objects. For import operations, CSVDE uses a comma-delimited file as the import source. You can run CSVDE by using these general parameters:

- **-i** Turns on import mode (rather than export, which is the default mode)
- **-f filename** Sets the source for an import or the output file for an export
- **-s servername** Sets the server to use for the import or export (rather than the default domain controller for the domain)
- **-v** Turns on verbose mode

For import operations, the source file's first row defines the list of Lightweight Directory Access Protocol (LDAP) attributes for each object defined. Each successive line of data provides the details for a specific object to import and must contain exactly the attributes listed. Here is an example:

```
DN,objectClass,sAMAccountName,sn,givenName,userPrincipalName
"CN=William Stanek,OU=Eng,DC=tvpress,DC=com",user,williams,
William,Stanek,williams@tvpress.com
```

Given this listing, if the import source file is named newusers.csv, you could import the file into Active Directory by entering the following command at an elevated command prompt:

```
csvde -i -f newusers.csv
```

For export operations, CSVDE writes the exported objects to a comma-delimited text file. You can also run CSVDE by using the general parameters listed previously as export-specific parameters, which include the following:

- **-d RootDN** Sets the starting point for the export, such as **-d "OU=Sales,DC=domain,DC=local"**. The default is the current naming context.
- **-l list** Provides a comma-separated list of attributes to output.
- **-r Filter** Sets the LDAP search filter, such as **-r "(objectClass=user)"**.
- **-m** Configures output for the Security Accounts Manager (SAM) rather than Active Directory.

To create an export file for the current naming context (the default domain), you could enter the following at an elevated command prompt:

```
csvde -f newusers.csv
```

However, this could result in a very large export dump. Thus, in most cases, you should specify at a minimum the RootDN and an object filter, as shown here:

```
csvde -f newusers.csv -d "OU=Service,DC=tvpress,DC=com" -r  
"(objectClass=user)"
```



## Renaming Users and Groups

When you rename a user account, you give the account a new label. As discussed in Chapter 9, “Creating Accounts,” user names are meant to make managing and using accounts easier. Behind the scenes, Windows Server uses security identifiers (SIDs) to identify, track, and handle accounts independently from user names. SIDs are unique identifiers that are generated when accounts are created.

Because SIDs are mapped to account names internally, you don’t need to change the privileges or permissions on renamed accounts. Windows Server simply maps the SIDs to the new account names as necessary.

One common reason for changing the name of a user account is that the user gets married and decides to change her last name. For example, if Heidi Steen (heidis) gets married, she might want her user name to be changed to Heidi Jensen (heidij). When you change the user name from heidis to heidij, all associated privileges and permissions will reflect the name change. If you view the permissions on a file that heidis had access to, heidij now has access (and heidis is no longer listed).

To simplify the process of renaming user accounts, Active Directory Users And Computers provides a Rename User dialog box you can use to rename a user’s account and all the related name components. This dialog box currently isn’t in Active Directory Administrative Center, so you need to open the Properties dialog box and enter the new name properties for each box as appropriate.

To rename an account, follow these steps:

1. Find the user account you want to rename in Active Directory Users And Computers.
2. Right-click the user account, and then click Rename. Active Directory Users And Computers highlights the account name for editing. Press Backspace or Delete to erase the existing name, and then press Enter to open the Rename User dialog box.



Rename User

Full name: Tina Stevens

First name: Tina

Last name: Stevens

Display name: Tina Stevens

User logon name:  @imaginedlands.local

User logon name (pre-Windows 2000):

OK Cancel

3. Make the necessary changes to the user's name information, and then click OK. If the user is logged on, a warning prompt is displayed telling you that the user should log off and then log back on by using the new account logon name.
4. The account is renamed, and the SID for access permissions remains the same. You might still need to modify other data for the user in the account Properties dialog box, including the following:
  - **User Profile Path** Change the Profile Path in Active Directory Users And Computers, and then rename the corresponding directory on disk.
  - **Logon Script Name** If you use individual logon scripts for each user, change the Logon Script Name in Active Directory Users And Computers, and then rename the logon script on disk.
  - **Home Directory** Change the home directory path in Active Directory Users And Computers, and then rename the corresponding directory on disk.

**NOTE** Changing directory and file information for an account when a user is logged on might cause problems. You might want to update this information after hours or ask the user to log off for a few minutes and then log back on. You can usually write a simple Windows script that can perform the tasks for you automatically.

## Managing Multiple Accounts

You can use Active Directory Users And Computers to modify the properties of multiple accounts simultaneously. Any changes you make to the property settings are applied to all the selected accounts. When you right-click the selected accounts, the following options are available:

- **Add To A Group** Displays the Select Group dialog box, which you can use to designate the groups the selected users should be members of
- **Disable Account** Disables all the selected accounts
- **Enable Account** Enables all the selected accounts
- **Move** Moves the selected accounts to a new container or organizational unit
- **Cut** Moves the selected accounts to a new container or organizational unit when you later select Paste
- **Delete** Deletes the selected accounts from the directory
- **Properties** Allows you to configure a limited set of properties for multiple accounts

In Active Directory Administrative Center, the options are similar. Add To Group, Disable All, Enable All, Unlock All, Move, Delete, and Properties are displayed. Select Properties to open the Properties For Multiple Items dialog box, which has a different interface from the Properties dialog box for standard users. When working with properties for multiple accounts, you should note the following:

- Account name and password boxes are no longer available. You can, however, set the Domain Name System (DNS) domain name (user principal name [UPN] suffix), logon hours, computer restrictions, account options, account expiration, and profiles. • You must specifically select properties you want to work with by selecting the related properties' check boxes. After you do this, the value you enter in the box is applied to all the selected accounts.

One of the best reasons to edit multiple accounts is that you can set the environment profiles for all of the accounts by using a single interface. To do this, you usually rely on the %UserName% environment variable, which lets you assign paths and file names that are based on individual user names. For example, if you assign the logon script name as %UserName%.cmd, Windows replaces this value with the user name, and it does so for each user you're managing.



Properties for Multiple Items

General Account Address Profile Organization

 Multiple users selected

To change a property for multiple objects, first select the checkbox to enable the change, and then type the change.

Depending on the number of objects selected, you might have to wait while the changes are applied.

☐ Description:

☐ Office:


☐ Telephone number:

☐ Fax:

☐ Web page:

☐ E-mail:

OK Cancel Apply



Properties for Multiple Items

General Account Address Profile Organization

To change a property for multiple objects, first select the checkbox to enable the change, and then type the change.

User profile

☒ Profile path:

☐ Logon script:


☒ Home folder

☐ Local path:

☒ Connect: Z: To:

Although you might want all users to have unique file names and paths, sometimes you want users to share this information. For example, if you're using mandatory profiles for users, you might want to assign a specific user profile path rather than one that's dynamically created.

Select the Account tab if you want to manage restrictions collectively for log on hours and log on workstations. However, Active Directory Users And Computers doesn't tell you the previous logon hour designations for the selected accounts, and it doesn't warn you if the logon hours for the accounts are different.



User accounts have many options that control logon, passwords, and account expiration. You set these values on the Account tab as well. When you work with multiple accounts, you must enable the option you want to work with by selecting the corresponding check box in the leftmost column. You now have two choices:

- Enable the option by selecting its check box. For example, if you are working with the Password Never Expires option, a flag is set so that the password for the selected users won't expire when you click OK.
- Don't set the option, which effectively clears the option. For example, if you are working with the Account Is Disabled option, the accounts for the selected users are reenabled when you click OK.



If you want to set the expiration date of the selected accounts, start by selecting Account Expires, and then select the appropriate expiration value. The Never option removes any current account expiration values. Select the End Of option to set a specific expiration date.

# Changing and Resetting Passwords

As an administrator, you often have to change or reset user passwords. This usually happens when users forget their passwords or when their passwords expire.

To change or reset a password, follow these steps:

1. Open Active Directory Users And Computers, Active Directory Administrative Center, or Local Users And Groups (whichever is appropriate).
2. Right-click the account name, and then click Reset Password or Set Password.
3. Enter a new password for the user and confirm it. The password should conform to the password-complexity policy set for the computer or domain.



4. User Must Change Password At Next Logon forces the user to change his password when the user logs on next. If you don't want the user to have to change his password, clear this check box.
5. The Account Lockout Status On This Domain Controller property shows whether the account is locked or unlocked. If the account is locked, select Unlock The User's Account to unlock it. Click OK.

## Enabling User Accounts

User accounts can become disabled for several reasons. If a user forgets her password and tries to guess it, the user might exceed the account policy for bad logon attempts. Another administrator could have disabled the account while the user was on vacation, or the account could have expired.

Active Directory Users And Computers and Active Directory Administrative Center depict disabled accounts with an arrow next to the user icon in the main view. When an account is disabled, follow these steps to enable it:





1. Open Active Directory Users And Computers, Active Directory Administrative Center, or Local Users And Groups (whichever is appropriate).
2. Right-click the user's account name. Select the appropriate option for the tool you are using, either Enable or Enable Account.

**TIP** To quickly search the current domain for disabled accounts, enter **dsquery user -disabled** at a command prompt.

You can select multiple accounts at the same time and then use the options on the shortcut menu to enable or disable them. In Active Directory Users And Computers, enable all selected accounts by using the Enable Account option, or disable them by selecting Disable Account. In Active Directory Administrative Center, enable the accounts by using the Enable All option, or disable them by using Disable All.

When an account is locked out, follow these steps to unlock it:

1. Open Active Directory Users And Computers, Active Directory Administrative Center, or Local Users And Groups (whichever is appropriate).
2. Double-click the user's account name, and then select the Unlock Account check box. In Active Directory Users And Computers, this check box is on the Account tab.

The screenshot shows the 'Account' tab of a user's properties in Active Directory. The 'User logon name' is 'edg' and the domain is '@imaginedlands.local'. The 'User logon name (pre-Windows 2000)' is 'IMAGINEDLANDS\edg'. There are buttons for 'Logon Hours...' and 'Log On To...'. At the bottom, the 'Unlock account' checkbox is checked.

In Active Directory Administrative Center, you can unlock multiple accounts at the same time. Simply select the accounts and then use the Unlock All option on the shortcut menu to unlock the accounts.

**NOTE** If users frequently get locked out of their accounts, consider adjusting the account policy for the domain. You might want to increase the value for acceptable bad logon attempts and reduce the duration for the associated counter. For more information about setting account policy, see Chapter 8 “Managing Account Policies.”



Only domain accounts have an expiration date. (Local user accounts don't have an expiration date.) When a domain account expires, follow these steps to change the expiration date:

1. Open Active Directory Users And Computers or Active Directory Administrative Center.
  2. Double-click the user's account name. Open the Account tab or panel.
  3. Under Account Expires, do one of the following:
- Select End Of, and then click the arrow on the related list box. With Active Directory Users And Computers, this displays a calendar you can use to set a new expiration date. With Active Directory Administrative Center, enter the date in the format shown.



Select Never to specify that the account no longer has a predefined expiration date.

## Setting Active Directory Permissions

Users, groups, computers and all other objects in Active Directory have standard and advanced security permissions. These permissions grant or deny access to the objects.


Permissions for Active Directory objects aren't as straightforward as other permissions. Different types of objects can have sets of permissions that are specific to the type of object. They can also have general permissions that are specific to the container they're defined in.

You can view and set standard security permissions for objects by following these steps:

1. Start Active Directory Users And Computers, and then display advanced options by choosing Advanced Features from the View menu. Next, right-click the user, group, or computer account you want to work with, and then click Properties.



2. In the Properties dialog box, click the Security tab. A list of groups and users that have been assigned permissions on the object you previously selected is displayed. If the permissions are unavailable, it means the permissions are inherited from a parent object.
3. Users or groups with access permissions are listed in the Group Or User Names list box. You can change permissions for these users and groups by doing the following:
  - Select the user or group you want to change.
  - Grant or deny access permissions in the Permissions list.
  - When inherited permissions are not available, override inherited permissions by selecting the opposite permissions.
4. To set access permissions for additional users, computers, or groups, click Add. In the Select Users, Computers, Service Accounts, Or Groups dialog box, add users, computers, or groups.



5. In the Group Or User Names list, select the user, computer, or group you want to configure. In the check boxes in the Permissions area, allow or deny permissions. Repeat this step for other users, computers, or groups.
6. Click OK when you have finished.

**CAUTION** Only administrators who have a solid understanding of Active Directory and Active Directory permissions should manipulate object





permissions. Setting object permissions incorrectly can cause problems that are very difficult to track down.

One way to view and set advanced security permissions for objects is by following these steps:


1. Start Active Directory Users And Computers, and then display advanced options by choosing Advanced Features from the View menu. Next, right-click the user, group, or computer account you want to work with, and then click Properties.
2. In the Properties dialog box, click the Security tab, and then click Advanced. A list of individual permission entries for the previously selected object is displayed. Permission entries that are inherited are listed as being inherited from a specific parent object.
3. To view and set the individual permissions associated with a permission entry, select the entry, and then click Edit. You can change advanced permissions for the selected user or group by granting or denying access permissions in the Permissions list. When inherited permissions are not available, override inherited permissions by selecting the opposite permissions.
4. Click OK twice when you have finished.



## Deleting Users and Groups

Deleting an account permanently removes the account. After you delete an account, you can't create an account with the same name to restore the same permissions. That's because the SID for the new account won't match the SID for the old account.

Because deleting built-in accounts can have far-reaching effects on the domain, Windows Server 2016 doesn't let you delete built-in user accounts or group accounts. You can remove other types of accounts by selecting them and pressing the Delete key, or by right-clicking and selecting Delete. When prompted, click Yes.



With Active Directory Users And Computers, one way to work with multiple accounts is by doing one of the following:

- Select multiple user names for editing by holding down the Ctrl key and clicking each account you want to select.
- Select a range of user names by holding down the Shift key, selecting the first account name, and then clicking the last account in the range.

**NOTE** When you delete a user account, Windows Server 2016 doesn't delete the user's profile, personal files, or home directory. If you want to delete these files and directories, you have to do it manually. If this is a task you perform routinely, you might want to create a script that performs the necessary procedures for you. However, don't forget to back up files or data that might be needed before you do this.

## Troubleshooting Logon Problems

The previous section listed ways in which accounts can become disabled. Active Directory Users And Computers shows disabled accounts with a red warning icon next to the account name. To enable a disabled account, right-click the account in Active Directory Users And Computers, and then click Enable Account.

You can also search the entire domain for users with disabled accounts by entering **dsquery user –disabled** at a command prompt. To enable a disabled account from the command line, enter **dsmod user UserDN–disabled no** .

When a user account has been locked out by the Account Lockout policy, the account cannot be used for logging in until the lockout duration has elapsed or an administrator resets the account. If the account lockout duration is indefinite, the only way to unlock the account is to have an administrator reset it as discussed previously.

Windows Server 2016 can record logon success and failure through auditing. When you enable account logon failure auditing, logon failure is recorded in the security log on the login domain controller. Auditing policies for a site, domain, or organizational unit GPO are found under Windows Settings\Security Settings under Computer Configuration\Local Policies\Audit Policy.

When a user logs on to the network by using his domain user account, the account credentials are validated by a domain controller. By default, users can log on by using their domain user accounts even if the network connection is down or no domain controller is available to authenticate the user's logon.

The user must have previously logged on to the computer and have valid, cached credentials. If the user has no cached credentials on the computer, and the network connection is down or no domain controller is available, the user will not be able to log on. Each member computer in a domain can cache up to 10 credentials by default.

Authentication can also fail if the system time on the member computer deviates from the logon domain controller's system time by more than is allowed in the Kerberos Policy: Maximum Tolerance For Computer Clock Synchronization. The default tolerance is 5 minutes for member computers.

Beyond these typical reasons for an account being disabled, some system settings can also cause access problems. Specifically, you should look for the following:

- **A user gets a message that says that the user can't log on interactively** The user right to log on locally isn't set for this user, and the user isn't a member of a group that has this right.

The user might be trying to log on to a server or domain controller. If so, keep in mind that the right to log on locally applies to all domain controllers in the domain. Otherwise, this right applies only to the single workstation.

If the user is supposed to have access to the local system, configure the Logon Locally user right as described in the "Configuring User Rights Policies" section in Chapter 8.

- **A user gets a message that the system could not log on the user** If you've already checked the password and account name, you might want to check the account type. The user might be trying to access the domain with a local account. If this isn't the problem, the global catalog server might be unavailable, which means that only users with administrator privileges can log on to the domain.
- **A user has a mandatory profile and the computer storing the profile is unavailable** When a user has a mandatory profile, the computer storing the profile must be accessible during the logon process. If the computer is shut down or otherwise unavailable, users with mandatory profiles might not be able to log on. See the "Using Roaming and Mandatory Profiles" section earlier in this chapter.
- **A user gets a message saying the account has been configured to prevent the user from logging on to the workstation** The user is trying to access a workstation that isn't defined as a permitted logon workstation. If the user is supposed to have access to this workstation, change the logon workstation information as described in the "Restricting Workstation Logons" section earlier in the chapter.

## Chapter 12. Maintaining Your Servers

In this chapter, you'll learn how to use the standard options and tools in Windows Server 2016 to maintain your servers. Ensuring servers you are responsible for are running and operational is essential. In IT, everything tends to run smoothly until it doesn't and then it's either all downhill from there because you are poorly prepared or a steady march to resolution because you are ready and prepared to do what needs to be done.

The status and usage of system resources can change dramatically over time. Things can and do go wrong. If you are prepared, however, you don't need to worry. Windows Server provides everything you need to weather a storm, whether the storm is caused by services that have stopped, file systems that have run out of space, applications that have thrown exceptions or some other wild card that's shown up and decided to make your day a difficult one.





# Working with Services

When you install roles and features on a server, those roles and features often operate as services. Services are facilitators, providing the key functions to workstations and servers.

All services have a primary executable, and can be started and stopped, paused and resumed. Some services have dependencies. These are either components the service depends on, components that depend on the service, or both.

Services can be configured to started automatically with the operating system, automatically with a delay, or manually. Some manually started services can be triggered to start by the startup of other services.






# Navigating Services in Server Manager

When you are working with Server Manager and select the Local Server node, the All Servers node, or a server group node, the right pane will have a Services panel. If you select the server you want to work with in the Servers panel, its services are listed in the Services panel. You can use this panel as follows:

- For a server you are logged on to locally, you can use the Services panel in the Local Server node.
- For a local or remote server, you can use the Services panel in the All Servers node to work with services.
- Automatically created server group nodes are organized by server roles, such as Active Directory Domain Services (AD DS) or Domain Name System (DNS), and you'll be able to manage the services running on servers that role depends on.
- For custom server groups created by you or other administrators, you'll be able to use the related Services panel to manage services on any remote servers that have been added to the group.



The columns on the Services panel can be adjusted by right-clicking any column header and then clicking the columns to add or remove. The columns you can use include the following:

- **Server Name** The name of the server on which the service is running.
- **FQDN** The fully qualified domain name (FQDN) of the server on which the service is running.
- **Display Name** The descriptive name of the service.
- **Service Name** The internal name of the service.
- **Description** A short description of the service and its purpose.
- **Status** Whether the status of the service is running, paused, or stopped.
- **Type** The startup setting for the service. Automatic services are started at



system start up. Users or other services start manual services. Disabled services are turned off and can't be started while they remain disabled.

When you are working with many servers, use the service grouping options to help you more easily manage services. You can group services by server name, FQDN, display name, service name, status, and start type by right-clicking any column header, selecting Group By, and then clicking your grouping option.



# Navigating Services in Computer Management


For quick and easy management of any service on a remote server, you can use the Services node in Computer Management. You can open Computer Management and automatically connect to a remote server from Server Manager. To do this, follow these steps:

1. Select All Servers or any server group node in the left pane.
2. On the Servers panel, right-click the server to which you want to connect.
3. Click Computer Management.

**TIP** If the user account you are currently using doesn't have the appropriate credentials for working with the remote server, you won't be able to connect to the server in Computer Management. To use alternate credentials, right-click the server to which you want to connect, select **Manage As**, enter your alternate credentials, and then click **OK**. Optionally, you can select **Remember My Credentials** before clicking **OK** to save the credentials for each time you log on and want to work with the server remotely. After you set your credentials, rightclick the server to which you want to connect and then click **Computer Management**. Now Computer Management will open and connect to the server using the credentials you specified.

When you are working with Computer Management, you view and work with services by expanding the **Services And Applications** node and then selecting the **Services** node. The columns in the Services pane are slightly different from those shown when you are working with the Services node in Computer Management:

- **Name** The name of the service. Only services installed on the system are listed here. Double-click an entry to configure its startup options. If a service you need isn't listed, you can install it by installing the related role or feature, as discussed in "Managing Roles and Features" in Chapter 2.
- **Description** A short description of the service and its purpose.
- **Status** The status of the service, indicated as running, paused, or stopped. (Stopped is indicated by a blank entry.)
- **Startup Type** The startup setting for the service. Automatic services are started at system start up. Users or other services start manual services. Disabled services are turned off and can't be started while they remain disabled.
- **Log On As** The account the service logs on as. The default in most cases is the local system account.



The Services pane has two views: Extended and Standard. To change the view, use the tabs at the bottom of the Services pane. In Extended view, quick links are provided for managing services. Click Start to start a stopped service. Click Restart to stop and then start a service—essentially resetting that service. If you select a service when the Services pane is in Extended view, a description that details the service’s purpose is displayed.

**NOTE** Both the operating system and a user can disable services. Generally, Windows Server 2016 disables a service if a possible conflict with another service exists.

## Tracking Service Status

Many problems with servers originate from services that are configured to start automatically but don’t or shut down during operations due to some issue with the operating system.. If a service has an automatic startup type but isn’t started, you will need to start the service manually.

When services that are set to start automatically fail, the status is listed as blank rather than the operational state, such as Running. Service failures also may be logged to the system’s event logs.

To start, stop, or pause a service, right-click the service you want to manage and then select Start, Stop, or Pause as appropriate. You can also choose Restart to have Windows stop and then start the service after a brief pause. Restarting a service that supports a role or function can sometimes resolve issues the server is experiencing with the role or function. Additionally, if you pause a service, you can use the Resume option to resume normal operation.

NOTE You can configure actions to handle service failure automatically. For example, you can have Windows Server 2016 attempt to restart the service for you. For details, see the section “Configuring Service Recovery” later in this chapter.

When you are having problems with a server it’s helpful to know the typical status of services on the server and be able to compare this to the current status of services. For this reason, I recommend creating a baseline record of service status for each of your servers that you can use for troubleshooting comparisons. Creating the baseline record and using it during troubleshooting are much easier than you might imagine.

To create a baseline record of service status, enter the following at a PowerShell prompt:

```
get-service > baseline-s.txt
```

This command stores the output of Get-Service in a file called Baseline-S.txt. If you open the file, you’ll see its contents look similar to the following:

| Status  | Name     | DisplayName                       |
|---------|----------|-----------------------------------|
| -----   | ----     | -----                             |
| Running | ADWS     | Active Directory Web Services     |
| Stopped | AJRouter | AllJoyn Router Service            |
| Stopped | ALG      | Application Layer Gateway Service |
| Stopped | AppIDSvc | Application Identity              |
| Running | Appinfo  | Application Information           |
| Stopped | AppMgmt  | Application Management            |

As you can see, each service running on the server is listed alphabetically with its current status. Once you create a baseline, you have it available for troubleshooting. The basic steps for diagnosing a problem with services follow:

1. Start your troubleshooting by creating a new status record for system services. To do this, enter the following at a PowerShell prompt:

```
get-service > current-s.txt
```

2. Compare the original baseline record with the current record by entering the following at a PowerShell prompt:

```
compare-object (get-content baseline-s.txt) (get-content current-s.txt)
```

3. The output of the compare-object command shows the differences between the baseline record and the current record. If there is no output, there are no changes in status. Otherwise, output marked with the sideindicator => is from the current record and output marked with the sideindicator <= is from the baseline.



Consider the following sample output from a comparison:

| InputObject            | SideIndicator                        |
|------------------------|--------------------------------------|
| -----                  | -----                                |
| Running AppXSvc        | AppX Deployment Service (AppXSVC) => |
| Running ClipSVC        | Client License Service (ClipSVC) =>  |
| Running LicenseManager | Windows License Manager Service =>   |
| Stopped MpsSvc         | Windows Firewall =>                  |
| Running Netman         | Network Connections =>               |
| Running NetSetupSvc    | Network Setup Service =>             |
| Stopped NgcCtnrSvc     | Microsoft Passport Container =>      |
|                        |                                      |
| Stopped AppXSvc        | AppX Deployment Service (AppXSVC) <= |
| Stopped ClipSVC        | Client License Service (ClipSVC) <=  |
| Stopped LicenseManager | Windows License Manager Service <=   |
| Running MpsSvc         | Windows Firewall <=                  |
| Stopped Netman         | Network Connections <=               |
| Stopped NetSetupSvc    | Network Setup Service <=             |
| Running NgcCtnrSvc     | Microsoft Passport Container <=      |

Here, there are a number of changes between the current status report and the original baseline. Only two services have stopped running, however. Windows Firewall and Microsoft Passport Container. In the original baseline, both services were running.

As services that were running but are stopped now can be an indicator of a problem, you can zero in on the stopped services as part of your troubleshooting. Check the Startup Type of the stopped services. If you determine that one or more of the services should have started automatically, you should start the services manually.


## Configuring Service Startup

You can set services to start manually or automatically. You can also turn them off permanently by disabling them. You configure service startup in Computer Management by following these steps:

1. Right-click the service you want to configure, and then choose Properties.
  2. On the General tab, use the Startup Type list to choose a startup option from the following choices and then click OK:
- **Automatic** Select Automatic to start services at bootup.
  - **Automatic (Delayed Start)** Select Automatic (Delayed Start) to delay the start of the service until all nondelayed automatic services have started.



- **Manual** Select Manual to allow the services to be started manually.
- **Disabled** Select Disabled to turn off the service.



## Configuring Service Logon

You can configure services to log on as a system account or as a specific user. To do either, follow these steps:

1. In Computer Management, right-click the service you want to configure, and then click Properties.
2. Click the Log On tab.





3. Click Local System Account if you want the service to log on by using the system account (the default for most services). If the service provides a user interface that can be manipulated, select Allow Service To Interact With Desktop to allow users to control the service's interface.
4. Click This Account if you want the service to log on by using a specific user account. Be sure to enter an account name and password in the text boxes provided. Use the Browse button to search for a user account if necessary.
5. Click OK.


**SECURITY ALERT** You should keep track of any accounts that are used with services. These accounts can be the source of security problems if they're not configured properly. Service accounts should have the strictest security settings and as few permissions as possible while allowing the service to perform necessary functions. Typically, accounts used with services don't need many of the permissions you would assign to a typical user account. For example, most service accounts don't need the right to log on locally. Every administrator should know what service accounts are used (so that they can better track the use of these accounts) and should treat the accounts as if they were administrator accounts. This means using secure passwords, carefully monitoring account usage, carefully applying account permissions and privileges, and so on.

## Configuring Service Recovery

You can configure services to take specific actions when a service fails. For example, you can attempt to restart the service or run an application. To configure recovery options for a service, follow these steps:

1. In Computer Management, right-click the service you want to configure, and then click Properties.
2. Click the Recovery tab.





**NOTE** Windows Server 2016 automatically configures recovery for critical system services during installation. In most cases, you'll find that critical services are configured to restart automatically if the service fails. Some extremely critical services are configured to restart the computer if the service fails. You cannot change these settings because they are not modifiable for critical services.

3. You can now configure recovery options for the first, second, and subsequent recovery attempts. The following options are available:
  - **Take No Action** The operating system won't attempt recovery for this failure but might still attempt recovery of previous or subsequent failures.
  - **Restart The Service** Stops and then starts the service after a brief pause.
  - **Run A Program** Allows you to run a program or a script in case of failure. The script can be a batch program or a Windows script. If you select this option, set the full file path to the program you want to run, and then set any necessary commandline parameters to pass in to the program when it starts.
  - **Restart The Computer** Shuts down and then restarts the computer. Before you choose this option, double-check the computer's Startup and Recovery options. You want the system to select defaults quickly and automatically.



**BEST PRACTICES** When you configure recovery options for critical services, you might want to try to restart the service on the first and second attempts and then restart the server on the third attempt.

4. Configure other options based on your previously selected recovery options. If you elected to run a program as a recovery option, you need to set options in the Run Program panel. If you elected to restart the service, you need to specify the restart delay. After stopping the service, Windows Server waits for the specified delay before trying to start the service. In most cases, a delay of 1 to 2 minutes should be sufficient.
5. Click OK.



## Disabling Unnecessary Services

As an administrator, you need to ensure that servers and the network are secure, and unnecessary services are a potential source of security problems. For example, in many organizations that I've reviewed for security problems, I've found servers running Worldwide Web Publishing Service, Simple Mail Transfer Protocol (SMTP), and File Transfer Protocol (FTP) Publishing Service when these services weren't needed. Unfortunately, these services can make it possible for anonymous users to access servers and can also open the server to attack if not properly configured.

If you find unnecessary services, you have several options. With services installed through roles, role services, or features, you can remove the related role, role service, or feature to remove the unnecessary component and its related services. Or you can simply disable the services that aren't being used. Typically, you'll want to start by disabling services rather than uninstalling components. This way, if you disable a service and another administrator or a user says he or she can't perform a particular task anymore, you can enable the related service again if necessary.

To disable a service, follow these steps:

1. In Computer Management, right-click the service you want to configure, and then click Properties. On the General tab, click Disabled in the Startup Type list.
2. Disabling a service doesn't stop a running service. It only prevents the service from being started the next time the computer is started, meaning that the security risk still exists. To address this, click Stop on the General tab in the Properties dialog box, and then click OK.

## Troubleshooting Services

When you think a server has a problem related to services—their status or configuration, for example—one of the fastest ways to troubleshoot is to use Windows PowerShell. Windows PowerShell provides several cmdlets for working with system services, including:

- **Get-Service** Gets information about the services.
- **Stop-Service** Stops one or more running services.
- **Start-Service** Starts one or more stopped services.
- **Suspend-Service** Suspends (pauses) one or more running services.
- **Resume-Service** Resumes one or more suspended (paused) services.
- **Restart-Service** Stops and then starts one or more services.



**Set-Service** Changes the properties or status of a service on a local or remote computer.

• **New-Service** Creates a new Windows service in the registry and in the services database.

To get a list of all services configured on a system, type **get-service** at the command prompt. Optionally, use the -ComputerName parameter to specify a remote computer to work with, as shown in this example.

```
Get-Service -ComputerName CorpServer18
```

If you want to check multiple computers, simply enter the computer name in a commaseparated list.

```
Get-Service -ComputerName Server18, Server24, Server30
```

Regardless of whether you check the local computer or remote computers, the standard output shows the status, name, and display name of each configured service, as shown in the following example.

| Status  | Name               | DisplayName                       |
|---------|--------------------|-----------------------------------|
| -----   | ----               | -----                             |
| Running | ADWS               | Active Directory Web Services     |
| Stopped | AeLookupSvc        | Application Experience            |
| Stopped | ALG                | Application Layer Gateway Service |
| Stopped | AllUserInstallA... | Windows All-User Install Agent    |
| Stopped | AppIDSvc           | Application Identity              |
| Running | Appinfo            | Application Information           |

Knowing this, you can display information about a specific service by using its name or display name, such as the following.

```
Get-Service -DisplayName "Application Experience" -ComputerName CorpServer18
```

To match partial names, you can use wildcard characters. In this example, you want to check the status of key services used by Hyper-V.

```
Get-Service -DisplayName Hyper* -ComputerName CorpServer18
```

As part of troubleshooting, you might also want to locate all services that are disabled or stopped on a computer. To find disabled services, you can enter the following.

```
Get-Service | where {$_.status -eq 'disabled'}
```

To find stopped services, you can enter the following.

```
Get-Service | where {$_.status -eq 'stopped'}
```

A service can start only if the services it depends on also are started. You can use the RequiredServices parameter to confirm the status of the services that are required for





a particular service to start. Alternatively, you can check the status of services that depend on a particular service by using the `-DependentServices` parameter.

As part of troubleshooting, you might need to determine the start mode of services. If a service isn't configured to start automatically, that would explain why it isn't running. Although `Get-Service` doesn't display the `StartMode` of services, you can use `Get-WMIObject Win32_Service` to get extended information about services, including their start mode.

The syntax for displaying detailed information about a particular service follows.

```
Get-WMIObject Win32_Service -Filter "name= 'ServiceName' " | Select *
```

Or you could use the following.

```
Get-WMIObject Win32_Service -Filter "displayname= 'DisplayName' " | Select *
```

In these examples, you use `Get-WMIObject Win32_Service` to examine all services on a computer and filter the results by either the name or display name of the service. You then use `Select-Object` (which has `Select` as an alias) to list all associated properties.

In your troubleshooting, you might want to list the start mode of all enabled services on a computer. The following example shows one way to do this.

```
get-wmiobject win32_service -filter "StartMode <>'disabled'" | sort  
StartMode | format-table -GroupBy StartMode -Property Name,State
```

Here, you list all enabled services on a computer, sorting the output by the start mode, and then alphabetically list services in each group by name and state. The output is similar to the following.

| StartMode: Auto   |         |
|-------------------|---------|
| Name              | State   |
| ----              | -----   |
| Netlogon          | Running |
| NlaSvc            | Running |
| MpsSvc            | Running |
| ...               |         |
| StartMode: Manual |         |
|                   |         |
| Name              | State   |



| ----               | -----   |
|--------------------|---------|
| AppReadiness       | Stopped |
| vmicguestinterface | Stopped |
| UmRdpService       | Running |
| ...                |         |

Rather than list all services, you might instead want to quickly identify services that should be running but aren't. To do this, you can use the following command.

```
get-wmiobject win32_service -filter "startmode='auto' AND state<> 'Running'" | Select Name,State
```

Here, you use `Get-WMIObject Win32_Service` to examine all services that have the Auto start mode (which includes services that have a delayed automatic start) but aren't running. You look for services that have a state other than 'Running' to ensure services with any state other than 'Running' are included in the results, as shown in this sample output.

| name           | state   |
|----------------|---------|
| ----           | -----   |
| RemoteRegistry | Stopped |
| sppsvc         | Stopped |
| workfolderssvc | Stopped |

To determine whether a service that should have started automatically but didn't start because a required service isn't running, you can enter the following command.

```
get-wmiobject win32_service -filter "startmode='auto' AND state<> 'Running'" | get-service -RequiredServices
```

Here, you use `Get-WMIObject Win32_Service` to examine all services that have the Auto start mode but aren't running. You then display the details regarding services that are required for the stopped service to start. If a required service is stopped, that could explain why a service was stopped. You then typically would begin deeper troubleshooting by trying to determine why the required service didn't start. Otherwise, if required services are running or if there are no required services, you would begin deeper troubleshooting by trying to determining why the original service didn't start.

## Managing System Processes

Any time you start an application or type a command at the command line, Windows Server 2016 starts one or more processes to handle the related program. Generally, processes you start in this manner are called interactive processes—that is, you start the processes interactively with the keyboard or mouse. If the application or program is active and selected, the interactive process has control over the keyboard and mouse



until you switch control by terminating the program or selecting a different one. When a process has control, it's said to be running in the foreground .

Processes can also run in the background . For processes started by users, this means that programs that aren't currently active can continue to operate, however, they generally aren't given the same priority as active processes. You can also configure background processes to run independently of the user logon session; the operating system usually starts such processes. An example of this type of background process is a scheduled task run by the operating system. The configuration settings for the task tell the system to execute a command at a specified time.





## Working with Processes

The key tool you use to manage system processes and applications is Task Manager. You can use any of the following techniques to display Task Manager:

- Press Ctrl+Shift+Esc.
- Press Ctrl+Alt+Del, and then click Task Manager.
- Press the Windows logo key, type **taskmgr**, and then press Enter.
- Right-click the taskbar, and then click Task Manager on the shortcut menu.

**NOTE** When you press the Windows logo key and type **taskmgr**, two matches are displayed. One match is the full name, Task Manager. The other match is the command name, taskmgr.

Task Manager has two general views:

- **Summary** Shows only applications running in the foreground, which lets you quickly select and work with foreground applications
- **Expanded** Expands the view, providing additional tabs that you can use to get information about all running processes, system performance, connected users, and configured services

If you are in summary view, you can switch to expanded view by clicking More Details.

If you are in the expanded view, you can switch to summary view by clicking Fewer Details. When you close and reopen Task Manager, the view that you last used is displayed.

Generally, as an administrator, you'll work with the expanded view, which has multiple tabs you can select to work with running processes, system performance, connected users, and configured services. The Processes tab shows the general status of processes. Processes are grouped by type and listed alphabetically within each type by default.

There are three general types:

- Apps, which are programs running in the foreground
- Background processes, which are programs running in the background
- Windows processes, which are processes run by the operating system




| Name                                | Status | PID  | Process name      | 39% CPU | 14% Memory |
|-------------------------------------|--------|------|-------------------|---------|------------|
| <b>Apps (8)</b>                     |        |      |                   |         |            |
| Microsoft Management Console        |        | 5008 | mmc.exe           | 0%      | 11.6 MB    |
| Notepad                             |        | 4404 | notepad.exe       | 0%      | 1.5 MB     |
| Notepad                             |        | 4692 | notepad.exe       | 0%      | 1.5 MB     |
| Server Manager                      |        | 4932 | ServerManager.exe | 0%      | 69.6 MB    |
| Task Manager                        |        | 1296 | Taskmgr.exe       | 0%      | 8.0 MB     |
| Windows Command Processor           |        | 4180 | cmd.exe           | 0%      | 0.4 MB     |
| Windows Explorer                    |        | 1692 | explorer.exe      | 0%      | 20.9 MB    |
| Windows PowerShell                  |        | 2816 | powershell.exe    | 0%      | 20.2 MB    |
| <b>Background processes (17)</b>    |        |      |                   |         |            |
| Distributed File System Replication |        | 2196 | dfsrs.exe         | 0%      | 9.1 MB     |
| Domain Name System (DNS) Server     |        | 2204 | dns.exe           | 0%      | 227.6 MB   |

**NOTE** The Group By Type option on the View menu controls whether grouping is used. If you clear this option, all processes are listed alphabetically without grouping by type. Note also that you can start a new program from within Task Manager by clicking Run New Task on the File menu and then entering a command to run the application. Options are included for running the task with Administrator privileges and for browsing to find the executable you want to work with.

**REAL WORLD** Many Windows processes also are grouped by the service host they are running under, which can include Local Service, Local System, and Network Service. The number of grouped processes is shown in parentheses, and you can expand the related node to view the actual processes. Select Expand All on the View menu to expand all process groups for easy viewing.

The Status column tells you whether an application is running normally or has stopped responding. A blank status is normal and indicates the process is running normally. Any other status indicates a problem, such as when an application might be frozen and you might want to end the task related to it. However, some applications might not respond to the operating system during certain process-intensive tasks. Because of this, you should be certain the application is really not responding before you end its related task.

Other columns on the Processes tab provide a lot of information about running processes. You can use this information to determine which processes are overconsuming system resources, such as CPU time and memory. Although only CPU and Memory columns are displayed by default, others columns can be added by rightclicking any column header and then selecting checkboxes for the additional columns to display.




In addition to name and status, the other available columns include the following:

- **CPU** The percentage of CPU utilization for the process (across all cores). The bold value in the column header represents the total CPU utilization for the server (across all cores).
- **Memory** The total physical memory reserved for the process. The bold value in the column header represents the total physical memory utilization for the server.
- **Command Line** The full file path to the executable running the process, and any command-line arguments passed in when the process was started.
- **PID** The numeric identifier for the process.
- **Process Name** The name of the process or executable running the process.
- **Publisher** The publisher of the process, such as Microsoft Corporation.
- **Type** The general process type displayed as app, background process, or Windows process. This information is useful if you deselect the Group By Type option on the View menu.

Right-clicking an application's listing in Task Manager displays a shortcut menu that you can use to do the following:

- End the application's task.
- Create a dump file for debugging the process.
- Go to the related process on the Details tab.
- Open the file location for the related executable.
- Open the Properties dialog box for the related executable.




The Go To Details option is very helpful when you're trying to find the primary process for a particular application. Selecting this option highlights the related process on the Details tab.

**NOTE** You can stop a process by selecting the process and then clicking End Task. You shouldn't try to stop critical Windows processes by using this technique. If you try to stop a critical Windows process, Task Manager displays a warning prompt stating that ending this process will cause Windows to become unusable or to shut down. If you proceed Windows may crash and display a blue screen.

# Troubleshooting Processes

Task Manager's Details tab provides detailed information about the processes that are running. The columns displayed by default on the Details tab are similar to those provided on the Processes tab:

- **Name** The name of the process or executable running the process
- **User Name** The name of the user or system service running the process
- **CPU** The percentage of CPU utilization for the process
- **Memory (Private Working Set)** The amount of physical memory reserved by the process
- **Status** The run status of the process
- **Description** A description of the process



The screenshot shows the Windows Task Manager window with the 'Details' tab selected. The table displays the following data:

| Name                | PID  | Status  | User name      | CPU | Memory (private w... | Description                                             |
|---------------------|------|---------|----------------|-----|----------------------|---------------------------------------------------------|
| System Idle Process | 0    | Running | SYSTEM         | 95  | 4 K                  | Percentage of time the processor is idle                |
| System              | 4    | Running | SYSTEM         | 04  | 28 K                 | NT Kernel & System                                      |
| Taskmgr.exe         | 1296 | Running | williams       | 00  | 8,188 K              | Task Manager                                            |
| taskhostw.exe       | 2100 | Running | SYSTEM         | 00  | 1,400 K              | Host Process for Windows Tasks                          |
| dwm.exe             | 904  | Running | DWM-2          | 00  | 22,996 K             | Desktop Window Manager                                  |
| lsass.exe           | 696  | Running | SYSTEM         | 00  | 53,072 K             | Local Security Authority Process                        |
| System Interrupts   | -    | Running | SYSTEM         | 00  | 0 K                  | Deferred procedure calls and interrupt service routines |
| MsMpEng.exe         | 2244 | Running | SYSTEM         | 00  | 42,904 K             | Antimalware Service Executable                          |
| svchost.exe         | 1108 | Running | NETWORK SER... | 00  | 24,228 K             | Host Process for Windows Services                       |
| csrss.exe           | 3528 | Running | SYSTEM         | 00  | 1,036 K              | Client Server Runtime Process                           |
| cmd.exe             | 4180 | Running | williams       | 00  | 428 K                | Windows Command Processor                               |
| conhost.exe         | 4532 | Running | williams       | 00  | 2,712 K              | Console Window Host                                     |

Other columns can be added by right-clicking any column header and then clicking Select Columns. Next, add columns by selecting the related option.



The 'Select columns' dialog box is shown with the following options:

- ☐ Package name
- ☒ PID
- ☒ Status
- ☒ User name
- ☐ Session ID
- ☐ Job Object ID
- ☒ CPU
- ☐ CPU time
- ☐ Cycle
- ☐ Working set (memory)
- ☐ Peak working set (memory)

Buttons: OK, Cancel

When you're trying to troubleshoot system problems by using process information, you might want to add the following columns to the view:

- **Base Priority** Displays how much of the system's resources are allocated to a process. To set the priority for a process, right-click the process, choose Set Priority, and then select the new priority from these options: Low, Below Normal, Normal, Above Normal, High, and RealTime. Most processes have a normal priority by default. The highest priority is given to real-time processes.
- **CPU Time** Shows the total amount of CPU cycle time used by a process since it was started. To quickly display the processes that are using the most CPU time, display this column, and then click the column header to sort process entries by CPU time.
- **Data Execution Protection** Specifies whether DEP is enabled or disabled for the process.
- **Elevated** Specifies whether the process is running with elevated, administrator privileges.
- **Handles** Shows the total number of file handles maintained by the process. Use the handle count to gauge how dependent the process is on the file system. Some processes, such as those used by Microsoft Internet Information Services (IIS), have thousands of open file handles. Each file handle requires system memory to be maintained.
- **I/O Reads, I/O Writes** Shows the total number of disk input/output (I/O) reads or writes since the process was started. Together, the number of I/O reads and writes tells you how much disk I/O activity has occurred. If the number of I/O reads and writes is growing disproportionately to actual activity on the server, the process might not be caching files or file caching might not be properly configured. Ideally, file caching reduces the need for I/O reads and writes.
- **Page Faults** Shows whether there are page faults. A page fault occurs when a process requests a page in memory and the system can't find it at the requested location. If the requested page is elsewhere in memory, the fault is called a soft page fault. If the requested page must be retrieved from disk, the fault is called a hard page fault. Most processors can handle large numbers of soft faults. Hard faults, however, can cause significant delays.
- **Paged Pool, NP Pool** Shows the paged pool and nonpaged pool. A paged pool is an area of system memory for objects that can be written to disk when they aren't used. NP pool, or nonpaged pool, is an area of system memory for objects that can't be written to disk. You should note processes that require a large amount of nonpaged



pool memory. If there isn't enough free memory on the server, these processes might be the reason for a high level of page faults.

- **Peak Working Set** Shows the highest amount of memory used by the process. The change, or delta, between current memory usage and peak memory usage is also important to note. Applications that have a high delta between base memory usage and peak memory usage, such as Microsoft SQL Server, might need to be allocated more memory when they start so that they perform better.
- **Platform** Specifies whether the process is running on the 64-bit or 32-bit platform. Windows 64-bit editions support both 64-bit and 32-bit applications by using the Windows on Windows 64 (WoW64) x86 emulation layer. The WoW64 subsystem isolates 32-bit applications from 64-bit applications. This prevents file-system and registry problems. The operating system provides interoperability across the 32bit/64-bit boundary for the Component Object Model (COM) and for basic operations. However, 32-bit processes cannot load 64-bit dynamic-link libraries (DLLs), and 64-bit processes cannot load 32-bit DLLs.
- **Process ID (PID)** Shows the numeric identifier for the process.
- **Session ID** Shows the identifier for the session under which the process is running.
- **Threads** Shows the current number of threads the process is using. Most server applications are multithreaded. Multithreading allows concurrent execution of process requests. Some applications can dynamically control the number of concurrently executing threads to improve application performance. Too many threads, however, can actually reduce performance, because the operating system has to switch thread contexts too frequently.
- **UAC Virtualization** Indicates whether User Account Control (UAC) virtualization is enabled, disabled, or not allowed in the process. UAC virtualization is needed for legacy applications written for early releases of Windows. When UAC virtualization is enabled for these applications, error notifications and error logging related to virtualized files and registry values are written to the virtualized location rather than to the actual location to which the process was trying to write. If virtualization is required but disabled or not allowed, the process will silently fail when trying to write to protected folders or protected areas of the registry.


If you examine processes running in Task Manager, you'll notice a process called System Idle Process. You can't set the priority of this process. Unlike processes that track resource usage, System Idle Process tracks the percentage of system resources that aren't used. Thus, a 99 in the CPU column of the System Idle Process means that 99 percent of system resources currently aren't being used.





Processes that are waiting to use a resource that is locked by another process are in a wait state and can continue only when the locked resource is released. As part of normal operations, resources are locked for one process or another and then released to be used by another process. Sometimes, though, with poorly architected programs, a process can get stuck waiting for a resource that never gets released.

You can view the wait chain for processes by right-clicking the process and then clicking Analyze Wait Chain. If the process is waiting for a resource to be released, the wait chain is displayed for that process. The root node in the wait tree is the process using, or waiting to use, the required resource. A process waiting on another process for a resource might explain why a process doesn't seem as responsive as you might expect.



If you suspect there's a locking problem, you can select one or more processes in the wait chain and then click End Process. Task Manager then stops the processes, which should free the locked resource. However, keep in mind that it is routine and normal for processes to lock resources while those resources are being used and then to free the resources when they are no longer being used. A problem occurs when a process fails to release a resource, as can happen with a poorly architected program.

As you examine processes, keep in mind that a single application might start multiple processes. Generally, these processes are dependent on a central process. From this main process, a process tree containing dependent processes is formed. You can find the main process for an application by right-clicking the application on the Processes tab and selecting Go To Details. When you terminate processes, you'll usually want to



target the main application process or the application itself rather than dependent processes. This ensures that the application is stopped cleanly.

To stop the main application process and dependent processes, you have several choices:

- Right-click the application on the Processes tab, and then click End Task.
- Right-click the main application process on the Details tab, and then click End Task.
- Right-click the main or a dependent process on the Processes tab, and then click End Process Tree.

## Tracking System and Security Problems

Event logs provide historical information that can help you track down system and security problems. To work with services on remote servers, remote management and inbound exceptions for Remote Service Management must be enabled.

The Windows Event Log service controls whether events are tracked. When you start this service, you can track user actions and resource usage events through the event logs. Two general types of log files are used:

- **Windows logs** Logs that the operating system uses to record general system events related to applications, security, setup, and system components.
- **Applications and services logs** Logs that specific applications and services use to record application-specific or service-specific events.

Windows logs that are used include:

- **Application** This log records events logged by applications, such as the failure of SQL Server to access a database. The default location is %SystemRoot%\System32\Winevt\Logs\Application.evtx.
- **Forwarded Events** When event forwarding is configured, this log records forwarded events from other servers. The default location is %SystemRoot%\System32\Config\ForwardedEvents.evtx.
- **Security** This log records events you've set for auditing with local or global group policies. The default location is %SystemRoot%\System32\Winevt\Logs\Security.evtx.

**NOTE** Any user who needs access to the security log must be granted the user right to the following: Manage Auditing and the Security Log. By default, members of the Administrators group have this user right. To learn how to assign user rights, see the section “Configuring User Rights Policies” in Chapter 8.





- **Setup** This log records events logged by the operating system or its components during setup and installation. The default location is %SystemRoot%\System32\Winevt\Logs\Setup.evtx.
- **System** This log records events logged by the operating system or its components, such as the failure of a service to start at system startup. The default location is %SystemRoot%\System32\Winevt\Logs\System.evtx.


**SECURITY ALERT** If you're an administrator, you tend to monitor the application and system logs the most—but you shouldn't forget about the security log. The security log is one of the most important logs, and you should monitor it closely. If the security log on a server doesn't contain events, the likeliest reason is that local auditing hasn't been configured, or that domainwide auditing is configured—in which case, you should monitor the security logs on domain controllers rather than on member servers.

Applications and services logs available include the following:

- **DFS Replication** This log records Distributed File System (DFS) replication activities. The default location is %SystemRoot%\System32\Winevt\Logs\Dfs Replication.evtx.
- **Directory Service** This log records events logged by Active Directory Domain Services (AD DS) and its related services. The default location is %SystemRoot%\System32\Winevt\Logs\Directory Service.evtx.
- **DNS Server** This log records DNS queries, responses, and other DNS activities. The default location is %SystemRoot%\System32\Winevt\Logs\DNS Server.evtx.
- **File Replication Service** This log records file replication activities on the system. The default location is %SystemRoot%\System32\Winevt\Logs\File Replication Service.evtx.
- **Hardware Events** When hardware subsystem event reporting is configured, this log records hardware events reported to the operating system. The default location is %SystemRoot%\System32\Config\Hardware.evtx.
- **Microsoft\Windows** This provides logs that track events related to specific Windows services and features. Logs are organized by component type and event category. Operational logs track events generated by the standard operations of the related component. In some cases, supplemental logs are displayed for analysis, debugging, and recording administration-related tasks.
- **Windows PowerShell** This log records activities related to the use of Windows PowerShell. The default location is %SystemRoot%\System32\Winevt\Logs\Windows PowerShell.evtx.

## Accessing Events in Server Manager

When you are working with Server Manager and select the Local Server node, the All Servers node, or a server group node, the right pane will have an Events panel for reviewing recent warning and error events from the application and system logs. When you select the server you want to work with in the Servers panel, its events are listed in the Events panel.



The columns on the Events panel can be adjusted by right-clicking any column header and then clicking the columns to add or remove. The columns you can use include:

- **Server Name** The name of the server on which the service is running
- **FQDN** The fully qualified domain name of the server on which the service is running
- **ID** Generally, a numeric identifier for the specific event, which could be helpful when searching knowledge bases
- **Severity** The event level as an error or warning
- **Source** The application, service, or component that logged the event
- **Log** The log in which the event was recorded
- **Date And Time** The date and time the event was recorded


**TIP** When you are working with many servers, use the grouping options to help you more easily manage events. You can group events by server name, FQDN, ID, severity, source, log, and date and time by right-clicking any column header, clicking Group By, and then selecting your grouping option.

## Accessing Events in Event Viewer

To work with event logs on remote servers, remote management and inbound exceptions for Remote Event Log Management must be enabled. You access the event logs by following these steps:

1. In Server Manager, select All Servers or any server group node in the left pane.

2. On the Servers panel, right-click the server to which you want to connect.
3. Click Computer Management to automatically connect to the selected server.
4. In Computer Management, you view and work with the event logs by expanding the System Tools node and then selecting the Event Viewer node.



5. Expand the Event Viewer node. You can work with the server's event logs in the following ways:
  - To view all errors and warnings for all logs, expand Custom Views, and then select Administrative Events. In the main pane, a list of all warning and error events for the server is displayed.
  - To view all errors and warnings for a specific server role, expand Custom Views, expand Server Roles, and then select the role to view. In the main pane, a list of all events for the selected role is displayed.
  - To view events in a specific log, expand the Windows Logs node, the Applications And Services Logs node, or both nodes. Select the log you want to view, such as Application or System.
6. Use the information in the Source column to determine which service or process logged a particular event.

Entries in the main pane of Event Viewer provide a quick overview of when, where, and how an event occurred. To obtain detailed information about an event, review the details provided on the General tab in the lower portion of the main pane. The event level or keyword precedes the date and time of the event. Event levels include the following:

- **Information** An informational event, which is generally related to a successful action.

- **Audit Success** An event related to the successful execution of an action.
- **Audit Failure** An event related to the failed execution of an action.
- **Warning** A warning. Details for warnings are often useful in preventing future system problems.
- **Error** A noncritical error, such as the failure of a zone transfer request on a DNS server.
- **Critical** A critical error, such as the Cluster service not responding because a quorum was lost.

**NOTE** Warnings and errors are the two key types of events you'll want to examine closely. Whenever these types of events occur and you're unsure of the cause, review the detailed event description.

In addition to level, date, and time logged, the summary and detailed event entries provide the following information:

- **Source** The application, service, or component that logged the event
- **Event ID** Generally, a numeric identifier for the specific event, which could be helpful when searching knowledge bases
- **Task Category** The category of the event, which is almost always set to None, but is sometimes used to further describe the related action, such as a process or a service
- **User** The user account that was logged on when the event occurred, if applicable
- **Computer** The name of the computer on which the event occurred
- **Description** In the detailed entries, a text description of the event
- **Data** In the detailed entries, any data or error code output by the event

## Filtering Event Logs

Event Viewer creates several filtered views of the event logs for you automatically. Filtered views are listed under the Custom Views node. When you select the Administrative Events node, a list of all errors and warnings for all logs is displayed. When you expand the Server Roles node and then select a role-specific view, a list of all events for the selected role is displayed.

If you want to create a custom view of your own, you can do so in Computer Management by following these steps:

1. In the left pane, right-click the Custom Views node, and then click Create Custom View.



2. Use the Logged list to select a time frame for logging events. You can choose to include events from the last hour, last 12 hours, last 24 hours, last 7 days, or last 30 days. Alternatively, you can set a custom range.
3. Use the Event Level check boxes to specify the level of events to include. Select Verbose to display additional details for events.
4. Create a custom view for either a specific set of logs or a specific set of event sources:
  - Use the Event Logs list to select event logs to include. You can select multiple event logs by selecting their check boxes. If you select specific event logs, all other event logs are excluded.
  - Use the Event Sources list to select event sources to include. You can select multiple event sources by selecting their check boxes. If you select specific event sources, all other event sources are excluded.
5. Optionally, use the User and Computer(s) boxes to specify users and computers that should be included. If you do not specify users and computers to include, events generated by all users and computers are included.
6. When you click OK, Windows displays the Save Filter To Custom View dialog box.



Create Custom View

Filter XML

Logged: Any time

Event level: ☒ Critical ☒ Warning ☐ Verbose  
☒ Error ☐ Information

☒ By log Event logs: Application, Security, System, DNS Server

☐ By source Event sources:

Includes/Excludes Event IDs: Enter ID numbers and/or ID ranges separated by commas. To exclude criteria, type a minus sign first. For example 1,3,5-99,-76

<All Event IDs>

Task category:

Keywords:


User: <All Users>

Computer(s): <All Computers>

Clear

OK Cancel

7. Type a name and description for the custom view.
8. Select where to save the custom view. By default, custom views are saved under the Custom Views node. You can create a new node by clicking New Folder, entering a name for the folder, and then clicking OK.
9. Click OK to close the Save Filter To Custom View dialog box, which displays a filtered list of events. Review these events carefully, and take steps to correct any problems that exist.



If you want to review a particular type of event, you can filter the log in Computer Management by following these steps:

1. Expand Windows Logs or Applications And Services Logs as appropriate for the type of log you want to configure to display a list of event logs.
2. Right-click the log you want to work with, and then click Filter Current Log.
3. Use the Logged list to select the time frame for logging events. You can choose to include events from the last hour, last 12 hours, last 24 hours, last 7 days, or last 30 days.
4. Use the Event Level check boxes to specify the level of events to include. Select Verbose to get additional details.
5. Use the Event Sources list to select event sources to include. If you select specific event sources, all other event sources are excluded.
6. Optionally, use the User and Computer(s) boxes to specify users and computers that should be included. If you do not specify users and computers, events generated by all users and computers are included.
7. Click OK, which displays a filtered list of events. Review these events carefully, and take steps to correct any problems that exist. To clear the filter and display all events for the log, click Clear Filter in the Actions pane or on the Action menu.




# Configuring Event Logging

Log options allow you to control the size of event logs and also how logging is handled. By default, event logs are set with a maximum file size. When a log reaches this limit, events are overwritten to prevent the log from exceeding the maximum file size.

To set log options in Computer Management, follow these steps:

1. Expand Windows Logs or Applications And Services Logs as appropriate for the type of log you want to configure. This displays a list of event logs.
2. Right-click the event log whose properties you want to set, and then click Properties on the shortcut menu.



3. Type or set a maximum size in kilobytes (KB) in the Maximum Log Size box. Make sure that the drive containing the operating system has enough free space for the maximum log size you specify. Log files are stored in the %SystemRoot%\System32\Winevt\Logs directory by default.
4. Select an event log–wrapping mode. The following options are available:
  - **Overwrite Events As Needed (Oldest Events First)** Events in the log are overwritten when the maximum file size is reached. Generally, this is the best option on a low-priority system.





- **Archive The Log When Full, Do Not Overwrite Events** When the maximum file size is reached, Windows archives the events by saving a copy of the current log in the default directory. Windows then creates a new log for storing current events.
- **Do Not Overwrite Events (Clear Logs Manually)** When the maximum file size is reached, the system generates error messages telling you the event log is full.

5. Click OK when you are finished.

**NOTE** On critical systems where security and event logging is very important, you should use Archive The Log When Full, Do Not Overwrite Events. When you use this method, you ensure that event history is preserved in archives automatically. You may also want to consider a centralized logging solution.





## Clearing Event Logs

When an event log is full, you need to clear it. To do that in Computer Management, follow these steps:

1. Expand Windows Logs or Applications And Services Logs as appropriate for the type of log you want to configure. This displays a list of event logs.
2. Right-click the event log whose properties you want to set, and then click Clear Log on the shortcut menu.
3. Click Save And Clear to save a copy of the log before clearing it. Click Clear to continue without saving the log file.





# Archiving the Event Logs

You'll want to keep several months' worth of logs for domain controllers, application servers and other servers that provide essential services. However, it usually isn't practical to set the maximum log size to accommodate this. Instead, you should allow Windows to periodically archive the event logs, or you should manually archive the event logs.





# Using Archiving

Logs can be archived as the following four types of files:

- Event (.evtx) for access in Event Viewer
- Tab-delimited text (.txt) for access in text editors or word processors or to import into spreadsheets and databases
- Comma-delimited text (.csv) for importing into spreadsheets or databases
- XML (.xml) for saving as an XML file

When you export log files to a comma-delimited file, a comma separates each column in the event entry. The event entries look like this.

Information,08/29/16 3:43:24 PM,DNS Server,2,None,The DNS server has started. Error,08/29/14 3:40:04 PM,DNS Server,4015,None,The DNS server has encountered a critical error from the Directory Service (DS). The data is the error code.


The format for the entries is as follows.

Level,Date and Time,Source,Event ID,Task Category,Description



# Creating Log Archives


Windows creates log archives automatically when you select the event log–wrapping mode named Archive The Log When Full, Do Not Overwrite Events. To set this mode, right-click the log in Event Viewer, and then select Properties. Next, in the Properties dialog box, select the archive option and then click OK.




You also can create log archives manually. If you plan to archive logs regularly, you might want to create an archive directory in which you can easily locate the log archives. You should also name the log file so that you can easily determine the log file type, the period of the archive and the originating computer. For example, if you're archiving the system log file for FileServer82, you might want to use the file name **System Log October 2016 - FileServer82**.

To archive logs manually, follow these steps:

1. Expand Windows Logs or Applications And Services Logs as appropriate for the type of log you want to configure. This displays a list of event logs.
2. Right-click the event log you want to archive, and then click Save All Events As on the shortcut menu.
3. In the Save As dialog box, select a directory, and then enter a log file name.
4. In the Save As Type list, Event Files (\*.evtx) is the default file type. Select the log file type you want to use, and then click Save.



5. If you plan to view the log on other computers, you might need to include display information. To save display information, select Display Information For These Languages, choose the language in the list provided, and then click OK.  
Otherwise, just click OK to save the log without display information.




**TIP** The best file type to use for archiving is .evtx. Use this extension if you plan to review old logs in Event Viewer. However, if you plan to review logs in other applications, you might need to save the logs in a tab-delimited or commadelimited format. With the tab-delimited or comma-delimited format, you sometimes need to edit the log file in a text editor for the log to be properly interpreted. If you saved the log with the .evtx extension, you can always save another copy in tab-delimited or comma-delimited format later by doing another Save As after opening the archive in Event Viewer. Note that you might not be able to use the .evtx file type to save events from a remote computer to a local folder. In this case, you need to save the events to the local computer with a different file extension, such as .xml. Otherwise, save the events with the .evtx extension on the remote computer.


## Viewing Log Archives

You can view log archives in text format in any text editor or word processor. You should view log archives in the event log format in Event Viewer. You can view log archives in Event Viewer by following these steps:


1. In Computer Management, select and then right-click the Event Viewer node. On the shortcut menu, click Open Saved Log.



2. In the Open Saved Log dialog box, select a directory and a log file name. By default, the Event Logs Files format is selected. This ensures that logs saved as .evtx, .evt, and .etl are listed. You can also filter the list by selecting a specific file type.



3. Click Open. If you are prompted about converting the log to the new event log format, click Yes.
4. Windows displays the Open Saved Log dialog box. Enter a name and description for the saved log.



5. Specify where to save the log. By default, saved logs are listed under Saved Logs. You can create a new node by clicking New Folder, entering a name for the folder, and then clicking OK.
6. Click OK to close the Open Saved Log dialog box. This displays the contents of the saved log.

**TIP** To remove the saved log from Event Viewer, click Delete in the Actions pane or on the Action menu. When prompted to confirm, click Yes. The saved log file still exists in its original location.







## Chapter 13. Optimizing Server Performance

As part of routine administration and maintenance, you should periodically monitor your servers, track their resource usage, and determine whether there are ways you can optimize performance. Before you start monitoring a server, you might want to establish baseline performance metrics. To do this, you measure server performance at various times and under different load conditions. You can then compare the baseline performance with subsequent performance to determine how the server is performing. Performance metrics that are well above the baseline measurements might indicate areas where the server needs to be optimized or reconfigured.


Optimizing a server isn't something you should do haphazardly. You need to have a clear plan—a set of goals you hope to achieve. Workload planning or troubleshooting a performance problem is a key reason for optimization. For example, users might be having problems connecting to an application running on a server, and you might want to track resource usage and perform other monitoring to identify the source of the problem. Often, you may be able to improve performance by reducing or eliminating the source of performance bottlenecks. As the number of users accessing a server grows, however, you might not be able to provide additional optimizations to improve performance, but you might be able to improve performance by shifting workloads or distributing parts of the workload to other servers.

### Tracking Performance

Before you can optimize a server, you need to understand how the server is using its resources. The easiest way to get a basic understanding of server resource usage is to use Task Manager, which was introduced in the “Managing System Processes” section of Chapter 12 “Maintaining Your Servers.” Open Task Manager by right-clicking Start, and then selecting Task Manager on the shortcut menu.

In Task Manager, the Performance tab provides an overview of CPU, memory and Ethernet usage. The graphs and statistics displayed give you a quick check of how a server is using its resources.





Click a summary entry in the left pane to view detailed information and a related graph in the right pane, including the following information:


- **CPU** A graph of CPU usage plotted over time
- **Memory** A graph of memory usage plotted over time
- **Ethernet** A graph of network throughput plotted over time

To view a close-up of any graph, double-click the graph. Double-clicking again returns you to normal viewing mode.

The Update Speed option on the View menu allows you to change the speed of graph updating and to pause the graph. Updates occur once every 4 seconds for Low, once every 2 seconds for Normal, and twice per second for High.

## Understanding Processor Utilization

When you select CPU in Task Manager, the % Utilization graph shows overall processor utilization for the last 60 seconds. If a system has multiple physical CPUs, a graph is displayed for each CPU by default. You also can view logical processors or NUMA nodes by right-clicking a CPU graph, clicking Change Graph To, and then clicking Logical Processors or NUMA Nodes as appropriate.



You can view kernel times by right-clicking a CPU graph, and then selecting Show Kernel Times. Because usage by the kernel is plotted separately, you can more easily track the amount of CPU time used by the operating system kernel.

**TIP** Tracking the kernel usage can be handy for troubleshooting. For example, if you are using Microsoft Internet Information Services (IIS) with output caching in kernel mode, you can get a better understanding of how kernel caching might be affecting CPU usage and overall performance by showing kernel times. Kernel usage tracking isn't enabled by default because it adds to the overhead of monitoring a server in Task Manager.

The CPU information provided can help you quickly determine the up time for the server, the number of physical processors, the number of logical processors, whether hardware virtualization is enabled, and the amount of on-processor cache for each available register (L1, L2, L3). Keep the following in mind:

- **Handles** Shows the number of I/O handles in use; I/O handles act as tokens that let programs access resources. I/O throughput and disk performance affect a system more than a consistently high number of I/O handles.



- **Threads** Shows the number of threads in use; threads are the basic units of execution within processes.
- **Processes** Shows the number of processes in use; processes are running instances of applications or executable files.
- **Up Time** Shows how long the system has been up since it was last started.

If CPU usage is consistently high, even under average usage conditions, you might want to perform more detailed performance monitoring to determine the cause of the problem. Memory is often a source of performance problems, and you should rule it out before upgrading or adding CPUs. For more details, see “Fine-Tuning System Performance” later in this chapter.





## Tracking Memory Allocation and Usage

When you select Memory in Task Manager, the Memory Usage graph shows overall usage of the private working set for the last 60 seconds. The Memory Composition histogram shows the following:


- **In-Use Memory** The amount of memory being used by processes
- **Modified Memory** The amount of memory whose contents must be written to disk before the memory can be used for another purpose
- **Standby Memory** The amount of memory with cached data and code not actively being used
- **Free Memory** The amount of memory that is not currently allocated for any purpose

You can use the memory information provided to quickly determine the speed of the memory, the number of memory slots used and available, and the memory form factor.

The total amount of physical RAM configured on the server is listed in the upper-right corner when you are working with the memory graphs. Other memory statistics shown below the memory graphs provide the following information:

- **In Use** Shows the amount of physical RAM that is in use on the server.
- **Available** Shows the amount of physical RAM that is available for use (includes memory marked as standby and free ). If a server has very little physical memory free, you might need to add memory to the system. In general, you want the free memory to be no less than 5 percent of the total physical memory on the server.






- **Committed** Lists the virtual memory currently in use followed by the total amount of virtual memory available. If the current page file usage is consistently within 10 percent of the maximum value (meaning consistent usage of 90 percent or more), you might want to add physical memory, increase the amount of virtual memory, or do both.
- **Cached** Shows the amount of memory used for system caching.
- **Paged Pool** Provides information on noncritical kernel memory used by the operating system kernel.
- **Non-paged Pool** Provides information on critical kernel memory used by the operating system kernel.

Critical portions of kernel memory must operate in RAM and can't be paged to virtual memory. Because of this, this type of kernel memory is listed as being in the nonpaged pool. The rest of kernel memory can be paged to virtual memory and is listed as being in the paged pool.

## Understanding Network Utilization

When you select Ethernet, Task Manager provides an overview of the network adapters used by the system. You can use the information provided to quickly determine the

percent utilization, link speed, and operational status of each network adapter configured on a system.



The name of the active network adapter in the Network Connections folder is shown in the upper-right corner. Keep the following in mind:

- If a system has one network adapter, the summary graph shows details of the network traffic on this adapter over time.
- If a system has multiple network adapters, the graph displays a composite index of all network connections, which represents all network traffic.
- If the network adapters are teamed, the graph shows details for Microsoft's multiplexor driver and the network traffic presented is representative of all team interfaces.

To view detailed information about link speed, link state, bytes sent, bytes received, and more, right-click the Network Throughput graph and then select View Network Details. When working with network details, keep the following in mind:

- **Network Utilization** Percentage of network usage based on the initial connection speed for the interface or the combined speed of teamed interfaces. For





example, an adapter with an initial link speed of 10 gigabits per second (Gbps) and current traffic of 100 megabits per second (Mbps) is utilized at 1 percent.

- **Link Speed** Connection speed of the interface as determined by the initial connection speed, such as 1 Gbps or 10 Gbps.
- **State** Operational status of network adapters, such as Connected or Disconnected.
- **Bytes Sent Throughput** Percentage of current connection bandwidth used by traffic sent from the system.
- **Bytes Received Throughput** Percentage of current connection bandwidth used by traffic received by the system.
- **Bytes Throughput** Percentage of current connection bandwidth used for all traffic on the network adapter.
- **Bytes Sent** Cumulative total bytes sent on the connection to date.
- **Bytes Received** Cumulative total bytes received on the connection to date.
- 

| Network Details                   |               |
|-----------------------------------|---------------|
| Property                          | Ethernet      |
| Network utilization               | 28%           |
| Link speed                        | 20 Gbps       |
| State                             | Connected     |
| Bytes sent throughput             | 41%           |
| Bytes received throughput         | 35%           |
| Bytes throughput                  | 78%           |
| Bytes sent                        | 886,699,787   |
| Bytes received                    | 888,998,354   |
| Bytes                             | 1,775,698,141 |
| Bytes sent per interval           | 884,409       |
| Bytes received per interval       | 881,071       |
| Bytes per interval                | 1,765,480     |
| Unicasts sent                     | 813,404       |
| Unicasts received                 | 815,421       |
| Unicasts                          | 1,728,825     |
| Unicasts sent per interval        | 813           |
| Unicasts received per interval    | 815           |
| Unicasts per interval             | 1,628         |
| Nonunicasts sent                  | 339,958       |
| Nonunicasts received              | 67            |
| Nonunicasts                       | 340,025       |
| Nonunicasts sent per interval     | 330           |
| Nonunicasts received per inter... | 14            |
| Nonunicasts per interval          | 344           |

Bytes

Cumulative total bytes on the connection to date.






Any time usage is consistently approaching or exceeding 50 percent of total capacity, you should start monitoring the server more closely, and you might want to consider adding and then teaming network adapters.

## Tracking Remote User Sessions

You can use Remote Desktop to connect to remote systems and manage them remotely, as if you were sitting at the console. Windows Server 2016 allows up to two active console sessions at a time.

As remote sessions can impact performance, you'll want to determine what, if any, impact remote users are have on servers you are monitoring. One way to view and manage remote desktop connections is to use the Users tab in Task Manager. The Users tab shows interactive user sessions for both local and remote users.



The screenshot shows the Windows Task Manager window with the 'Users' tab selected. The window title is 'Task Manager' and it has a menu bar with 'File', 'Options', and 'View'. Below the menu bar are tabs for 'Processes', 'Performance', 'Users', 'Details', and 'Services'. The 'Users' tab is active, displaying a table of user sessions. Above the table, there are summary statistics: '14%' for CPU and '13%' for Memory. The table has columns for 'User', 'Status', 'CPU', and 'Memory'. There are two users listed: 'edg (27)' and 'williams (14)'. The 'edg (27)' user is highlighted with a grey background and shows 6% CPU and 254.6 MB Memory. The 'williams (14)' user is highlighted with a yellow background and shows 8% CPU and 227.8 MB Memory. At the bottom of the window, there is a 'Fewer details' button on the left and 'Disconnect' and 'Sign out' buttons on the right.

| User          | Status | CPU | Memory   |
|---------------|--------|-----|----------|
| edg (27)      |        | 6%  | 254.6 MB |
| williams (14) |        | 8%  | 227.8 MB |

Each user connection is listed with user name, status, CPU utilization, and memory usage by default. Other columns can be added by right-clicking any column header and then clicking the columns to add. Available columns include

- **ID** The session ID. The first logon has a session ID of 1. The second logon has an ID of 2.
- **Session** The session type. A user logged on to the local system is listed with Console as the session type. Other users have a session type that indicates the connection type and protocol being used, such as RDP-TCP for a connection that is using the Remote Desktop Protocol (RDP) with TCP as the transport protocol.
- **Client name** For remote connections, lists the name of the originating client computer.

CPU and memory utilization details are convenient for troubleshooting performance issues related to logged-on users. The combined utilization value is listed above the column heading, and individual utilization values for each logged-on user are listed below it.



In the example, the server’s CPU is 95% utilized by the logged-on users. This high usage level could affect the overall performance of the server, and the server might not be as responsive when performing other tasks.


If you right-click a user session, the options available include:

- **Connect** Connects a remote user session if it’s inactive.
- **Disconnect** Disconnects a local or remote user session, halting all user-started applications without saving application data.
- **Sign Off** Logs off a user by using the normal logoff process. Application data and system state information are saved just as they are during a typical logoff.



**Send Message** Sends a console message to a logged-on user.

When you are working with the Users tab, the user’s name is followed by the number of processes that user is running. If you double-click the user’s name or right-click the user name and select Expand, an entry for each running process is displayed. Processes are listed by name, CPU usage, and memory usage.



# Preparing for Monitoring

Tracking baseline performance metrics for your servers is an important part of routine administration and maintenance. You establish baselines by measuring server performance at various times and under different load conditions, so you can compare these baselines with subsequent performance to identify areas where the server might need to be optimized or reconfigured.





## Creating a Monitoring Plan

After you establish baseline metrics, you should formulate a monitoring plan. A comprehensive monitoring plan includes the following steps:

1. Determine which server events should be monitored to help you accomplish your goal.
2. Set filters to reduce the amount of information collected.
3. Configure performance counters to watch resource usage.
4. Log the event data so that it can be analyzed.
5. Analyze the event data to help find solutions to problems.

Although each step is important in a comprehensive plan, there are times when you might not want to go through all the monitoring steps. For example, you might want to monitor and analyze activity as it happens rather than log and analyze the data later.

The primary tools you use to monitor servers include the following:

- **Performance Monitor** Displays resource usage over time via counters that you configure. You can use this information to gauge the performance of the server and determine areas that can be optimized.
- **Reliability Monitor** Tracks changes to the system and compares them to changes in system stability. This gives you a graphical representation of the relationship between changes in the system configuration and changes in system stability.
- **Resource Monitor** Provides detailed information about resource usage on the server. The information provided is similar to that provided by Task Manager (though more extensive).
- **Event logs** Provides information that you can use to troubleshoot systemwide problems, including those from the operating system and configured applications. The primary logs you work with are the system, security, and application event logs, as well as logs for configured server roles.


## Using the Monitoring Consoles

Resource Monitor, Reliability Monitor, and Performance Monitor are the tools of choice for performance tuning. You can access Resource Monitor by pressing Ctrl+Shift+Esc and then clicking the Open Resource Monitor button on the Performance tab of Task Manager. In Server Manager, you can access Resource Monitor in a standalone console by clicking Tools and then clicking Resource Monitor.




Resource usage statistics are broken down into four categories:

- **CPU usage** The summary details show the current CPU utilization and the maximum CPU frequency (as related to processor idling). Expanding the CPU entry (by clicking the options button) displays a list of currently running executables by name, process ID, description, status, number of threads used, current CPU utilization, and average CPU utilization.
- **Disk usage** The summary details show the number of kilobytes per second being read from or written to disk, and the highest percentage of usage. Expanding the Disk entry (by clicking the options button) displays a list of currently running executables that are performing or have performed I/O operations is displayed by name, process ID, file being read or written, average number of bytes being read per second, average number of bytes being written per second, total number of bytes being read and written per second, I/O priority, and the associated disk response time.
- **Network usage** The summary details show the current network bandwidth utilization in kilobytes and the percentage of total bandwidth utilization. Expanding the Network entry (by clicking the options button) displays a list of currently running executables that are transferring or have transferred data on the network is displayed by name, process ID, server or IP address being contacted, average number of bytes being sent per second, average number of bytes received per second, and total bytes sent or received per second.
- **Memory usage** The summary details show the current memory utilization and the number of hard faults occurring per second. Expanding the Memory entry (by clicking the options button) displays a list of currently running executables by name, process ID, hard faults per second, commit memory in KB, working set memory in KB, shareable memory in KB, and private (nonshareable) memory in KB.



**TIP** You also can open Resource Monitor by typing **perfmon /rel** in the Search box and then pressing Enter.

Performance Monitor displays statistics graphically for the set of performance parameters you've selected for display. These performance parameters are referred to as counters. When you install certain applications on a system, Performance Monitor might be updated with a set of counters for tracking the server's performance. You also can update these counters when you install additional services and add-ons for the application.



In Server Manager, you can access Performance Monitor in a stand-alone console by clicking Tools and then clicking Performance Monitor. In Computer Management, you can access the tool as a snap-in in the System Tools node. Expand System Tools\Performance\Monitoring Tools, and then click Performance Monitor.


Performance Monitor creates a graph depicting the counters you're tracking. The update interval for this graph is set to 1 second by default, but it can be configured with a different value. As is apparent when you work with Performance Monitor, the tracking information is most valuable when you record performance information in a log file so that the information can be played back. Also, Performance Monitor is helpful when you configure alerts to send messages when certain events occur.

Windows Server 2016 also includes Reliability Monitor. To access Reliability Monitor, follow these steps:



1. In Control Panel, under the System And Security heading, click Review Your Computer's Status And Resolve Issues.
2. In Action Center, expand the Maintenance panel, and then click View Reliability History.

Alternatively, you can run Reliability Monitor by entering **perfmon /rel** at a command prompt or in the Search box.



Reliability Monitor tracks changes to the server and compares them to changes in system stability, which provides a graphical representation of the relationship between changes in the system configuration and changes in system stability. Recording software installation, software removal, application failures, hardware failures, Windows failures, and key events regarding the configuration of the server provides a timeline of changes in both the server and its reliability, which you can use to pinpoint changes that are causing problems with stability. For example, if you notice of a sudden drop in stability, you can click a data point and then expand the related data set to find the specific event that caused the drop in stability.

## Choosing Counters to Monitor

Performance Monitor displays information only for counters you're tracking. Thousands of counters are available, and you'll find counters related to just about every server role you've installed. The easiest way to learn about these counters is to read the explanations available in the Add Counters dialog box. Start Performance Monitor, click Add on the toolbar, and then expand an object in the Available Counters list. Select the Show Description check box, and then scroll through the list of counters for this object.


When Performance Monitor is monitoring a particular object, it can track all instances of all counters for that object. Instances are multiple occurrences of a particular counter. For example, when you track counters for the Processor object on a multiprocessor system, you have a choice of tracking all processor instances or specific processor instances. If you think a particular processor is going bad or experiencing other problems, you could monitor just that processor instance.

You need to be at least a member of the Performance Monitor Users group in the domain or on the local computer to perform remote monitoring. When you use performance logging, you need to be at least a member of the Performance Log Users group in the domain or on the local computer to work with performance logs on remote computers.

Performance Monitor has several views and view types. Be sure that current activity is displayed by clicking View Current Activity on the toolbar or pressing Ctrl+T. You can switch between the view types (Line, Histogram Bar, and Report) by clicking Change Graph Type or pressing Ctrl+G.

To select which counters you want to monitor, follow these steps:

1. Open Performance Monitor and then click Add on the toolbar, or press Ctrl+I. This displays the Add Counters dialog box.
2. In the Select Counters From Computer list, enter the Universal Naming Convention (UNC) name of the server you want to work with, such as \\CorpServer84, or choose <Local Computer> to work with the local computer.



3. In the Available Counters panel, performance objects are listed alphabetically. If you select an object entry by clicking it, all related counters are selected. Expanding an object entry displays all the related counters, and you can select individual counters by clicking them. For example, you could expand the entry for the Active Server Pages object and then select the Requests Failed Total, Requests Not Found, Requests Queued, and Requests Total counters.
4. When you select an object or any of its counters, the related instances are displayed. Choose All Instances to select all counter instances for monitoring, or select one or more counter instances to monitor. For example, you could select instances of Anonymous Users/Sec for individual websites or for all websites.
5. When you've selected an object or a group of counters for an object in addition to the object instances, click Add to add the counters to the graph.
6. Repeat steps 3–5 to add other performance parameters.
7. Click OK when you are finished.

**TIP** Don't try to chart too many counters or counter instances at once. You'll make the display too difficult to read, and you'll use system resources—namely, CPU time and memory—that might affect server responsiveness.


## Performance Logging

You use data collector sets to specify groups of performance objects and counters to track. After you create a data collector set, you can easily start or stop monitoring the



performance objects and counters included in the set. In a way, this makes data collector sets similar to the performance logs used in early releases of Windows. However, data collector sets are much more sophisticated. You can use a single data set to generate multiple performance counters and trace logs. You can also do the following:

- Assign access controls to manage who can access collected data
- Create multiple run schedules and stop conditions for monitoring
- Use data managers to control the size of collected data and reporting
- 



Generate reports

based on collected data

While you are working with Performance Monitor, you can review currently configured data collector sets and reports under the Data Collector Sets and Reports nodes, respectively. You'll find data sets and reports that are user-defined and system-defined. User-defined data sets are created by users for general monitoring and performance tuning. System-defined data sets are created by the operating system to aid in automated diagnostics.

## Creating and Managing Data Collector Sets

You can view the currently configured data collector sets by clicking the Performance Monitor option in the Administrative Tools program group, and then expanding the Data Collector Sets node. Work with data collectors in a variety of ways:

- You can view currently defined user or system data collector sets by selecting either User Defined or System, as appropriate. When you select a data collector set in the left pane, related data collectors are displayed in the main pane listed by name and type. The Trace type is for data collectors that record performance data whenever related events occur. The Performance Counter type is for data collectors that record



data on selected counters when a predetermined interval has elapsed. The Configuration type is for data collectors that record changes to particular registry paths.

- You can view running event traces by selecting Event Trace Sessions. You can then stop a data collector that is running a trace by right-clicking it and then selecting Stop.
- You can view the enabled or disabled status of event traces configured to run automatically when you start the computer by selecting Startup Event Trace Sessions. You can start a trace by right-clicking a startup data collector and then selecting Start As Event Trace Session. You can delete a startup data collector by right-clicking it and then clicking Delete.
- You can save a data collector as a template that can be used as the basis of other data collectors by right-clicking the data collector and then selecting Save Template. In the Save As dialog box, select a directory, enter a name for the template, and then click Save. The data collector template is saved as an XML file that can be copied to other systems.
- You can delete a user-defined data collector by right-clicking it and then selecting Delete. If a data collector is running, you need to stop collecting data first and then delete the collector. Deleting a collector also deletes the related reports.

## Collecting Counter Data


You use data collectors to record performance data on the selected counters at a specific sampling interval. For example, you could sample performance data for the CPU every 15 minutes.

The default location for logging is %SystemDrive%\PerfLogs\Admin. Log files can grow in size quickly. If you plan to log data for an extended period, be sure to place the log file on a drive with lots of free space. Remember, the more frequently you update the log file, the greater the drive space and CPU resource usage on the system.

Follow these steps to collect performance counter data:

1. Open Performance Monitor. Under the Data Collector Sets node, right-click the User Defined node in the left pane, point to New, and then click Data Collector Set.
2. In the Create New Data Collector Set Wizard, enter a name for the data collector, such as **System Performance Monitor** or **Processor Status Monitor**. Note that if you enter an invalid name, such as one with a nonalphanumeric character, you won't be able to continue. Select the Create Manually option, and then click Next.





← Create new Data Collector Set.

How would you like to create this new data collector set?

Name:  
System Performance Monitor

☐ Create from a template (Recommended)

☒ Create manually (Advanced)

3. On the What Type Of Data Do You Want To Include page, the Create Data Logs option is selected by default. Select the Performance Counter check box, and then click Next.



← Create new Data Collector Set.

What type of data do you want to include?

☒ Create data logs


☒ Performance counter

☐ Event trace data

☐ System configuration information

☐ Performance Counter Alert

4. On the Which Performance Counters Would You Like To Log page, click Add. This displays the Add Counters dialog box, which you can use as previously discussed to select the performance counters to track. When you are finished selecting counters, click OK.
5. On the Which Performance Counters Would You Like To Log page, enter a sampling interval and select a time unit in seconds, minutes, hours, days, or weeks. The sampling interval specifies when new data is collected. For example, if you sample every 15 minutes, the data log is updated every 15 minutes. Click Next when you are ready to continue.



← Create new Data Collector Set.

Which performance counters would you like to log?

Performance counters:


\Processor Information[\_Total]\% Processor Time

Add...

Remove

Sample interval: 15 Units: Minutes

6. On the Where Would You Like The Data To Be Saved page, type the root path to use for logging collected data. Alternatively, click Browse, and then use the Browse For Folder dialog box to select the logging directory. Click Next when you are ready to continue.



← Create new Data Collector Set.

Where would you like the data to be saved?

Root directory:

%systemdrive%\PerfLogs\Admin\System Performance Monitor

Browse...

7. On the Create Data Collector Set page, the Run As box lists <Default> to indicate that the log will run under the privileges and permissions of the default system account. To run the log with the privileges and permissions of another user, click Change. Enter the user name and password for the account, and then click OK. User names can be entered in domain\username format, such as imaginedl\williams for the Williams account in the Imaginedl domain.



8. Click the Open Properties For This Data Collector Set option, and then click Finish. This saves the data collector set, closes the wizard, and then opens the related Properties dialog box.
9. By default, logging is configured to start manually. To configure a logging schedule, click the Schedule tab, and then click Add. You can now set the Active Range, Start Time, and run days for data collection.
10. By default, logging stops only if you set an expiration date as part of the logging schedule. Using the options on the Stop Condition tab, you can configure the log file to stop automatically after a specified period of time, such as seven days, or when the log file is full (if you set a maximum size limit).
11. By default, Windows doesn't perform any additional operations when monitoring stops. However, if you want to run a script or perform other tasks when data collection stops, you can use the options on the Task tab to do this.
12. Click OK when you've finished setting the logging schedule and stop conditions.

You can manage the data collector as explained previously.


## Collecting Trace Data

You can use data collectors to record performance trace data whenever events related to their source providers occur. A source provider is an application or operating system service that has traceable events.

To collect performance trace data, follow these steps:

1. In Performance Monitor, under the Data Collector Sets node, right-click the User Defined node in the left pane, point to New, and then click Data Collector Set.

2. In the Create New Data Collector Set Wizard, type a name for the data collector, such as **Logon Trace** or **Disk IO Trace** . Note that if you type an invalid name, such as one with a nonalphanumeric character, you won't be able to continue. Select the Create Manually option, and then click Next.



← Create new Data Collector Set.

How would you like to create this new data collector set?

Name:

☐ Create from a template (Recommended)

☒ Create manually (Advanced)

3. On the What Type Of Data Do You Want To Include page, the Create Data Logs option is selected by default. Select the Event Trace Data check box, and then click Next.



← Create new Data Collector Set.


What type of data do you want to include?

☒ Create data logs

- ☐ Performance counter
- ☒ Event trace data
- ☐ System configuration information

☐ Performance Counter Alert

4. On the Which Event Trace Providers Would You Like To Enable page, click Add. Select an event trace provider to track, and then click OK. By selecting individual properties in the Properties list and clicking Edit, you can track particular property values rather than all values for the provider. Repeat this process to select other event trace providers to track. Click Next when you are ready to continue.




5. Complete steps 6–12 from the procedure in the previous section, “Collecting Counter Data.”

# Collecting Configuration Data

You can use data collectors to record changes in registry configuration. To collect configuration data, follow these steps:

1. In Performance Monitor, under the Data Collector Sets node, right-click the User Defined node, point to New, and then click Data Collector Set.



← Create new Data Collector Set.

How would you like to create this new data collector set?

Name:  
AD Registry

☐ Create from a template (Recommended)

☒ Create manually (Advanced)

2. In the Create New Data Collector Set Wizard, enter a name for the data collector, such as **AD Registry** or **Registry Adapter Info**.
3. Select the Create Manually option, and then click Next.
4. On the What Type Of Data Do You Want To Include page, the Create Data Logs option is selected by default. Select the System Configuration Information check box, and then click Next.



← Create new Data Collector Set.

What type of data do you want to include?

☒ Create data logs

☐ Performance counter

☐ Event trace data

☒ System configuration information

☐ Performance Counter Alert

5. On the Which Registry Keys Would You Like To Record page, click Add. Type the registry path to track. Repeat this process to add other registry paths to track.  
Click Next when you are ready to continue.





← Create new Data Collector Set.

Which registry keys would you like to record?

Registry keys:

Add...

Remove


6. Complete steps 6–12 from the procedure in the section “Collecting Counter Data.”

## Viewing Data Collector Reports

During troubleshooting, you'll often want to log performance data over an extended period of time and then review the data to analyze the results. For each data collector that has been or is currently active, you'll find related data collector reports. As with data collector sets themselves, data collector reports are organized into two general categories: user-defined and system.

You can view data collector reports in Performance Monitor. Expand the Reports node, and then expand the individual report node for the data collector you want to analyze. Under the data collector's report node, you'll find individual reports for each logging session. A logging session begins when logging starts and ends when logging is stopped.

The most recent log is the one with the highest log number. If a data collector is actively logging, you won't be able to view the most recent log. You can stop collecting data by right-clicking a data collector set and clicking Stop. For performance counters, collected data is shown by default in a graph view from the start of data collection to the end of data collection.



You can modify the report details by using the following process:

1. In the monitor pane, press Ctrl+Q or click the Properties button on the toolbar. This displays the Performance Monitor Properties dialog box.
2. On the Source tab, specify data sources to analyze. Under Data Source, click Log Files and then click Add to open the Select Log File dialog box. You can now select additional log files to analyze.
3. Specify the time window you want to analyze. Click Time Range, and then drag the Total Range bar to specify the appropriate starting and ending times. Drag the left edge to the right to move up the start time. Drag the right edge to the left to make the end time later.
4. Click the Data tab. You can now select counters to view. Select a counter, and then click Remove to remove it from the graph view. Click Add to display the



Add Counter dialog box, which you can use to select the counters you want to analyze.

NOTE Only counters you selected for logging are available. If a counter you want to work with isn't displayed, you need to modify the data collector properties, restart the logging process, and then check the logs at a later date.

5. Click OK. In the monitor pane, click the Change Graph Type button to select the type of graph.

# Configuring Performance Counter Alerts

You can configure alerts to notify you when certain events occur or when certain performance thresholds are reached. You can send these alerts as network messages and as events that are logged in the application event log. You can also configure alerts to start applications and performance logs.

To configure an alert, follow these steps:

1. In Performance Monitor, under the Data Collector Sets node, right-click the User Defined node in the left pane, point to New, and then click Data Collector Set.
2. In the Create New Data Collector Set Wizard, enter a name for the data collector, such as **Processor Alert** or **Disk IO Alert** . Click the Create Manually option, and then click Next.



← Create new Data Collector Set.

How would you like to create this new data collector set?

Name:  
Processor Alert

☐ Create from a template (Recommended)

☒ Create manually (Advanced)

3. On the What Type Of Data Do You Want To Include page, click the Performance Counter Alert option, and then click Next.



← Create new Data Collector Set.

What type of data do you want to include?

☐ Create data logs


- ☐ Performance counter
- ☐ Event trace data
- ☐ System configuration information

☒ Performance Counter Alert

4. On the Which Performance Counters Would You Like To Monitor page, click Add to display the Add Counters dialog box. This dialog box is identical to the

Add Counters dialog box discussed previously. Use the dialog box to add counters that trigger the alert.

5. In the Performance Counters panel, select the first counter, and then use the Alert When box to set the occasion when an alert for this counter is triggered. Alerts can be triggered when the counter is above or below a specific value. Select Above or Below, and then set the trigger value. The unit of measurement is whatever makes sense for the currently selected counter or counters. For example, to generate an alert if processor time is over 90percent, select Above, and then type **90** . Repeat this process to configure other counters you've selected.



6. Complete steps 6–12 from the procedure in the section “Collecting Counter Data.”

## Fine-Tuning System Performance

Now that you know how to monitor your system, let's look at how you can tune the operating system and hardware performance. Let's examine the following areas:

- Memory usage and caching
- Processor utilization
- Disk I/O
- Network bandwidth and connectivity



## Tuning Physical and Virtual Memory

You should always rule out memory problems before examining other areas of the system. Servers use both physical and virtual memory. To rule out memory problems with a system, you should configure application performance, memory usage, and data throughput settings, and then monitor the server's memory usage to check for problems.

Application performance and memory usage settings determine how system resources are allocated. In most cases, you want to give the operating system and background applications the lion's share of resources. This is especially true for Active Directory, file, print, and network and communications servers. On the other hand, for application, database, and streaming media servers, you should give the programs a server is running the most resources.

You control application performance by following these steps:

1. Open the System Properties dialog box with the Advanced tab selected by typing **SystemPropertiesAdvanced** in the Search box, and then pressing Enter.
2. In the Performance Options dialog box, click the Advanced tab.
3. In the Processor Scheduling panel, you have the following options:
  - **Programs** Use this option to give the active application the best response time and the greatest share of available resources. Generally, you'll want to use this option only on development servers or when you are using Windows Server 2016 as your desktop operating system.
  - **Background Services** Use this option to give background applications a better response time than the active application. Generally, you'll want to use this option for production servers.
4. Click OK.

Using the monitoring techniques discussed previously in this chapter, you can determine memory is being used and check for problems. Counters you'll want to track to uncover memory, caching, and virtual memory (paging) bottlenecks include:

- Memory\Available Kbytes
  - Memory\Committed Bytes
  - Memory\Pages Input/sec
  - Memory\Page Reads/sec
  - Memory\Pool Paged Bytes
- 



## Memory\Pool Nonpaged Bytes

Memory\Available Kbytes is the amount of physical memory available to processes running on the server. Memory\Committed Bytes is the amount of committed virtual memory. If the server has very little available memory, you might need to add memory to the system. In general, you want the available memory to be no less than 5 percent of the total physical memory on the server. If the server has a high ratio of committed bytes to total physical memory on the system, you might also need to add memory. In general, you want the committed bytes value to be no more than 75 percent of the total physical memory.

A page fault occurs when a process requests a page in memory and the system can't find it at the requested location. If the requested page is elsewhere in memory, the fault is called a soft page fault. If the requested page must be retrieved from disk, the fault is called a hard page fault. Most processors can handle large numbers of soft faults. Hard faults, however, can cause significant delays.

Page Faults/sec is the overall rate at which the processor handles all types of page faults. Pages Input/sec is the total number of pages read from disk to resolve hard page faults. Page Reads/sec is the total disk reads needed to resolve hard page faults. Pages Input/sec will be greater than or equal to Page Reads/sec and can give you a good idea of your hard page fault rate. A high number of hard page faults could indicate that you need to increase the amount of memory or reduce the cache size on the server.

Memory\Pool Paged Bytes and Memory\Pool Nonpaged Bytes track the number of bytes in the paged and nonpaged pool. The paged pool is an area of system memory for objects that can be written to disk when they aren't used. The nonpaged pool is an area of system memory for objects that can't be written to disk. If the size of the paged pool is large relative to the total amount of physical memory on the system, you might need to add memory to the system. If the size of the nonpaged pool is large relative to the total amount of virtual memory allocated to the server, you might want to increase the virtual memory size.

## Tuning Processors and Disks

A server's CPU does the actual processing of information on your server. As you examine a server's performance, you should focus on the CPU after you eliminate memory bottlenecks. If the server's processors are the performance bottleneck, adding memory, drives, or network connections won't overcome the problem. Instead, you might need to upgrade the processors to faster clock speeds or add processors to





increase the server's upper capacity. You could also move processor-intensive applications, such as Microsoft SQL Server, to another server.

Before you make a decision to upgrade CPUs or add CPUs, you should rule out problems with memory and caching. If signs still point to a processor problem, you should monitor performance counters for each CPU installed on the server, including System\Processor Queue Length and Processor\% Processor Time.

System\Processor Queue Length displays the number of threads waiting to be executed. These threads are queued in an area shared by all processors on the system. If this counter has a sustained value of more than 10 threads per processor, you need to upgrade or add processors.

Processor\% Processor Time displays the percentage of time the selected CPU is executing a nonidle thread. You should track this counter separately for all processor instances on the server. If the % Processor Time values are high while the network interface and disk I/O throughput rates are relatively low, you need to upgrade or add processors.

Accessing memory is much faster than accessing disks, so if a server has to do a lot of disk reads and writes, a server's overall performance can be degraded. To reduce the amount of disk I/O, you want the server to manage memory efficiently and page to disk only when necessary. You monitor and tune memory usage as discussed in the section "Tuning Physical and Virtual Memory" earlier in this chapter.

In addition to memory tuning, you can monitor some counters to gauge disk I/O activity. Counters you should monitor include:

- PhysicalDisk\% Disk Time
- PhysicalDisk\Disk Writes/sec
- PhysicalDisk\Disk Reads/sec
- PhysicalDisk\Avg. Disk Write Queue Length
- PhysicalDisk\Avg. Disk Read Queue Length
- PhysicalDisk\Current Disk Queue Length

Track PhysicalDisk\% Disk Time in conjunction with Processor\% Processor Time and

Network Interface Connection\Bytes Total/sec. If the % Disk Time value is high and the processor and network connection values aren't high, the system's hard disk drives might be creating a bottleneck. Be sure to monitor % Disk Time for all hard disk drives on the server.







Track PhysicalDisk\Disk Writes/sec and PhysicalDisk\Disk Reads/sec in conjunction with PhysicalDisk\Avg. Disk Write Queue Length, PhysicalDisk\Avg. Disk Read Queue Length and PhysicalDisk\Current Disk Queue Length. The number of writes and reads per second tell you how much disk I/O activity there is. The write and read queue lengths tell you how many write or read requests are waiting to be processed. In general, you want very few waiting requests. Keep in mind that the request delays are proportional to the length of the queues minus the number of drives in a redundant array of independent disks (RAID) set.





## Tuning Networking

Few factors matter more to the way a user perceives your server's performance than the network that connects your server to the user's computer. The delay, or latency, between when a request is made and the time it's received can make all the difference. With a high degree of latency, it doesn't matter if you have the fastest server on the planet: the user experiences a delay and perceives that your servers are slow.

Generally speaking, much of the latency the user experiences is beyond your control. It's a function of the type of connection the user has and the route the request takes to your server. The total capacity of your server to handle requests, the amount of bandwidth available to your servers and the service provider you use are factors under your control, however. Network bandwidth availability is a function of your organization's network infrastructure. Network capacity is a function of the network cards and interfaces configured on the servers. Network responsiveness is influenced by the way your service provider connects your servers to the Internet.

The capacity of your network card can be a limiting factor in some instances. Although 10-Gbps networking is increasingly being used, most servers use 100-Mbps or 1-Gbps network cards, which can be configured in many ways. Someone might have configured a 1-Gbps card for 100 Mbps, or the card might be configured for half duplex instead of full duplex. If you suspect a capacity problem with a network card, you should always check the configuration.

As modern servers often ship with multiple network cards, be sure to check all enabled network adapters. If a server has multiple adapters, you might want to enable and configure NIC teaming to ensure the available bandwidth is used optimally. When NIC teaming is being used, you'll also want to ensure the configuration is optimized for the way the server is currently being used. In Server Manager, you can configure NIC teaming as a Local Server option, which means you must log on locally to the server you want to configure or access the server through a Remote Desktop Connection.

You can determine the throughput and current activity on a server's network cards using the following counters:

- [Network Adapter\Bytes Received/sec](#)
- [Network Adapter\Bytes Sent/sec](#)
- [Network Adapter\Bytes Total/sec](#)
- [Network Adapter\Current Bandwidth](#)

If the total bytes-per-second value is more than 50 percent of the total capacity under average load conditions, your server might have problems under peak load conditions.





You might want to ensure that operations that take a lot of network bandwidth, such as network backups, are performed on a separate interface card. Compare these values in conjunction with `PhysicalDisk\% Disk Time` and `Processor\% Processor Time`.

If the disk time and processor time values are low but the network values are very high, you might have a capacity problem. Solve the problem by optimizing the network card settings or by adding a network card. Remember, planning is everything—it isn't always as simple as inserting a card and plugging it into the network.





## Chapter 14. Using Group Policy

Even when starting up, computers running business and server editions of Windows start reaching for Group Policy, looking for policies to apply. At first the policies are only for the computer itself; but after a user logs on, Windows applies policies to the user as well. Windows also checks periodically for updates to policy and refreshes policy in specific intervals.

Business and server editions of Windows are so policy-dependent that they rely on policies to establish the core security parameters and on the Group Policy client service to isolate Group Policy notification and processing from the Windows logon process. Separating Group Policy from the Windows logon process reduces the resources used for background processing of policy while increasing overall performance and allowing delivery and application of new Group Policy files as part of the update process without requiring a restart, which ensures Windows can be responsive to policy changes as they occur.

This chapter explains how to use Group Policy, create policy objects, and work policy editors. For ways to optimize, maintain, and troubleshoot policy, see Chapter 15.





# Working with Group Policy

As an IT professional, you can use Group policies to simplify administration and centralize control over the privileges, permissions, and capabilities of both user accounts and computers. You can apply group policies to multiple domains, individual domains, subgroups within a domain, or individual systems. Policies that apply to individual systems are referred to as local group policies and are stored on the local system only. Other group policies are linked as objects in the Active Directory data store and apply throughout various areas of the enterprise, according to their applicable scope.

Group policies apply only to systems running business editions of Windows-based operating systems. To understand Group Policy itself, you need to understand the physical and logical structure of Active Directory, which provides the framework for policy application in the enterprise.

In Active Directory, sites represent the physical structure of your network. A site is a group of TCP/IP subnets, with each subnet representing a physical network segment. A domain is a logical grouping of objects for centralized management, and subgroups within a domain are called organizational units. Your network might have sites called NewYorkMain, CaliforniaMain, and WashingtonMain. Within the WashingtonMain site, you could have domains called SeattleEast, SeattleWest, SeattleNorth, and SeattleSouth. Within the SeattleEast domain, you could have organizational units called Information Services (IS), Engineering, and Sales.

Group Policy settings are stored in a Group Policy Object (GPO). You can think of a GPO as a container for the policies you apply and their settings. You can apply multiple GPOs to a single site, domain, or organizational unit. Because Group Policy is described by using objects, many object-oriented concepts apply. If you know a bit about object-oriented programming, you might expect the concepts of parent-child relationships and inheritance to apply to GPOs—and you'd be right.

A container is simply an object that contains other objects. Through inheritance, a policy applied to a parent container is inherited by a child container. Essentially, this means that a policy setting applied to a parent object is passed down to a child object. For example, if you apply a policy setting in a domain, the setting is inherited by organizational units within the domain. In this case, the GPO for the domain is the parent object, and the GPOs for the organizational units are the child objects.

The order of inheritance flows from site to domain to organizational unit. This means that the Group Policy settings for a site are passed down to the domains within that





site, and the settings for a domain are passed down to the organizational units within that domain.

To change the way policy is applied, you can override inheritance. When you do this, you assign a policy setting for a child container that is different from the policy setting for the parent and prevent the child container from inheriting the settings of a parent object. As long as overriding the policy is allowed (that is, overriding isn't blocked), the child container's policy setting will be applied appropriately. To learn more about overriding and blocking GPOs, see "Blocking, Overriding, and Disabling Policies" in Chapter 15.





# Understanding Processing Order

The processing order of policies is important. Policies are applied in the following order:

1. Local group policies
2. Site group policies
3. Domain group policies
4. Organizational unit group policies
5. Child organizational unit group policies

Any time policy settings conflict the settings applied later have precedence and overwrite previously applied policy settings. For example, organizational unit policies have precedence over domain group policies. As you might expect, there are exceptions

to the precedence rule. These exceptions are discussed in section “Blocking, Overriding, and Disabling Policies” in Chapter 15.



# Understanding Policy Application

Policy settings are divided into two broad categories:

- Those that apply to computers
- Those that apply to users

Computer policies are normally applied during system startup, and user policies are normally applied during logon. The exact sequence of events is often important in troubleshooting system behavior. The events that take place during startup and logon are as follows:

1. The network starts, and then Windows Server applies computer policies. By default, computer policies are applied one at a time in the previously specified order. No user interface is displayed while computer policies are being processed.
2. Windows Server runs startup scripts. By default, startup scripts are executed one at a time, with each completing or timing out before the next one starts. Script execution isn't displayed to the user unless specified.
3. A user logs on. After the user is validated, Windows Server loads the user profile.
4. Windows Server applies user policies. By default, user policies are applied one at a time in the previously specified order. The user interface is displayed while user policies are being processed.
5. Windows Server runs logon scripts. Logon scripts for Group Policy are executed simultaneously by default. Script execution isn't displayed to the user unless specified. Scripts in the Netlogon share run last in a normal command shell window.
6. Windows Server displays the start shell interface configured in Group Policy.
7. By default, Group Policy is refreshed when a user logs off or a computer is restarted automatically within a 90 to 120 minute period. You can change this behavior by setting a Group Policy refresh interval, as discussed in the section "Refreshing Group Policy" later in Chapter 15. To do this, open a command prompt and enter **gpupdate** .

It's important to keep in mind that some user settings, such as Folder Redirection, can't be updated when a user is logged on. The user must log off and then log back on for these settings to be applied. You can enter **gpupdate /logoff** at a command prompt or in the Search box to log off the user automatically after the refresh. Similarly, some computer settings can be updated only at startup. The computer must be restarted for





these settings to be applied. You can enter **gpupdate /boot** at a command prompt or in the Search box to restart the computer after the refresh.



# Understanding Policy Versioning


Each new version of the Windows operating system has brought with it changes to Group Policy. Sometimes these changes have made earlier policies obsolete on newer versions of Windows. In this case, the policy is stated as only working on specific versions of Windows, such as only on Windows 7 and Windows Server 2008 R2.

Generally, policies are forward-compatible but not backward-compatible. This means policies introduced in a previous release of Windows typically can be used in the current release of Windows, while policies introduced in the current release usually aren't applicable to earlier releases. If a policy isn't applicable to a particular version of the Windows operating system, you can't enforce the policy on computers running those versions of Windows.

The Properties dialog box for each policy setting has a Supported On box. This textonly field lists the policy's compatibility with various versions of Windows.



You don't have to open it if you select a policy in any of the Group Policy editors and also have selected the Extended tab (rather than the Standard tab). A Requirements entry is displayed that lists compatibility.



## Managing Local Group Policies

As long as the computer is not a domain controller, computers running current versions of Windows allow the use of multiple local Group Policy Objects on a single computer. This allows you to assign a different local GPO to each local user or



general user type, supporting a more flexible, wider array of implementation scenarios.

With computers in a stand-alone configuration rather than a domain configuration, you might find that multiple local GPOs are useful because you no longer have to explicitly disable or remove settings that interfere with your ability to manage a computer before performing administrator tasks. Instead, you can implement one local GPO for administrators and another local GPO for nonadministrators. In a domain configuration, however, you might not want to use multiple local GPOs. In domains, most computers and users already have multiple GPOs applied to them—adding multiple local GPOs to this already varied mix can make managing Group Policy confusing.





## Using Group Policy Objects

Windows supports three layers of local Group Policy Objects:

- **Local Group Policy** Local Group Policy is the only local Group Policy Object that allows both computer configuration and user configuration settings to be applied to all users of the computer.
- **Administrators and Non-Administrators local Group Policy** Administrators and Non-Administrators Local Group Policy contains only user configuration settings. This policy is applied based on whether the user account being used is a member of the local Administrators group.
- **User-specific local Group Policy** User-Specific Local Group Policy contains only user configuration settings. This policy is applied to individual users and groups.

These layers of local Group Policy Objects are processed in the following order:

1. Local Group Policy.
2. Administrators and Non-Administrators Local Group Policy.
3. User-Specific Local Group Policy.

When multiple local policy objects are used, Windows resolves conflicts in settings by overwriting any previous setting with the last-read and most-current settings. Thus, the final setting is the one that Windows uses.

**NOTE** In resolving conflicts, only the enabled or disabled state of settings matters. A setting set as Not Configured has no effect on the state of the setting from a previous policy application.

**REAL WORLD** To simplify domain administration, you can disable processing of local Group Policy Objects by enabling the Turn Off Local Group Policy Objects Processing policy setting in a domain GPO. In Group Policy, this setting is located in the Administrative Templates for Computer Configuration under System\Group Policy.

## Accessing the Top-Level Local Policy Settings


All computers running current releases of Windows have an editable local Group Policy Object. That said, although a domain controller has a local Group Policy Object, you shouldn't edit its settings. You manage policy for domain controllers using default policy, as discussed in "Understanding Default Policies" later in this chapter.

A quick way to access the local GPO on a computer is to enter the following command in the Search box.



gpedit.msc

This command starts the GPOE in a Microsoft Management Console (MMC) with its target set to the local computer.




To access the top-level local GPO on a remote computer, enter the following in the Search box.

gpedit.msc /gpcomputer: RemoteComputer

Here RemoteComputer is the host name or fully qualified domain name of the remote computer. Note the colon, which is required, as shown in the following example.

gpedit.msc /gpcomputer: corpsvr82

When you are connected to a remote computer, the root node lists the name of the remote computer. Otherwise, the root node label is shown as Local Computer Policy.




You can also manage the top-level local GPO on a computer by following these steps:

1. At a command prompt or in the Search box, type **mmc.exe**, and then press Enter.

2. In the Microsoft Management Console, click File, and then click Add/Remove Snap-In. This opens the Add Or Remove Snap-Ins dialog box.



3. In the Add Or Remove Snap-Ins dialog box, click Group Policy Object Editor, and then click Add.
4. In the Select Group Policy Object dialog box, click Finish because the local computer is the default object, and then click OK.



**TIP** You can use the same MMC snap-in to manage more than one local Group Policy Object. In the Add Or Remove Snap-Ins dialog box, you simply add one instance of the Local Group Policy Object Editor for each object you want to work with.

Whether you access the Local GPO directly or by adding GPOE to an MMC, you are then able to manage local policy settings by using the options provided. On computers where local group policies are applied, you'll find the following subfolders in the %SystemRoot%\System32\GroupPolicy folder:

- **Machine** Stores computer scripts in the Script folder and registry-based policy information for HKEY\_LOCAL\_MACHINE (HKLM) in the Registry.pol file
- **User**

Stores user scripts in the Script folder and registry-based policy information for HKEY\_CURRENT\_USER (HKCU) in the Registry.pol file


The files stored in these folders contain details about configured settings.

**CAUTION** Don't edit these folders and files directly. Instead, you should use the appropriate features of one of the Group Policy management tools. By default, these files and folders are hidden. If you want to view hidden files and folders in File Explorer, click the View tab, and then select Hidden Items. You also might want to select File Name Extensions.

## Accessing Tailored Local Group Policies


The only local policy object that exists by default on a computer is the Local Group Policy Object. You can create and manage other local objects as necessary (except on domain controllers) by following these steps:

1. At a command prompt or in the Search box, type **mmc.exe**, and then press Enter. In the Microsoft Management Console, click File, and then click Add/Remove Snap-In. This opens the Add/Remove Snap-In dialog box.




2. In the Add Or Remove Snap-Ins dialog box, click Group Policy Object Editor, and then click Add.





3. In the Select Group Policy Object dialog box, click Browse. In the Browse For A Group Policy Object dialog box, click the Users tab.



4. On the Users tab, the entries in the Group Policy Object Exists column specify whether a particular local policy object has been created. Do one of the following:
  - Select Administrators to create or access the Administrators Local Group Policy Object.
  - Select Non-Administrators to create or access the Non-Administrators Local Group Policy Object.
  - Select the local user whose user-specific local Group Policy Object you want to create or access.
5. Click OK. If the selected object doesn't exist, it will be created. Otherwise, the existing object opens for review and editing.





Policy settings for administrators, nonadministrators, and users are stored in the %SystemRoot%\System32\GroupPolicyUsers folder on each computer running Windows Server. Because these local GPOs apply only to user configuration settings, user-specific policy settings under %SystemRoot%\System32\GroupPolicyUsers have only a User subfolder, and this subfolder stores user scripts in the Script folder and registry-based policy information for HKEY\_CURRENT\_USER in the Registry.pol file.


## Navigating Enterprise Policies

In the enterprise (AD DS), you can use Active Directory–based Group Policy any time you deploy Active Directory Domain Services. Each site, domain, and organizational unit can have one or more group policies. Group policies listed higher in the Group Policy list have higher precedence than policies listed lower in the list. This ensures that policies are applied appropriately throughout the related sites, domains, and organizational units.



# Getting Started with GPMC

Microsoft provides the Group Policy Management Console (GPMC) as an installable feature for centrally managing Group Policy. Add this feature to any installation of Windows Server by using the Add Roles And Features Wizard. The GPMC also is available on Windows desktops when you install the Remote Server Administration Tools (RSAT). After you add the GPMC to a computer, it is available on the Tools menu in Server Manager. At a prompt or in the Search box, type **gpmc.msc**, and then press Enter.



In GPMC, the console root node is labeled Group Policy Management, and below this node is the Forest node. The Forest node represents the forest to which you are currently connected and is named after the forest root domain for that forest. If you have appropriate credentials, you can add connections to other forests. To do this, right-click the Group Policy Management node, and then click Add Forest. In the Add Forest dialog box, enter the name of the forest root domain in the Domain box, and then click OK.

Expanding the Forest node displays the following nodes:

- **Domains** Provides access to the policy settings for domains in the related forest. You are connected to your logon domain by default. If you have appropriate credentials, you can add connections to other domains in the related forest. To do this, right-click the Domains node, and then click Show Domains. In the Show Domains dialog box, select the check boxes for the domains you want to add, and then click OK.
- **Sites** Provides access to the policy settings for sites in the related forest. Sites are hidden by default. If you have appropriate credentials, you can add connections for



sites. To do this, right-click the Sites node, and then click Show Sites. In the Show Sites dialog box, select the check boxes for the sites you want to add, and then click OK.

- **Group Policy Modeling** Provides access to the Group Policy Modeling Wizard, which helps you plan policy deployment and simulate settings for testing purposes. Any saved policy models are also available.
- **Group Policy Results** Provides access to the Group Policy Results Wizard. For each domain you are connected to, all related GPOs and organizational units are available to work with in one location.

GPOs listed under the domain, site, and organizational unit containers in the GPMC are

GPO links and not the GPOs themselves. You can access the actual GPOs through the Group Policy Objects container of the selected domain. Note that the icons for GPO links have small arrows at the bottom left, similar to shortcut icons, whereas GPOs themselves do not.

When you start the GPMC, the console connects to Active Directory running on the domain controller that is acting as the PDC emulator for your logon domain and obtains a list of all GPOs and organizational units in that domain. It does this by using Lightweight Directory Access Protocol (LDAP) to access the directory store and the Server Message Block (SMB) protocol to access the SYSVOL directory. If the PDC emulator isn't available for some reason when you start the GPMC, such as when the server is offline, the GPMC displays a prompt so that you can choose to work with policy settings on the domain controller you are currently connected to or on any available domain controller.

You can manually change the domain controller you are connected to as well. To do this, right-click the domain node for which you want to set the domain controller focus, and then click Change Domain Controller. In the Change Domain Controller dialog box, the domain controller you are currently connected to is listed under Current Domain Controller. Using the Change To options, specify the domain controller to use, and then click OK.






## Using Policy Editors

When you want to edit a GPO in the GPMC, the GPMC opens the Group Policy Management Editor, which you use to manage the policy settings. If Microsoft had stopped with these two tools, we'd have a wonderful and easy-to-use policymanagement environment. Unfortunately, several other, nearly identical editors also exist:

- **Group Policy Starter GPO Editor** An editor you can use to create and manage starter policy objects. As the name implies, starter GPOs are meant to provide a starting point for policy objects you'll use throughout your organization. When you create a policy object, you can specify a starter GPO as the source or basis of the object.
- **Local Group Policy Object Editor** An editor you can use to create and manage policy objects for the local computer. As the name implies, local GPOs are meant to provide policy settings for a specific computer as opposed to settings for a site, domain, or organizational unit.





Group Policy Management Editor, Group Policy Starter GPO Editor, and Local Group Policy Object Editor are essentially identical except for the set of policy objects you have access to. For this reason, and because you use these tools to manage individual policy objects in the same way, I won't differentiate between them unless necessary. As a matter of preference, I refer to these tools collectively as policy editors . Sometimes, I might use the acronym GPOE to refer to policy editors in general because it is more easily distinguished from the management console, the GPMC.

## Understanding Default Policies

Each domain in your organization has two default GPOs:

- **Default Domain Controllers Policy GPO** A default GPO created for and linked to the Domain Controllers organizational unit. This GPO is applicable to all domain controllers in a domain (as long as they aren't moved from this organizational unit). Use it to manage security settings for domain controllers in a domain.
- **Default Domain Policy GPO** A default GPO created for and linked to the domain itself within Active Directory. Use this GPO to establish baselines for a wide variety of policy settings that apply to all users and computers in a domain.


Although these GPOs aren't meant for general management of Group Policy, they can be used to manage specific areas of policy that you can't manage using other GPOs. Specifically, you use the Default Domain Controllers Policy GPO to manage specific User Rights Assignment and Security Options settings that limit the ways domain controllers can be used. You use the Default Domain Policy GPO to manage default Account Policies settings and, in particular, three specific areas of Account Policies: password policy, account lockout policy, and Kerberos policy.

**NOTE** Several security options are also managed through the Default Domain Policy GPO. These include Accounts: Rename Administrator Account, Accounts: Administrator Account Status, Accounts: Guest Account Status, Accounts: Rename Guest Account, Network Security: Force Logoff When Logon Hours Expire, Network Security: Do Not Store LAN Manager Hash Value On Next Password Change, and Network Access: Allow Anonymous SID/Name Translation.

The Default Domain Policy GPO typically is the highest-precedence GPO linked to the domain level, and the Default Domain Controllers Policy GPO is the highestprecedence GPO linked to the Domain Controllers container. You can link additional GPOs to the domain level and to the Domain Controllers container. When

you do this, the settings in the highest-precedence GPO override settings in lower-precedence GPOs. Thus:

- One way to override settings applied in the Default Domain Policy GPO is to create a GPO with the overriding settings and link it with a higher precedence to the domain container.
- One way to override settings applied in the Default Domain Controllers Policy GPO is to create a GPO with the overriding settings and link it with a higher precedence to the Domain Controllers container.



To manage any other areas of policy, you should create a GPO and link it to the domain or to an appropriate site or organizational unit within the domain. Site, domain, and organizational unit group policies are stored in the %SystemRoot%\Sysvol\Domain\Policies folder on domain controllers. In this folder, you'll find one subfolder for each policy you defined on the domain controller. The policy folder name is the policy's globally unique identifier (GUID). You can find the policy's GUID on the policy's Properties page on the General tab in the Summary frame.

Within these individual policy folders, you'll find the following subfolders:

- **Machine** Stores computer scripts in the Script folder and registry-based policy information for HKEY\_LOCAL\_MACHINE (HKLM) in the Registry.pol file
- **User** Stores user scripts in the Script folder and registry-based policy information for HKEY\_CURRENT\_USER (HKCU) in the Registry.pol file





Do not edit these folders or the related files directly. Instead, use the appropriate features of one of the Group Policy management tools.

## Understanding Registry-based Settings

In Group Policy, the administrative templates are where registry-based policy settings are stored. Current releases of Windows use an XML-based administrative templates format called ADMX. Because ADMX files are created by using XML, the files are strictly structured and can be easily and rapidly parsed during initialization of the policy-processing engine. This can help to improve performance when the operating system processes Group Policy during the startup, logon, logoff, and shutdown phases, in addition to during policy refreshes.

ADMX files are divided into language-neutral files ending with the .admx file name extension and language-specific files ending with the .adml extension. The languageneutral files ensure that a GPO has identical core policies. The language-specific files allow policies to be viewed and edited in multiple languages. Because the languageneutral files store a policy's core settings, policies can be edited in any language for which a computer is configured, thus allowing one user to view and edit policies in English and another to view and edit policies in Spanish, for example. The mechanism that determines the language used is the language pack installed on the computer.

Language-neutral ADMX files are stored in the %SystemRoot%\PolicyDefinitions folder while language-specific ADMX files are stored in the %SystemRoot%\PolicyDefinitions\ LanguageCulture folder. Each subfolder is named by using the corresponding International Organization for Standardization (ISO) language/culture name, such as EN-US for U.S. English.

When you start a policy editor, it automatically reads ADMX files from the policy definitions folders. Because of this, you can copy ADMX files you want to use to an appropriate policy definitions folder to make them available when you are editing GPOs. If the policy editor is running when you copy the file or files, you must restart the policy editor to force it to read the file or files.


In domains, ADMX files can be stored in a central store—the domainwide directory created in the SYSVOL directory (%SystemRoot%\Sysvol\Domain\Policies). When you use a central store, administrative templates are no longer stored with each GPO. Instead, only the current state of the setting is stored in the GPO, and the ADMX files are stored centrally. This reduces the amount of storage space used as the number of

GPOs increases and also reduces the amount of data being replicated throughout the enterprise.

## Editing GPOs

With the GPMC, you can edit a GPO by right-clicking it and then selecting Edit on the shortcut menu. The policy editor has two main nodes:

- **Computer Configuration** Allows you to set policies that should be applied to computers, regardless of who logs on
- **User Configuration** Allows you to set policies that should be applied to users, regardless of which computer they log on to



Under the Computer Configuration and User Configuration nodes, you'll find the Policies and Preferences nodes. Settings for general policies are listed under the Policies node. Settings for general preferences are listed under the Preferences node.

The exact configuration of Computer Configuration and User Configuration depends on the add-ons installed and which type of policy you're creating. Still, you'll usually find that both Computer Configuration and User Configuration have subnodes for the following:

- **Software Settings** Sets policies for software settings and software installation. When you install software, subnodes might be added to Software Settings.
- **Windows Settings** Sets policies for folder redirection, scripts, and security.
- **Administrative Templates** Sets policies for the operating system, Windows components, and programs. Administrative templates are configured through template files. You can add or remove template files whenever you need to.

A complete discussion of all the available options is beyond the scope of this book. The sections that follow focus on using administrative templates.



# Using Administrative Templates to Set Policies

Administrative templates provide easy access to registry-based policy settings that you might want to configure. A default set of administrative templates is configured for users and computers in the policy editor. You can also add or remove administrative templates. Any changes you make to policies available through administrative templates are saved in the registry. Computer configurations are saved in HKEY\_LOCAL\_MACHINE, and user configurations are saved in HKEY\_CURRENT\_USER.


You can view the currently configured templates in the Administrative Templates node of the policy editor. This node contains policies you can configure for local systems, organizational units, domains, and sites. Different sets of templates are found under Computer Configuration and User Configuration. You can add templates containing new policies in the policy editor when you install new Windows components.

You can use administrative templates to manage the following:

- **Control Panel** Determine the available options and configuration of Control Panel and Control Panel utilities
- **Desktop** Configure the Windows desktop and the available options from the desktop
- **Network** Configure networking and network client options for offline files, DNS clients, and network connections
- **Printers** Configure printer settings, browsing, spooling, and directory options
- **Shared folders** Allow publishing of shared folders and Distributed File System (DFS) roots
- **Start screen and taskbar** Control the available options and configuration of the Start screen and the taskbar
- **System** Configure system settings for disk quotas, user profiles, user logon, system restore, error reporting, and so on
- **Windows components** Determine the available options and configuration of various Windows components, including Event Viewer, Internet Explorer, Task Scheduler, Windows Installer, and Windows Updates

The best way to get to know which administrative template policies are available is to browse the Administrative Templates nodes. As you browse the templates, you'll find that policies are in one of three states:

- **Not Configured** The policy isn't used, and no settings for it are saved in the registry.
- **Enabled** The policy is actively being enforced, and its settings are saved in the registry.
- **Disabled** The policy is turned off and isn't enforced unless overridden. This setting is saved in the registry.



You can enable, disable, and configure policies by following these steps:

1. In the policy editor, open the Administrative Templates folder in the Computer Configuration or User Configuration node, whichever is appropriate for the type of policy you want to set.
2. In the left pane, select the subfolder containing the policies you want to work with. The related policies are then displayed in the right pane.
3. Double-click a policy to display its related Properties dialog box.

You can read a description of the policy in the Help pane. The description is available only if one is defined in the related template file.

4. To set the policy's state, select one of the following options:

- **Not Configured** The policy isn't configured.
- **Enabled** The policy is enabled.
- **Disabled** The policy is disabled.

5. If you enable the policy, set any additional parameters, and then click OK.

**NOTE** Normally, computer policies have precedence in Windows Server. If there's a conflict between a computer policy setting and a user policy setting, the computer policy is enforced.

## Creating and Linking GPOs

Creating a Group Policy object and linking an object to a specific container within Active Directory are two different actions. You can create a GPO without linking it to any domain, site, or organizational unit. Then, as appropriate, you can link the GPO to




a specific domain, site, or organizational unit. You can also create a GPO and link it automatically to a domain, site, or organizational unit.


The technique you use for creating and linking a GPO primarily depends on your personal preference and how you plan to work with the GPO. Keep in mind that when you create and link a GPO to a site, domain, or organizational unit, the GPO is applied to the user and computer objects in that site, domain, or organizational unit according to the Active Directory options governing inheritance, the precedence order of GPOs, and other settings.

To create and then link a GPO to a site, domain, or organizational unit follow these steps:

1. In the GPMC, right-click the Group Policy Objects container, and then click New. In the New GPO dialog box, enter a descriptive name for the GPO, such as **Secure Workstation GPO** . If you want to use a starter GPO as the source for the initial settings, select the starter GPO to use in the Source Starter GPO list. When you click OK, the new GPO is added to the Group Policy Objects container.




2. In the GPMC, right-click the new GPO, and then click Edit. In the policy editor, configure the necessary policy settings, and then close the policy editor.
3. In the GPMC, select the site, domain, or organizational unit to which you want to link the GPO. In the right pane, the Linked Group Policy Objects tab shows the GPOs that are currently linked to the selected container (if any).



4. Right-click the site, domain, or organizational unit to which you want to link the GPO, and then click Link An Existing GPO. In the Select GPO dialog box, select the GPO you want to link with, and then click OK. When Group Policy




is refreshed for computers and users in the applicable site, domain, or organizational unit, the policy settings in the GPO are applied.



To create and link a GPO as a single operation, follow these steps:

5. In the GPMC, right-click the site, domain, or organizational unit for which you want to create and link the GPO, and then click **Create A GPO In This Domain, And Link It Here**.
6. In the New GPO dialog box, enter a descriptive name for the GPO, such as **Secure Workstation GPO** . If you want to use a starter GPO as the source for the initial settings, select the starter GPO to use in the Source Starter GPO list. When you click OK, the new GPO is added to the Group Policy Objects container and linked to the previously selected site, domain, or organizational unit.




7. Right-click the new GPO, and then click **Edit**. In the policy editor, configure the necessary policy settings, and then close the policy editor. When Group Policy is refreshed for computers and users in the applicable site, domain, or organizational unit, the policy settings in the GPO are applied.

# Using Starter GPOs

When you create a GPO in the GPMC, you can base the GPO on a starter GPO. The settings for the starter GPO are then imported into the new GPO, which allows you to use a starter GPO to define the base configuration settings for a new GPO. In a large organization, you should create different categories of starter GPOs based on the users and computers they will be used with or on the required security configuration.


You can create a starter GPO by following these steps:

1. In the GPMC, right-click the Starter GPOs container, and then click New. In the New Starter GPO dialog box, enter a descriptive name for the GPO, such as **General Management GPO** . You can also enter comments describing the GPO's purpose. Click OK.



The screenshot shows the 'New Starter GPO' dialog box. It has a title bar with a close button (X). Inside, there is a 'Name:' label followed by a text box containing 'General Management GPO'. Below that is a 'Comment:' label followed by a larger text box. At the bottom right, there are 'OK' and 'Cancel' buttons. The 'OK' button is highlighted with a blue border.

2. Right-click the new GPO, and then click Edit. In the policy editor, configure the necessary policy settings, and then close the policy editor.
3. After creating the starter GPO and configuring its settings, you can use the GPO as the basis of new GPOs. When you are working with the New GPO dialog box, specify that this GPO is the Source Starter GPO.



The screenshot shows the 'New GPO' dialog box. It has a title bar with a close button (X). Inside, there is a 'Name:' label followed by a text box containing 'New Group Policy Object'. Below that is a 'Source Starter GPO:' label followed by a dropdown menu showing 'General Management GPO'. At the bottom right, there are 'OK' and 'Cancel' buttons. The 'OK' button is highlighted with a blue border.



# Delegating Group Policy Management


All administrators have some level of privileges for performing Group Policy management tasks. Through delegation, other individuals can be granted permissions to perform any or all of the following tasks:

- Create GPOs and manage the GPOs they create
- View settings, modify settings, delete a GPO, and modify security
- Manage links to existing GPOs or generate Resultant Set of Policy (RSoP)




## Checking Delegations for Creating GPOs

Administrators can create GPOs, and anyone who has created a GPO has the right to manage that GPO. In the GPMC, you can determine who can create GPOs in a domain by selecting the Group Policy Objects node for that domain and then clicking the Delegation tab. On the Delegation tab is a list of groups and users that can create GPOs in the domain.



To grant GPO creation permission to a user or group, click Add. In the Select User, Computer, Or Group dialog box, select the user or group that should have the permission, and then click OK.




## Checking Delegations for Management

Users can also be assigned permission for Group Policy management. To determine delegations for domain, site, and organizational unit permissions, select the domain, site, or organizational unit you want to work with, and then click the Delegation tab in the right pane. In the Permission list, select the permission you want to check. The options are as follows:

- **Link GPOs** Lists users and groups that can create and manage links to GPOs in the selected site, domain, or organizational unit




- **Perform Group Policy Modeling Analyses** Lists users and groups that can determine RSoP for the purposes of planning
- **Read Group Policy Results Data** Lists users and groups that can determine RSoP that is currently being applied, for the purposes of verification or logging



To grant any of these permissions, complete the following steps:

1. In the GPMC, select the domain, site, or organizational unit you want to work with, and then click the Delegation tab in the right pane.
2. In the Permission list, select the permission you want to grant. The options are Link GPOs, Perform Group Policy Modeling Analyses, and Read Group Policy Results Data.
3. Click Add. In the Select User, Computer, Or Group dialog box, select the user or group, and then click OK.
4. In the Add Group Or User dialog box, specify how the permission should be applied. To apply the permission to the current container and all child containers, select This Container And All Child Containers. To apply the permission only to the current container, select This Container Only. Click OK.








# Checking Delegations for a GPO

Individual GPOs have delegations as well. For individual GPO permissions, select the GPO you want to work with in the GPMC, and then click the Delegation tab in the right pane. One or more of the following permissions are displayed for individual users and groups:

- **Read** Indicates that the user or group can view the GPO and its settings. **Edit**
- **Settings** Indicates that the user or group can view the GPO and change its settings. The user or group cannot delete the GPO or modify security.
- **Edit Settings, Delete, Modify Security** Indicates that the user or group can view the GPO and change its settings. The user or group can also delete the GPO and modify security.



To grant permissions for working with the GPO, complete the following steps:

1. In the GPMC, select the GPO you want to work with, click the Delegation tab in the right pane, and then click Add.
2. In the Select User, Computer, Or Group dialog box, select the user or group, and then click OK.
3. In the Add Group Or User dialog box, select the permission level to assign, such as Read, and then click OK.

## Chapter 15. Maintaining and Troubleshooting Group Policy

If Group Policy is broken, jump ahead to the section of this chapter on troubleshooting for assistance with fixing the problem. But if Group Policy is running reasonably well stay right here and continue reading to learn how to keep it running that way for as long as possible.





This chapter is a checklist of sorts, with each section explaining an essential and necessary task to keep Group Policy running at its best. You'll learn how to block, override and disable policies, for example. You'll also learn how to back up and restore Group Policy objects and more.





# Optimizing Group Policy

Windows Server provides a variety of features for keeping Group Policy running smoothly. Some allow you to change the way inheritance works or ensure the latest changes made to policies are applied as quickly as possible. Others help you model policy for planning or to determine the settings that are being applied to a given user or computer.



# Blocking, Overriding, and Disabling Policies

Inheritance ensures that every computer and user object in a domain, site, or organizational unit is affected by Group Policy. Most policies have three configuration options: Not Configured, Enabled, or Disabled. Not Configured is the default state for most policy settings. If a policy is enabled, the policy is enforced and is applied directly or through inheritance to all users and computers that are subject to the policy. If a policy is disabled, the policy is not enforced or applied.

Windows provides four key ways to change the way inheritance works:


- [Change the link order and precedence](#)
- [Override inheritance \(as long as there is no enforcement\)](#)
- [Block inheritance \(to prevent inheritance completely\)](#)
- [Enforce inheritance \(to supersede and prevent overriding or blocking\)](#)

The order of inheritance goes from the site level to the domain level and then to each nested organizational unit level. When multiple policy objects are linked to a particular level, the link order determines the order in which policy settings are applied. Linked policy objects are always applied in link-ranking order. Lower-ranking policy objects are processed first, and then higher-ranking policy objects are processed. The policy object processed last has priority, so any policy settings configured in this policy object are final and override those of other policy objects (unless you use inheritance blocking or enforcing).

When multiple policy objects can be inherited from a higher level, the precedence order shows exactly how policy objects are being processed. As with link order, lower-ranking policy objects are processed before higher-ranking policy objects. The policy object processed last has precedence, so any policy settings configured in this policy object are final and override those of other policy objects (unless you use inheritance blocking or enforcing).

You can change the link order (and thus the precedence order) of policy objects by following these steps:

1. In the GPMC, select the container for the site, domain, or organizational unit with which you want to work.
2. In the right pane, select the Linked Group Policy Objects tab. Click the policy object you want to work with.



3. Click the Move Link Up or Move Link Down buttons as appropriate to change the link order of the selected policy object.
4. When you are finished changing the link order, confirm that policy objects are being processed in the expected order by checking the precedence order on the Group Policy Inheritance tab.

You can override inheritance to change the way inheritance works. When a policy is enabled in a higher-level policy object, you can override inheritance by disabling the policy in a lower-level policy object. When a policy is disabled in a higher-level policy object, you can override inheritance by enabling the policy in a lower-level policy object. As long as a policy is not blocked or enforced, this technique achieves the effects you want.

Rather than overriding inheritance, you may want to block inheritance so that no policy settings from higher-level containers are applied to users and computers in a particular container. When inheritance is blocked, only configured policy settings from policy objects linked at that level are applied, and settings from all high-level containers are blocked (as long as there is no policy enforcement).

Domain administrators can use inheritance blocking to block inherited policy settings from the site level. Organizational unit administrators can use inheritance blocking to block inherited policy settings from both the domain and the site level. By using blocking to ensure the autonomy of a domain or organizational unit, you can ensure that domain or organizational unit administrators have full control over the policies that apply to users and computers under their administration.

In the GPMC, you can block inheritance by right-clicking the domain or organizational unit that should not inherit settings from higher-level containers, and then selecting Block Inheritance. If Block Inheritance is already selected, selecting it again removes the setting. When you block inheritance in the GPMC, a blue circle with an exclamation point is added to the container's node in the console tree. This notification icon provides



a quick way to tell whether any domain or organizational unit has the Block Inheritance setting enabled.



You can prevent administrators who have authority over a container from overriding or blocking inherited Group Policy settings by enforcing inheritance. When inheritance is enforced, all configured policy settings from higher-level policy objects are inherited and applied regardless of the policy settings configured in lower-level policy objects. Thus, enforcement of inheritance is used to supersede the overriding and blocking of policy settings.

Forest administrators can use inheritance enforcement to ensure that configured policy settings from the site level are applied and to prevent the overriding or blocking of policy settings by domain and organizational unit administrators. Domain administrators can use inheritance enforcement to ensure that configured policy settings from the domain level are applied and to prevent the overriding or blocking of policy settings by organizational unit administrators.

In the GPMC, you can enforce policy inheritance by expanding the top-level container from which to begin enforcement, right-clicking the link to the GPO, and then clicking Enforced. For example, if you want to ensure that a domain-level GPO is inherited by all organizational units in the domain, expand the domain container, right-click the domain-level GPO, and then click Enforced. If Enforced is already selected, selecting it again removes the enforcement.


In the GPMC, you can easily determine which policies are inherited and which policies are enforced. When you enforce inheritance, a lock symbol is added to the GPO link. This notification icon provides a quick way to tell whether any GPO has the Enforced setting enabled.





**NOTE** The lock symbol is only displayed on the GPO link. If you are working with the actual GPOs in the Group Policy Objects container, the symbol isn't displayed. Here, simply select a policy object in the console tree or double-click it in the main pane, and then view the related Scope tab in the right pane. If the policy object is enforced, the Enforced value is set to Yes. Otherwise, the Enforced value is set to No.

After you select a policy object, you can right-click a location entry on the Scope tab to display a shortcut menu that allows you to manage linking and policy enforcement. Enable or disable links by selecting or clearing the Link Enabled option, respectively. Enable or disable enforcement by selecting or clearing the Enforced option.





## Refreshing Group Policy

Any time you make changes to a policy those changes are immediate. However, they aren't propagated automatically. Client computers request policies at the following times:

- When the computer starts
- When a user logs on
- When an application or user requests a refresh
- When a refresh interval is set for Group Policy and the interval has elapsed

Computer configuration settings are applied during startup of the operating system. User configuration settings are applied when a user logs on to a computer. Typically, if there is a conflict between computer and user settings, computer settings have priority and take precedence.

After policy settings are applied, the settings are refreshed automatically to ensure that they are current. The default refresh interval for domain controllers is 5 minutes. For all other computers, the default refresh interval is 90 minutes, with up to a 30-minute variation to avoid overloading the domain controller with numerous concurrent client requests. This means that an effective refresh window for nondomain-controller computers is 90 to 120 minutes.

During a Group Policy refresh, the client computer contacts an available domain controller in its local site. If one or more of the policy objects defined in the domain have changed, the domain controller provides a list of the policy objects that apply to the computer and to the user who is currently logged on, as appropriate. The domain controller does this regardless of whether the version numbers on all the listed policy objects have changed. By default, the computer processes the policy objects only if the version number of at least one of the policy objects has changed. If any one of the related policies has changed, all the policies have to be processed again because of inheritance and the interdependencies between policies.

Security settings are a notable exception to the processing rule. By default, these settings are refreshed every 16 hours (960 minutes) regardless of whether policy objects contain changes. A random offset of up to 30 minutes is added to reduce the impact on domain controllers and the network during updates (making the effective refresh window 960 to 990 minutes). Also, if the client computer detects that it is connecting over a slow network connection, it informs the domain controller, and only the security settings and administrative templates are transferred over the network. This means that by default, only the security settings and administrative templates are






applied when a computer is connected over a slow link. You can configure the way slow-link detection works in Group Policy.

You must carefully balance the update frequency with the actual rate of policy change. If policy is changed infrequently, you might want to increase the refresh window to reduce resource usage. For example, you might want to use a refresh interval of 20 minutes on domain controllers and 180 minutes on other computers.

## Configuring the Refresh Interval

You can change the Group Policy refresh interval on a per-policy object basis. To set the refresh interval for domain controllers, follow these steps:

1. In the GPMC, right-click the Group Policy object you want to modify, and then click Edit. This GPO should be linked to a container that contains domain controller computer objects.
2. In the Administrative Templates for Computer Configuration under System\Group Policy, double-click the Set Group Policy Refresh Interval For Domain Controllers policy. This displays a Properties dialog box for the policy.
3. Define the policy by selecting Enabled.



4. Set the base refresh interval in the first Minutes box. You usually want this value to be between 5 and 59 minutes.
5. In the other Minutes box, set the minimum or maximum time variation for the refresh interval. The variation effectively creates a refresh window with the goal of avoiding overload resulting from numerous clients simultaneously requesting a Group Policy refresh. Click OK.

**NOTE** A faster refresh rate increases the likelihood that a computer has the most current policy configuration. A slower refresh rate reduces the frequency of policy refreshes, which can reduce overhead with regard to resource usage but also increase the likelihood that a computer won't have the most current policy configuration.

Options:

This setting allows you to customize how often Group Policy is applied to domain controllers. The range is 0 to 44640 minutes (31 days).

Minutes:

This is a random time added to the refresh interval to prevent all clients from requesting Group Policy at the same time. The range is 0 to 1440 minutes (24 hours)

Minutes:

To set the refresh interval for member servers and workstations, follow these steps:

1. In the GPMC, right-click the Group Policy Object you want to modify, and then click Edit. This GPO should be linked to a container that contains computer objects.
2. In the Administrative Templates for Computer Configuration under System\Group Policy, double-click the Set Group Policy Refresh Interval For Computers policy.
3. Define the policy by selecting Enabled.

Set Group Policy refresh interval for computers

Set Group Policy refresh interval for computers Previous Setting Next Setting

☐ Not Configured Comment:

☒ Enabled

☐ Disabled

Supported on: At least Windows 2000

4. In the first Minutes box, set the base refresh interval. You usually want this value to be between 60 and 240 minutes.
5. In the other Minutes box, set the minimum or maximum time variation for the refresh interval. The variation effectively creates a refresh window with the goal of avoiding overload resulting from numerous clients simultaneously requesting a Group Policy refresh. Click OK.

Options:

This setting allows you to customize how often Group Policy is applied to computers. The range is 0 to 44640 minutes (31 days).

Minutes:

This is a random time added to the refresh interval to prevent all clients from requesting Group Policy at the same time. The range is 0 to 1440 minutes (24 hours)

Minutes:

**REAL WORLD** You want to be sure that updates don't occur too frequently yet are timely enough to meet expectations or requirements. The more often a policy is refreshed, the more traffic is generated over the network. In a large installation, you typically want to set a refresh rate that is longer than the default to reduce network traffic, particularly if the policy affects hundreds of users or computers. In any installation where users complain about their computers periodically being sluggish, you also might want to increase the policy refresh interval. Consider that a once-a-day or once-a-week update might be all that it takes to keep policies current enough to meet your organization's needs.

## Refreshing Policy Manually

As an administrator, you might often need or want to refresh Group Policy manually. For example, you might not want to wait for Group Policy to be refreshed at the automatic interval, or you might be trying to resolve a problem with refreshes and want to force a Group Policy refresh. You can refresh Group Policy manually by using the Gpupdate command-line utility.

You can initiate a refresh in several ways. Entering **gpupdate** at a prompt or in the Search box refreshes settings in both Computer Configuration and User Configuration on the local computer. Only policy settings that have changed are processed and applied when you run Gpupdate. You can change this behavior by using the /force parameter to force a refresh of all policy settings.



You can refresh user and computer configuration settings separately. To refresh only computer configuration settings, enter **gpupdate /target:computer** at the command prompt. To refresh only user configuration settings, enter **gpupdate /target:user** at the command prompt.

You can also use Gpupdate to log off a user or restart a computer after Group Policy is refreshed. This is useful because some group policies are applied only when a user logs on or when a computer starts. To log off a user after a refresh, add the /Logoff parameter. To restart a computer after a refresh, add the /Boot parameter.



# Modeling Group Policy for Planning

Modeling Group Policy for planning is useful when you want to test various implementation and configuration scenarios. For example, you might want to model the effect of loopback processing or slow-link detection. You can also model the effect of moving users or computers to another container in Active Directory, or the effect of changing security group membership for users and computers.

All domain and enterprise administrators have permission to model Group Policy for planning, as do those who have been delegated the Perform Group Policy Modeling Analyses permission. To model Group Policy and test various implementation and update scenarios, follow these steps:

1. In the GPMC, right-click the Group Policy Modeling node, select Group Policy Modeling Wizard, and then click Next.
2. On the Domain Controller Selection page, select the domain you want to model in the Show Domain Controllers In This Domain list. By default, you will simulate policy on any available domain controller in the selected domain. If you want to use a specific domain controller, select This Domain Controller, and then click the domain controller to use. Click Next.

The screenshot shows the 'Domain Controller Selection' page of the Group Policy Modeling Wizard. At the top, there is a dropdown menu labeled 'Show domain controllers in this domain:' with 'imaginedlands.local' selected. Below this, there is a section titled 'Process the simulation on this domain controller:' with two radio button options. The first option, 'Any available domain controller running Windows Server 2003 or later', is selected. The second option, 'This domain controller:', is unselected. Below the second option is a table listing available domain controllers.

| Name                             | Site              |
|----------------------------------|-------------------|
| CorpServer17.imaginedlands.local | Boston-First-Site |
| CorpServer28.imaginedlands.local | Boston-First-Site |

3. On the User And Computer Selection page, you have the option of simulating policy based on containers or individual accounts. Use one of the following techniques to choose accounts, and then click Next:
- Use containers to simulate changes for entire organizational units or other containers. Under User Information, select Container, and then click Browse to display the Choose User Container dialog box. Use the dialog box to choose any of the available user containers in the selected domain. Under Computer Information, select Container, click Browse to display the Choose Computer Container dialog

box, and then choose any of the available computer containers in the selected domain.

- Select specific accounts to simulate changes for a specific user and computer. Under User Information, select User, click Browse to display the Select User dialog box, and then specify a user account. Under Computer Information, select Computer, click Browse to display the Select Computer dialog box, and then specify a computer account.

Simulate policy settings for the following:

User information

☒ Container:

☐ User:

Computer information

☒ Container:

☐ Computer:

4. On the Advanced Simulation Options page, select any advanced options for Slow Network Connections, Loopback Processing, and Site as necessary, and then click Next.

Simulate policy implementation for the following:

☐ Slow network connection (for example, a dial-up connection)


☐ Loopback processing

☐ Replace

☐ Merge

Site:

5. On the User Security Groups page, you can simulate changes to the security group membership of the applicable user or users. Any changes you make to group membership affect the previously selected user or user container. For example, to simulate what happens if a user in the designated user container is a member of the CorpManagers group, add this group to the Security Groups list. Click Next.



6. On the Computer Security Groups page, you can simulate changes to the applicable security group membership for a computer or computers. Any changes you make to group membership affect the previously selected computer or computer container. For example, to simulate what happens if a computer in the designated computer container is a member of the RemoteComputers group, add this group to the Security Groups list. Click Next.



7. You can link Windows Management Instrumentation (WMI) filters to Group Policy Objects. By default, the selected users and computers are assumed to meet all the WMI filter requirements, which is what you want in most cases for planning purposes. Click Next twice to accept the default options.
8. Review the selections you made, and then click Next.

To make changes to your selections, click Back. To process the simulation, click Next.

| Selection                 | Settings                               |
|---------------------------|----------------------------------------|
| User container            | OU=Finance,DC=imaginedlands,DC=local   |
| Computer container        | CN=Computers,DC=imaginedlands,DC=local |
| Slow network simulation   | No                                     |
| Loopback mode             | (None)                                 |
| Site name                 | (None)                                 |
| User security groups      | (Not specified)                        |
| Computer security groups  | (Not specified)                        |
| WMI filters for users     | (All linked WMI filters equal TRUE)    |
| WMI filters for computers | (All linked WMI filters equal TRUE)    |

Processing the simulation on this domain controller:

CorpServer28.imaginedlands.local

9. After the wizard gathers policy information, click Finish. When the wizard finishes generating the report, the report is selected in the left pane and the results are displayed in the right pane.
10. When you select the Details tab in the right pane, you can determine the settings that would be applied by browsing the report. Computer policy information is listed under Computer Details. User policy information is listed under User Details.





## Determining Current Settings and Refresh Status

You can use Group Policy modeling for logging Resultant Set of Policy (RSoP). When you use Group Policy modeling in this way, you can review all the policy objects that apply to a computer and the last time the applicable policy objects were processed (refreshed).

All domain and enterprise administrators have permission to model Group Policy for logging RSoP, as do those who have been delegated the permission Read Group Policy Results Data. In the GPMC, you can model Group Policy for the purpose of logging

RSoP by right-clicking the Group Policy Results node, and then clicking Group Policy Results Wizard. When the Group Policy Results Wizard starts, follow the prompts.





# Maintaining Group Policy Objects

The GPMC features built-in options to help you maintain GPOs. You can use the options to copy, paste, and import policy objects as well as to back up and restore policy objects. You also can use the options to modify how GPOs are processed.



# Copying, Pasting, and Importing Policy Objects


Using the GPMC, copy and paste is fairly straightforward. The Copy and Paste options are available when you right-click a GPO. You can copy a policy object and all its settings in one container and then browse to the container into which you want to paste the copy by right-clicking the container and selecting Paste.

When you are copying and pasting from one domain to another, the source and target domains can be any domains you can connect to in the GPMC and for which you have permission to manage related policy objects. In the source domain, you need Read permission to create a copy of a policy object. In the target domain, you need Write permission to write (paste) the copied policy object. Administrators have this privilege, as do those who have been delegated permission to create policy objects.

Copying policy objects between domains works well when you have connectivity between domains and the appropriate permissions. If you are an administrator at a remote office or have been delegated permissions, however, you might not have access to the source domain to create a copy of a policy object. In this case, another administrator can make a backup copy of a policy object for you and then send you the related data. When you receive the related data, you can import the backup copy of the policy object into your domain to create a policy object with the same settings.


Anyone with the Edit Settings Group Policy management privilege can perform an import operation. The import operation overwrites all the settings of the policy object you select. To import a backup copy of a policy object into a domain, follow these steps:

1. In the GPMC, right-click Group Policy Objects, and then click New. In the New GPO dialog box, enter a descriptive name for the new GPO, and then click OK.




2. The new GPO is now listed in the Group Policy Objects container. Right-click the new policy object, and then click Import Settings. This starts the Import Settings Wizard.

3. Click Next twice to bypass the Backup GPO page. You don't need to create a backup of the GPO at this time because it's new.
4. On the Backup Location page, click Browse. In the Browse For Folder dialog box, select the folder containing the backup copy of the policy object you want to import, and then click OK. Click Next to continue.




5. If multiple backups are stored in the designated backup folder, a list of them will be displayed on the Source GPO page. Click the one you want to use, and then click Next.



| Name                              | Time Stamp          | Description |
|-----------------------------------|---------------------|-------------|
| Default Application Security      | 9/8/2016 9:46:19 AM |             |
| Default Domain Controllers Policy | 9/8/2016 9:46:24 AM |             |
| Default Domain Policy             | 9/8/2016 9:46:25 AM |             |
| Sales                             | 9/8/2016 9:46:25 AM |             |
| Secure Workstation GPO            | 9/8/2016 9:46:25 AM |             |
| Technology                        | 9/8/2016 9:46:26 AM |             |


6. The Import Settings Wizard scans the policy object for references to security principals and Universal Naming Convention (UNC) paths that might need to be migrated. If any are found, you are given the opportunity to create migration tables or use existing migration tables.



7. Continue through the wizard by clicking Next, and then click Finish to begin the import process. When importing is complete, click OK. The import should




succeed. If there are problems, such as with your migration tables, you'll need to repeat this process and correct any issues.



# Backing Up and Restoring Policy Objects

As part of your periodic administration tasks, you should back up GPOs to protect them. You can use the GPMC to back up individual policy objects in a domain or all policy objects in a domain by following these steps:

1. In the GPMC, expand and then select the Group Policy Objects node. If you want to back up all policy objects in the domain, right-click the Group Policy Objects node, and then click Back Up All. If you want to back up a specific policy object in the domain, right-click the policy object, and then select Back Up.




The image shows a dialog box titled "Back Up Group Policy Object". It has a "Location:" label above a text box containing "C:\Logs\GPO-Backups". To the right of this text box is a "Browse..." button. Below the location section is a "Description:" label above an empty text box. At the bottom of the dialog are two buttons: "Back Up" and "Cancel".

2. In the Back Up Group Policy Object dialog box, click Browse. In the Browse For Folder dialog box, set the location where the GPO backup should be stored.
3. In the Description box, enter a description of the contents of the backup. Click Back Up to start the backup process.
4. The Backup dialog box shows the progress and status of the backup. Click OK when the backup is complete. If a backup fails, check the permissions on the policy and the folder to which you are writing the backup. You need Read permission on a policy and Write permission on the backup folder to create a backup. By default, members of the Domain Admins and Enterprise Admins groups should have these permissions.



Using the GPMC, you can restore a policy object to the state it was in when it was backed up. The GPMC tracks the backup of each policy object separately, even if you back up all policy objects at one time. Because version information is also tracked according to the backup time stamp and description, you can restore the last version of each policy object or a particular version of any policy object.



You can restore a policy object by following these steps:

1. In the GPMC, right-click the Group Policy Objects node, and then click Manage Backups. This displays the Manage Backups dialog box.
2. In the Backup Location box, click Browse. In the Browse For Folder dialog box, find the backup folder, and then click OK.

- 
- 
3. All policy object backups in the designated folder are listed under Backed Up GPOs. To show only the latest version of the policy objects according to the time stamp, select Show Only The Latest Version Of Each GPO.
  4. Select the GPO you want to restore. If you want to confirm its settings, click View Settings, and then use Internet Explorer to verify that the settings are as expected. When you are ready to continue, click Restore. Confirm that you want to restore the selected policy object by clicking OK.
  5. The Restore dialog box shows the progress and status of the restore operation. If a restore operation fails, check the permissions on the policy object and the folder from which you are reading the backup. To restore a GPO, you need Edit Settings, Delete, and Modify Security permissions on the policy object and Read permission on the folder containing the backup. By default, members of the Domain Admins and Enterprise Admins groups should have these permissions.






## Disabling Part of a GPO

Another way to disable a policy is to disable an unused part of the GPO. When you do this, you block computer configuration or user configuration settings, or both, and don't allow them to be applied. When you disable part of a policy that isn't used, the application of GPOs will be faster.

You can enable and disable policies partially or entirely by following these steps:

1. In the GPMC, select the container for the site, domain, or organizational unit with which you want to work.
  2. Select the policy object you want to work with, and then click the Details tab in the right pane.
  3. Choose one of the following status settings from the GPO Status list, and then click OK when prompted to confirm that you want to change the status of this GPO:
- **All Settings Disabled** Disallows processing of the policy object and all its settings.
  - **Computer Configuration Settings Disabled** Disables processing of computer configuration settings. This means that only user configuration settings are processed.
  - **Enabled** Allows processing of the policy object and all its settings.
  - **User Configuration Settings Disabled** Disables processing of user configuration settings. This means that only computer configuration settings are processed.



## Changing Processing Preferences

In Group Policy, computer configuration settings are processed when a computer starts and accesses the network. User configuration settings are processed when a user logs on to the network. In the event of a conflict between settings in Computer Configuration and User Configuration, the computer configuration settings win. It is also important to remember that computer settings are applied from the computer's GPOs, and user settings are applied from the user's GPOs.

In some special situations, you might not want this behavior. On a shared computer, you might want the user settings to be applied from the computer's GPOs, but you might also want to allow user settings from the user's GPOs to be applied. In a secure

lab or kiosk environment, you might want the user settings to be applied from the computer's GPOs to ensure compliance with strict security rules or guidelines for the lab. By using loopback processing, you can allow for these types of exceptions and obtain user settings from a computer's GPOs.

To change the way loopback processing works, follow these steps:

1. In the GPMC, right-click the Group Policy object you want to modify, and then click Edit.
  2. In the Administrative Templates for Computer Configuration under System\Group Policy, double-click the Configure User Group Policy Loopback Processing Mode policy. This displays a Properties dialog box for the policy.
  3. Define the policy by selecting Enabled, selecting one of the following processing modes from the Mode list, and then clicking OK:
- **Replace** Select the Replace option to ensure that user settings from the computer's GPOs are processed and that user settings in the user's GPOs are not processed. This means that the user settings from the computer's GPOs replace the user settings normally applied to the user.
  - **Merge** Select the Merge option to ensure that the user settings in the computer's GPOs are processed first, then user settings in the user's GPOs, and then user settings in the computer's GPOs again. This processing technique serves to combine the user settings in both the computer's and the user's GPOs. In the event of a conflict, the user settings in the computer's GPOs take precedence and overwrite the user settings in the user's GPOs.

Not Configured ☐ Comment:

Enabled ☒

Disabled ☐

Supported on: At least Windows 2000

Options: Mode: Replace

Help: This policy setting directs the system to apply the set of Group Policy objects for the computer to any user who logs on to a computer affected by this setting. It is intended for special-use computers, such as those in public places, laboratories, and classrooms, where you must modify the user setting based on the computer that is being used.

# Configuring Slow-Link Detection

Group Policy clients use slow-link detection to identify increased latency and reduced responsiveness on the network and to take corrective action to reduce the likelihood that processing of Group Policy will further saturate the network. After a slow link is detected, Group Policy clients reduce their network communications and requests, thereby reducing the overall network traffic load by limiting the amount of policy processing they do.

By default, if the connection speed is determined to be less than 500 kilobits per second (which could also be interpreted as high latency/reduced responsiveness on a fast network), the client computer interprets this as a slow network connection and notifies the domain controller. As a result, only security settings and administrative templates in the applicable policy objects are sent by the domain controller during a policy refresh.

You can configure slow-link detection by

Using the Configure Group Policy Slow Link Detection policy, you can control the way slow-link detection works. This policy is stored in the Administrative Templates for Computer Configuration under System\Group Policy.

Configure Group Policy slow link detection

Previous Setting Next Setting

☐ Not Configured Comment:

☒ Enabled

☐ Disabled

Supported on: At least Windows 2000

Options:

Connection speed (Kbps): 500

Enter 0 to disable slow link detection.

☐ Always treat WWAN connections as a slow link

Help:

This policy setting defines a slow connection for purposes of applying and updating Group Policy.

If the rate at which data is transferred from the domain controller providing a policy update to the computers in this group is slower than the rate specified by this setting, the system considers the connection to be slow.

If you disable this policy or do not configure it, clients use the default value of 500 kilobits per second to determine whether they are on a slow link. If you enable this policy, you can set a specific slow-link value, such as 384 kilobits per second. You also can specify that 3G connections should always be treated as slow links. Alternatively, if you want to disable slow-link detection completely, set the Connection Speed option to 0. This setting effectively tells clients not to detect slow links and to consider all links to be fast.

**REAL WORLD** Microsoft refers to connections on cellular and broadband as costed networks . Several policies are designed to help specify how networking should be used with mobile devices on costed networks. You can:

Control offline file synchronization on costed networks by using the Enable File Synchronization On Costed Networks policy found under Computer Configuration\Administrative Templates\Network\Offline Files.

Control background transfers on costed networks by using the Set Default Download Behavior For BITS Jobs On Costed Networks policy found under Computer Configuration\Administrative Templates\Network\Background Intelligent Transfer Services (BITS).

Specify that costed broadband networks have fixed, variable, or unrestricted usage charges by using the Set Cost policy found under Computer Configuration\Administrative Templates\Network\WLAN Service\WLAN Media Cost.

Specify that costed cellular networks have fixed, variable, or unrestricted usage charges by using the Set 3G Cost and Set 4G Cost policies found under Computer Configuration\Administrative Templates\Network\WWAN Service\WWAN Media Cost.

You can optimize slow-link detection for various areas of Group Policy processing as necessary. By default, policy areas that are not processed when a slow link is detected include:

- [Disk Quota Policy Processing](#)
- [EFS Recovery Policy Processing](#)
- [Folder Redirection Policy Processing](#)
- [Scripts Policy Processing](#)
- [Software Installation Policy Processing](#)

Security Policy Processing is always enabled automatically for slow links. By default, security policy is refreshed every 16 hours even if security policy has not changed. The only way to stop the forced refresh is to configure security policy processing so that it is not applied during periodic background refreshes. To do this, select the policy setting Do Not Apply During Periodic Background Processing. However, because security policy is so important, the Do Not Apply setting means only that security policy processing is stopped when a user is logged on and using the computer. One of the only reasons you'll want to stop security policy refreshes is applications failing during refresh operations.



You can configure slow-link detection and related policy processing by following these steps:

1. In the GPMC, right-click the policy object you want to modify, and then click Edit.
2. In the Administrative Templates for Computer Configuration under System\Group Policy, double-click the Configure Group Policy Slow Link Detection policy.
3. Select Enabled to define the policy. In the Connection Speed box, specify the speed that should be used to determine whether a computer is on a slow link. You also can specify that WWAN connections should always be treated as slow links. Click OK.

To configure slow-link and background policy processing of key areas of Group Policy, follow these steps:

1. In the GPMC, right-click the policy object you want to modify, and then click Edit.
2. Expand Computer Configuration\Administrative Templates\System\Group Policy.
3. Double-click the processing policy you want to configure. Click Enabled to define the policy, and then make your configuration selections. The options differ slightly depending on the policy selected and might include the following:
  - **Allow Processing Across A Slow Network Connection** Ensures that the related policy settings are processed even on a slow network.
  - **Do Not Apply During Periodic Background Processing** Overrides refresh settings when related policies change after startup or logon.
  - **Process Even If The Group Policy Objects Have Not Changed** Forces the client computer to process the related policy settings during a refresh even if the settings haven't changed.
4. Click OK to save your settings.



☐ Not Configured    Comment:

☒ Enabled

☐ Disabled

Supported on:

Options:

- ☒ Allow processing across a slow network connection
- ☒ Do not apply during periodic background processing
- ☒ Process even if the Group Policy objects have not changed

Help:

This policy setting determines when disk quota policies are updated.

This policy setting affects all policies that use the disk quota component of Group Policy, such as those in Computer Configuration\Administrative Templates\System\Disk Quotas.

This policy setting overrides customized settings that the program implementing the disk quota policy set when it was installed.

## Removing Links and Deleting GPOs

In the GPMC, you can stop using a linked GPO in two ways:

- [Remove a link to a GPO but not the GPO itself.](#)
- [Permanently delete the GPO and all links to it.](#)

Removing a link to a GPO stops a site, domain, or organizational unit from using the related policy settings but does not delete the GPO. Because of this, the GPO remains linked to other sites, domains, or organizational units as appropriate. In the GPMC, you can remove a link to a GPO by right-clicking the GPO link in the container that it is linked to, and then clicking Delete. When prompted to confirm that you want to remove the link, click OK. If you remove all links to the GPO from sites, domains, and organizational units, the GPO continues to exist in the Group Policy Objects container, but its policy settings have no effect in your organization.

Permanently deleting a GPO removes the GPO and all links to it. The GPO will not exist in the Group Policy Objects container and will not be linked to any sites, domains, or organizational units. The only way to recover a deleted GPO is to restore it from a backup (if one is available). In the GPMC, you can remove a GPO and all links to the object from the Group Policy Objects node. Right-click the GPO, and then select Delete. When prompted to confirm that you want to remove the GPO and all links to it, click Yes.

## Troubleshooting Group Policy

Group Policy should be carefully maintained to ensure proper operation. As problems occur, you should try to diagnose and resolve them. To troubleshoot Group Policy, you need a strong understanding of how policy is refreshed and processed. You also need a strong understanding of general maintenance and troubleshooting tasks.



# Diagnosing Group Policy Issues

When you are trying to determine why policy is not being applied as expected, one of first things you should do is examine the Resultant Set of Policy for the user and computer experiencing problems with policy settings. You can determine the GPO that a setting is applied from by following these steps:

1. In the GPMC, right-click the Group Policy Results node, and then click Group Policy Results Wizard. When the wizard starts, click Next.
2. On the Computer Selection page, select This Computer to view information for the local computer. To view information for a remote computer, select Another Computer, and then click Browse. In the Select Computer dialog box, enter the name of the computer, and then click Check Names. After you select the correct computer account, click OK, and then click Next.
3. On the User Selection page, select the user whose policy information you want to view. You can view policy information for any user who has logged on to the previously selected computer. Click Next.
4. Review the selections you made, and then click Next. After the wizard gathers policy information, click Finish. When the wizard finishes generating the report, the report is selected in the left pane, and the results are displayed in the right pane.
5. To determine the settings that are being applied, browse through the report. Computer and user policy information is listed separately. Computer policy information is listed under Computer Configuration Summary. User policy information is listed under User Configuration Summary.

Using the Gpresult command-line utility, you can also view RSoP. Gpresult provides details about the following:

- Special settings applied for folder redirection, software installation, disk quota, IPsec, and scripts
- The last time Group Policy was applied
- The domain controller from which policy was applied, and the security group memberships for the computer and user
- The complete list of GPOs that were applied, and the complete list of GPOs that were not applied because of filters

Gpresult has the following basic syntax.

```
gpresult /s ComputerName /user Domain \ UserName
```





Here ComputerName is the name of the computer you want to log policy results for, and Domain \ UserName indicates the user you want to log policy results for. For example, to view the RSoP for CorpPC85 and the user Tedg in the Imaginedl domain, you would enter the following command.

```
gpresult /s corppc85 /user imaginedl\tedg
```

You can view more detailed output by using one of the two verbose options. The /v parameter turns on verbose output, and results are displayed only for policy settings in effect. The /z parameter turns on verbose output with settings for policy settings in effect and all other GPOs that have the policy set. Because Gpresult output can be fairly long, you should create an HTML report by using the /h parameter, or an XML report by using the /x parameter. The following examples use these parameters.

```
gpresult /s corppc85 /user imaginedl\tedg /h gpreport.html
```

```
gpresult /s corppc85 /user imaginedl\tedg /x gpreport.xml
```



## Fixing Default Group Policy Objects

The Default Domain Policy and Default Domain Controller Policy GPOs are vital to the health of Active Directory Domain Services. If for some reason these policies become corrupted, Group Policy will not function properly. To resolve this, you must use the GPMC to restore a backup of these GPOs. If you are in a disaster-recovery scenario and do not have any backups of the Default Domain Policy or the Default Domain Controller Policy, you can use Dcgpofix to restore the security settings in these policies. The state that Dcgpofix restores these objects to depends on how you modified security and on the security state of the domain controller before you ran Dcgpofix. You must be a member of Domain Admins or Enterprise Admins to run Dcgpofix.

When you run Dcgpofix, both the Default Domain Policy and Default Domain Controller

Policy GPOs are restored by default, and you lose any base changes made to these GPOs. Some policy settings are maintained separately and are not lost, including Windows Deployment Services (WDS), Security Settings, and Encrypting File System (EFS). Nondefault Security Settings are not maintained, however, which means that other policy changes could also be lost. All other policy settings are restored to their previous values, and any changes you've made are lost.

To run Dcgpofix, log on to a domain controller in the domain in which you want to fix default Group Policy, and then enter **dcgpofix** at an elevated prompt. –Dcgpofix checks the Active Directory schema version number to ensure compatibility between the version of Dcgpofix you are using and the Active Directory schema configuration. If the versions are not compatible, Dcgpofix exits without fixing the default Group Policy Objects. By specifying the /ignoreschema parameter, you can enable Dcgpofix to work with different versions of Active Directory. However, default policy objects might not be restored to their original state. Because of this, you should always be sure to use the version of Dcgpofix that is installed with the current operating system.

You also have the option of fixing only the Default Domain Policy or only the Default Domain Controller Policy GPO. If you want to fix only the Default Domain Policy, enter **dcgpofix /target:domain** . If you want to fix only the Default Domain Controller Policy, enter **dcgpofix /target:dc** .

## Chapter 16. Optimizing Server Security

Windows Server 2016 is designed to allow users to access shared resources, files, applications, and services without letting anyone get into areas of the network they shouldn't have access to, or peek into things they shouldn't see.



The secret? Windows Server provides each person with his or her own user account, which has its own permissions, privileges and capabilities that work together to ensure a person can only access things they are permitted to access. As an IT professional, it's your job to ensure that sound security practices are followed and that the security settings in place isolate users from things they shouldn't get into.

This chapter explains how to lock down your servers. It also explains how to optimize security settings and disable or restrict features of the operating system that represent potential security risks.





# Locking Down Your Servers

One way to configure global security settings for computers running Windows-based operating systems is to use security templates, which provide a centralized way to manage security. You use security templates to apply customized sets of Group Policy definitions. These policy definitions generally affect:

- **Account policies** Control security for passwords, account lockout, and Kerberos security
- **Local policies** Control security for auditing, user rights assignment, and other security options
- **Event log policies** Control security for event logging
- **Restricted groups policies** Control security for local group membership administration
- **System services policies** Control security and startup mode for local services
- **File system policies** Control security for file and folder paths in the local file system
- **Registry policies** Control the permissions on security-related registry keys

Security templates are available in all Microsoft Windows Server installations and can be imported into any Group Policy object (GPO). Security templates apply only to the Computer Configuration area of Group Policy. They do not apply to the User Configuration area.

**NOTE** In Group Policy, you'll find settings you can manage with templates under Windows Settings\Security Settings in the Computer Configuration policies. Some security settings cannot be managed with templates. These include those that apply to wireless networks, public keys, software restrictions, and IP security.

## Managing Security Configurations

Working with security templates is a multipart process that involves the following steps:


1. Use the Security Templates snap-in to create a new template, or select an existing template that you want to modify.
2. Use the Security Templates snap-in to make necessary changes to the template settings, and then save the changes.
3. Use the Security Configuration And Analysis snap-in to analyze the differences between the template with which you are working and the current computer security settings.



4. Revise the template as necessary after you review the differences between the template settings and the current computer settings.
5. Use the Security Configuration And Analysis snap-in to apply the template and overwrite existing security settings.

When you first start working with security templates, you should determine whether you can use an existing template as a starting point. Other administrators might have created templates, or your organization might have baseline templates that should be used. You can also create a new template to use as your starting point.

If you select a template that you want to use as a starting point, you should go through each setting that the template applies and evaluate how the setting affects your environment. If a setting doesn't apply, you should modify it appropriately or delete it.




You use the Security Configuration And Analysis snap-in to apply templates rather than the Security Templates snap-in. You can also use the Security Configuration And Analysis snap-in to compare the settings in a template to the current settings on a computer. The results of the analysis highlight areas in which the current settings don't match those in the template. This is useful to determine whether security settings have changed over time.

## Accessing and Preparing the Security Tools

You can open the security snap-ins by following these steps:

1. Start the Microsoft Management Console (MMC) by pressing the Windows key, typing **mmc.exe**, and then pressing Enter.
2. In the Microsoft Management Console, click File, and then click Add/Remove Snap-In.
3. In the Add Or Remove Snap-Ins dialog box, click Security Templates, and then click Add.
4. Click Security Configuration And Analysis, and then click Add. Click OK.




By default, the Security Templates snap-in looks for security templates in the %SystemDrive%\Users\%UserName%\Documents\Security\Templates folder. You can add other search paths for templates by following these steps:

1. In the MMC, right click the Security Templates node and then choose New Template Search Path.
2. In the Browse For Folder dialog box, select the template location to add, such as %SystemRoot%\Security\Templates. Click OK.

Once you locate the template search path with which you want to work, you can select a template and expand the related nodes to review its settings. To create a template, follow these steps:

1. In the Security Templates snap-in, right-click the search path where the template should be created, and then click New Template.
2. Enter a name and description for the template in the text boxes provided. Click OK to create the template.





# Optimizing Security Settings

The template will not have settings configured, so you need to modify the settings carefully before the template is ready for use. You manage each type of template setting in a slightly different way.

Save any changes you make by right-clicking the template in the MMC and selecting Save. Alternatively, you can select Save As to assign a different name to the modified template.




# Changing Settings for Core Policies

Account policy settings control security for passwords, account lockout, and Kerberos security. Local policy settings control security for auditing, user rights assignment, and other security options. Event log policy settings control security for event logging. For detailed information on account policy, local policy settings, and configuring event logging, see Chapter 8, “Managing Account Policies.”

With account, local, and event log policies, you can change template settings by following these steps:

1. In the Security Templates snap-in, expand the Account Policies or Local Policies node as necessary, and then select a related subnode, such as Password Policy or Account Lockout Policy.
2. In the right pane, policy settings are listed alphabetically. The value in the Computer Setting column shows the current setting. If the template changes the setting so that it is no longer defined, the value is listed as Not Defined.
3. Double-click a setting to display its Properties dialog box. To determine the purpose of the setting, click the Explain tab. To define and apply the policy setting, select the Define This Policy Setting In The Template check box. To clear this policy and not apply it, clear this check box.



4. If you enable the policy setting, specify how the policy setting is to be used by configuring any additional options.
5. Click OK to save your changes.


**NOTE** After changing account or local policy settings, you might see the Suggested Value Changes dialog box. This dialog box informs you of other values that are changed to suggested values based on your setting change. For example, when you change the Account Lockout Threshold setting, Windows might also change the Account Lockout Duration and Reset Account Lockout Counter After settings.




# Configuring Restricted Groups

Restricted groups policy settings control the list of group members and the groups to which the configured group belongs. You can restrict a group by following these steps:

1. In the Security Templates snap-in, select the Restricted Groups node.
2. In the right pane, any currently restricted groups are listed by name. Members of the group are listed as well, and so are groups of which the restricted group is a member.
3. You can add a restricted group by right-clicking the Restricted Groups node in the left pane, and then clicking Add Group.




4. In the Add Group dialog box, click Browse. In the Select Groups dialog box, enter the name of a group you want to restrict, and then click Check Names. If multiple matches are found, select the account you want to use, and then click OK. If no matches are found, update the name you entered and try searching again.



5. When you click OK to close the Add Group dialog box, a properties dialog box for the previously selected group is displayed. Use the Add Members option to add members to the group. Click Add Members, and then specify the members of the group. If the group should not have any members, remove all members by clicking Remove. Any members who are not specified in the policy setting for the restricted group are removed when the security template is applied.

6. Click Add Groups to specify the groups to which this group belongs. If you specify membership in groups, the groups to which this group belongs are listed exactly as you've applied them (if the groups are valid in the applicable workgroup or domain). If you do not specify membership in groups, the groups to which this group belongs are not modified when the template is applied.
7. Click OK to save your settings.



You can remove a restriction on a group by following these steps:

1. In the Security Templates snap-in, select the Restricted Groups node.
2. In the right pane, any currently restricted groups are listed by name. Members of the group are listed along with the groups of which the restricted group is a member.
3. Right-click the group that should not be restricted, and then click Delete. When prompted to confirm the action, click Yes.

## Restricting System Services

Policy settings for system services control the general security and startup mode for local services. You can enable, disable, and configure system services by following these steps:

1. In the Security Templates snap-in, select the System Services node.



2. In the right pane, all currently installed services on the computer with which you are working are listed by name, startup setting, and permission configuration. Keep the following in mind when working with system services:
  - If the template does not change the startup configuration of the service, the value for the Startup column is listed as Not Defined. Otherwise, the startup configuration is listed as one of the following values: Automatic, Manual, or Disabled.
  - If the template does not change the security configuration of the service, the value for the Permission column is listed as Not Defined. Otherwise, the security configuration is listed as Configured.
3. Double-click the entry for a system service to display its Properties dialog box. To define and apply the policy setting, select the Define This Policy Setting In The Template check box. To clear this policy and not apply it, clear this check box.



4. If you enable the policy setting, specify the service startup mode by selecting Automatic, Manual, or Disabled. Keep the following in mind:
  - Automatic ensures that the service starts automatically when the operating system starts. Choose this setting for essential services that you know are secure and that you want to be sure are run if they are installed on the computer to which the template is being applied.
  - Manual prevents the service from starting automatically and allows the service only to be started manually, either by a user, application, or other service. Choose this setting when you want to restrict unnecessary or unused services or when you want to restrict services that you know are not entirely secure.






- Disabled prevents the service from starting automatically or manually. Choose this setting only with unnecessary or unused services that you want to prevent from running.
5. If you know the security configuration that the service should use, click Edit Security, and then set the service permissions in the Security For dialog box. You can set permissions to allow specific users and groups to start, stop, and pause the service on the computer.
  6. Click OK.



# Restricting Registry and File System Paths

Policy settings for the file system control security for file and folder paths in the local file system. Policy settings for the registry control the values of security-related registry keys. You can view or change security settings for currently defined registry and file system paths by following these steps:

1. In the Security Templates snap-in, select the Registry node or the File System node, depending on which type of file path with which you want work. In the right pane, all currently secured paths are listed.
2. Double-click a registry or file path to view its current settings.




3. To ensure that permissions on the path or key are not replaced, select Do Not Allow Permissions On This Key To Be Replaced, and then click OK. Skip the remaining steps in this procedure.
  4. To configure the path or key and replace permissions, select Configure This Key Then, and then choose one of the following options:
- **Propagate Inheritable Permissions To All Subkeys** Choose this option to apply all inheritable permissions to this registry or file path and to all registry and file paths below this path. Existing permissions are replaced only if they conflict with a security permission set for this path.
  - **Replace Existing Permissions On All Subkeys With Inheritable Permissions** Choose this option to replace all existing permissions on this registry or file path and on all registry and file paths below this path. Any existing permissions are removed, and only the current permissions remain.

5. Click Edit Security. In the Security For dialog box, configure security permissions for users and groups. You have the same options for permissions, auditing, and ownership as you do for files and folders used with NTFS.
6. Click OK twice to save the settings.

You can define security settings for registry paths by following these steps:

1. In the Security Templates snap-in, select and then right-click the Registry node, and then click Add Key to display the Select Registry Key dialog box.
2. In the Select Registry Key dialog box, select the registry path or value with which you want to work, and then click OK. Entries under CLASSES\_ROOT are for HKEY\_CLASSES\_ROOT. Entries under MACHINE are for HKEY\_LOCAL\_MACHINE. Entries under USERS are for HKEY\_USERS.




3. In the Database Security For dialog box, configure security permissions for users and groups. You have the same options for permissions, auditing, and ownership as you do for files and folders used with NTFS.
4. Click OK. The Add Object dialog box is displayed. To ensure that permissions on the path or key are not replaced, select Do Not Allow Permissions On This Key To Be Replaced, and then click OK. Skip the remaining steps in this procedure.
5. To configure the path or key and replace permissions, select Configure This Key Then, and then do one of the following:
  - Choose Propagate Inheritable Permissions To All Subkeys to apply all inheritable permissions to this registry path and all registry paths below this path. Existing

permissions are replaced only if they conflict with a security permission set for this path.

- Choose Replace Existing Permissions On All Subkeys With Inheritable Permissions to replace all existing permissions on this registry path and on all registry paths below this path. Any existing permissions are removed, and only the current permissions remain.

6. Click OK.



You can define security settings for file paths by following these steps:

1. In the Security Templates snap-in, select and then right-click the File System node, and then click Add File to display the Add A File Or Folder dialog box.
2. In the Add A File Or Folder dialog box, select the file or folder path or value with which you want to work, and then click OK.
3. In the Database Security For dialog box, configure security permissions for users and groups. You have the same options for permissions, auditing, and ownership as you do for files and folders used with NTFS.
4. Click OK. The Add Object dialog box is displayed. To ensure that permissions on the path are not replaced, select Do Not Allow Permissions On This File Or Folder To Be Replaced, and then click OK. Skip the remaining steps in this procedure.
5. To configure the path and replace permissions, select Configure This Path Then, and then do one of the following:



- Choose Propagate Inheritable Permissions To All Subfolders to apply all inheritable permissions to this file path and all file paths below this path. Existing permissions are replaced only if they conflict with a security permission set for this path.
- Choose Replace Existing Permissions On All Subfolders With Inheritable Permissions to replace all existing permissions on this file path and on all file paths below this path. Any existing permissions are removed, and only the current permissions remain.

6. Click OK.





## Reviewing, and Applying Security Settings


As discussed previously, the Security Configuration And Analysis snap-in is used to apply templates and to compare the settings in a template to the current settings on a computer. Applying a template ensures that a computer conforms to a specific security configuration. Comparing settings can help you identify any discrepancies between what is implemented currently and what is defined in a security template. This can also be useful to determine whether security settings have changed over time.

When you are working with the Security Configuration And Analysis snap-in, you can configure security only on the computer on which you are running the snap-in. If you want to use this tool to deploy security configurations on other computers, you must log on to and run the tool on each computer. Although this technique works for standalone computers, it is not the optimal approach in a domain. In a domain setting, you'll want to import the security template settings into a Group Policy object (GPO). You then use the GPO to deploy the security configuration to multiple computers. For more information, see "Deploying Security Configurations" later in this chapter.

The Security Configuration And Analysis snap-in uses a working database to store template security settings, and then applies the settings from this database. For analysis and comparisons, the template settings are listed as the effective database settings and the current computer settings are listed as the effective computer settings. Keep in mind that if you are actively editing a template in the Security Templates snap-in, you need to save the template so that the changes can be analyzed and used.

After you create a template or determine that you want to use an existing template, you can analyze and then configure the template by following these steps:

1. In the MMC, right-click the Security Configuration And Analysis node, and then click Open Database to display the Open Database dialog box.
2. A security database provides the work area for analyzing security settings. You can either open an existing database by specifying its name or create a new one by specifying a descriptive name for the database, such as **Current Config Comparison**. Afterward, click Open. By default, the Open Database dialog box's search path is set to %SystemDrive%\Users\%UserName%\Documents\Security\Database. The databases uses the Security Database Files format and has the .sdb file extension.





3. After the snap-in then either creates or opens a security database, the Import Template dialog box is displayed. Use the Import Template dialog box to specify the template that you want to work with and then click Open. The default search path set to  
%SystemDrive%\Users\%UserName%\Documents\Security\Templates.  
Security template files end with the .inf file extension.



4. Right-click the Security Configuration And Analysis node, and then click Analyze Computer Now. When prompted to set the error log path, enter a new path or click OK to use the default path.
5. Wait for the snap-in to complete the analysis of the template. If an error occurs during the analysis, you can view the error log by right-clicking the Security Configuration And Analysis node and choosing View Log File.
6. Review the differences between the template settings and the current computer settings. The template settings stored in the analysis database are listed in the Database Setting column, and the current computer settings are listed in the Computer Setting column. If a setting has not been analyzed, it is listed as Not Defined.

| Policy                                      | Database Setting   | Computer Setting        |
|---------------------------------------------|--------------------|-------------------------|
| Enforce password history                    | 5 passwords rem... | 24 passwords remembered |
| Maximum password age                        | 45 days            | 42 days                 |
| Minimum password age                        | 30 days            | 1 days                  |
| Minimum password length                     | 8 characters       | 7 characters            |
| Password must meet complexity requirements  | Disabled           | Enabled                 |
| Store passwords using reversible encryption | Disabled           | Disabled                |

To make changes to a setting stored in the database, follow these steps:

1. In the Security Configuration And Analysis snap-in, double-click the setting with which you want to work.
2. In the Properties dialog box, note the current computer setting. If information about the purpose of the setting is available, you can view this information by clicking the Explain tab.
3. To define and apply the policy setting, select the Define This Policy In The Database check box. To clear this policy and not apply it, clear this check box.



4. If you enable the policy setting, specify how the policy setting is to be used by configuring any additional options.
5. Repeat this process as necessary. To save your database changes to the template, right-click the Security Configuration And Analysis node, and then click Save.

You can also use the Secedit command-line utility to analyze, review, and apply security templates. The basic technique is as follows:

1. Open an elevated administrator prompt.
2. Use Secedit /Import to import a security template into a working database.
3. Use Secedit /Analyze to compare the template settings to a computer's current settings.
4. Use Secedit /Configure to apply the template settings.

Whether you are working with the graphical wizard or the command-line utility, you might want to create a rollback template before applying any settings. A rollback template is a reverse template that allows you to remove most settings applied with a template. The only settings that cannot be removed are those for access control lists on file system and registry paths.

At an elevated administrator prompt, you can create a rollback template by using the Secedit command-line utility. Enter the following:

```
secdit /generaterollback /db DatabaseName /cfg TemplateName  
/rbk RollBackName /log LogName
```

where DatabaseName is the name of a new database that will be used to perform the rollback, TemplateName is the name of an existing security template for which you are creating a rollback template, RollBackName sets the name of a new security template in which the reverse settings should be stored, and LogName sets the name of an optional file for tracking the status of the rollback process.

In the following example, you create a rollback template for the "File Servers" template:

```
secdit /generaterollback /db rollback.db /cfg "file servers.inf"  
/rbk fs-orig.inf /log rollback.log
```

When you're ready to apply the template, right-click the Security Configuration And Analysis node, and then click Configure Computer Now. When prompted to set the error log path, click OK because the default path should be sufficient. To view the



configuration error log, right-click the Security Configuration And Analysis node, and then click View Log File. Note any problems, and take action as necessary.

If you created a rollback template prior to applying a security template, you can restore the computer's security settings to its previous state. To apply a rollback template, follow these steps:

1. In the Security Configuration And Analysis snap-in, right-click the Security Configuration And Analysis node, and then click Import Template.
2. In the Import Template dialog box, select the rollback template.
3. Select the Clear This Database Before Importing check box, and then click Open.
4. Right-click the Security Configuration And Analysis node, and then click Configure Computer Now. Click OK.

The only settings that cannot be restored are for access control lists on file system and registry paths. After the permissions on file system and registry paths have been applied, you cannot reverse the process automatically and must instead manually reverse the changes one at a time.





# Locking Down Multiple Computers

Rather than applying security templates to one computer at a time, you can deploy your security configurations to multiple computers through Group Policy. To do this, you need to import the security template into a GPO processed by the computers to which the template settings should apply. Then, when policy is refreshed, all computers within the scope of the GPO receive the security configuration.





## Preparing for Configuration Deployment

Security templates apply only to the Computer Configuration portion of Group Policy. Before you deploy security configurations in this way, you should take a close look at the domain and organizational unit (OU) structure of your organization and make changes as necessary to ensure that the security configuration is applied only to relevant types of computers. Essentially, this means that you need to create OUs for the different types of computers in your organization, and then move the computer accounts for these computers into the appropriate OUs. Afterward, you need to create and link a GPO for each of the computer OUs. For example, you could create the following computer OUs:

- **Domain Controllers** An OU for your organization's domain controllers. This OU is created automatically in a domain.
- **High-Security Member Servers** An OU for servers that require higher than normal security configurations.
- **Member Servers** An OU for servers that require standard server security configurations.
- **Laptop and Mobile Devices** An OU for laptops and mobile devices which are inherently less secure and may require enhanced security configurations.
- **High-Security User Workstations** An OU for workstations that require higher than normal security configurations.
- **User Workstations** An OU for workstations that require standard workstation security configurations.
- **Remote Access Computers** An OU for computers that remotely access the organization's network.
- **Restricted Computers** An OU for computers that require restrictive security configurations, such as computers that are used in labs or kiosks.


**REAL WORLD** Be extremely careful when you deploy security templates through GPOs. If you haven't done this before, practice in a test environment first, and be sure to also practice recovering computers to their original security settings. If you create a GPO and link the GPO to the appropriate level in the Active Directory structure, you can recover the computers to their original state by removing the link to the GPO. This is why it's important to create and link a new GPO to apply security rather than use an existing GPO.

## Deploying Security Configurations

You can deploy a security template to a computer GPO by following these steps:



1. After you configure a security template and have tested it to ensure that it is appropriate, open the GPO you previously created and linked to the appropriate level of your Active Directory structure. In the Group Policy Management editor, open Computer Configuration\Windows Settings\Security Settings.
2. Right-click Security Settings, and then click Import Policy.
3. In the Import Policy From dialog box, select the security template to import, and then click Open. Security templates end with the .inf file extension.



After importing settings, you should check the configuration state of the security settings to verify that the settings were imported as expected, and then close the policy editor. Repeat the deployment process for each security template and computer GPO you've configured. In the default configuration of Group Policy, it will take 90 to 120 minutes for the settings to be pushed out to computers in the organization. Use Gpupdate to speed up the process.





# Chapter 17. Deploying Windows Server 2016

Welcome to Windows Server 2016: Server Infrastructure, your guide to managing Microsoft's newest server operating system. Before getting into the specifics of deploying and managing Windows servers, take a few seconds to familiarize yourself with the deployment options available, including

- **Physical server** A server deployment on a physical machine. You use a physical machine when you need a standalone installation of the operating system.
- **Virtual server** A server deployment to a virtual machine. You use a virtual server when you have an existing host server and want to install the operating system as a guest operating system.
- **Containerized server** A server deployment that uses containers. You use a container when you have an existing host server and want to establish the base image for a container stack as part of an application or service.

Physical servers are a standard deployment approach; virtual servers are the traditional type of virtualization. Virtual servers create virtual environments that run their own operating systems and manage their own memory, processors and resources.

Containers, on the other hand, virtualize the operating system layer rather than simulate the underlying physical machine.

Regardless of which deployment approach you use, you have three options for setting up the operating system:

- **Full-Server installation** This approach, referred to in the GUI as Server With Desktop Experience, configures a server with all user interfaces and management tools. You'll most often use a full-server installation with a physical server deployment and then primarily when you need to use roles or features not supported by the other installation options.
- **Server Core installation** This approach configures a server with a limited user interface that is designed to be managed locally using a command line and remotely using either graphical or command-line tools. Server Core supports a limited set of roles and features, primarily those for Active Directory, DHCP, DNS, and Hyper-V as well as file, print and media services. You use Server Core when you need a reduced footprint installation.
- **Nano Server installation** This approach configures a server with a minimal footprint that is designed to be managed remotely. For local management, there is only the Recovery Console. Nano Server supports a very limited set of roles and







features, primarily those for file servers, DNS servers and web servers. You use Nano Server when you need a minimal-footprint installation.

**NOTE** This book is designed to be used with Windows Server 2016: Administration Essentials. While this book focuses on infrastructure services, including those for file services, print services, DHCP and DNS, the latter book focuses on management and configuration of servers, setting up Active Directory and Group Policy and more. Chapter 1 of that book is where you'll find a complete discussion of configuration options, Nano Server and containers.

Together, the deployment and setup options give you many possibilities. Full Server is optimized for manageability; Server Core and Nano Server are optimized for performance. However, only a Full Server deployment provides the full graphical shell and the full desktop experience. As part of the graphical shell, user interfaces, graphical management tools, ink and handwriting services, and the server media foundation are installed. As part of the desktop experience, desktop themes, media player, Disk Cleanup and other Windows desktop tools are installed.





# Preparing for Windows Server 2016

You can install Windows Server 2016 by performing a clean installation or an upgrade. With a clean installation, Setup replaces any existing operating system on the computer, and all user or application settings are lost. With an upgrade, Setup performs a clean installation of the operating system and then migrates user settings, documents, and applications from the earlier version of Windows.





## Working with Windows Server 2016

Windows Server 2016 is available in two primary editions: Standard and Datacenter. While both editions provide the same core functionality, support for containers and nano servers, only Datacenter Edition supports storage replicas, shielded virtual machines and the extended networking stack. Further, with Standard Edition, servers are limited to two virtualized, Hyper-V environments and you must re-license the server for every two additional environments. In contrast, the license for Datacenter Edition allows an unlimited number of virtualized environments.

Before you install Windows Server 2016, you should be sure that your server meets the minimum requirements of the edition you plan to use. Windows Server 2016 requires a 64-bit processor that supports NX and DEP, which all current processors do. For core and nano server deployments, servers should have at least 512 MB of RAM and at least 32 GB of disk space for installation of the base operating system. For full-server deployments, servers should have at least 2 GB of RAM and at least 32 GB of disk space for the operating system.

**NOTE** For optimal performance, you should have at least 10 percent of free space on a server's disks at all times. Additional disk space is required for paging and dump files and also for the features, roles, and role services you install as well.

When you install Windows Server 2016, Setup automatically makes recovery options available on your server as an advanced boot option. In addition to a command line for troubleshooting and options for changing the startup behavior, you can use System Image Recovery to perform a full recovery of the computer by using a system image created previously. If other troubleshooting techniques fail to restore the computer and you have a system image for recovery, you can use this feature to restore the computer from the backup image.

After you install a server, you can configure it using the integrated management tools. The one you'll use the most for handling core configuration tasks is Server Manager. Server Manager provides setup and configuration options for the local server in addition to options for managing roles, features, and related settings on any remotely manageable server in the enterprise.






Tasks you can use Server Manager to perform include the following:

- Adding servers for remote management
- Initiating remote connections to servers
- Configuring the local server
- Managing installed roles and features
- Managing volumes and shares on file servers
- Configuring Network Interface Card (NIC) Teaming
- Viewing events and alerts
- Restarting servers

Server Manager is great for general system administration, but you also need a tool that gives you granular control over system environment settings and properties. This is where the Settings app comes into the picture. The Settings app is replacing the System utility as the preferred tool for managing a computer's basic configuration.



When you are working remotely, you'll need to use command-line tools for these tasks instead. Windows Server 2016 supports two command-line environments, regardless of whether you are working with a Server Core, nano server or full-server deployment:

- [Windows Command Processor \(cmd.exe\)](#)
- [Windows PowerShell \(powershell.exe\)](#)

Windows Server 2016 uses the same configuration architecture as Windows Server 2012 and Windows Server 2012 R2. You prepare servers for deployment by installing and configuring the following components:

- **Server roles** A server role is a related set of software components that allows a server to perform a specific function for users and other computers on a network. A computer can be dedicated to a single role, such as Active Directory Domain Services (AD DS), or provide multiple roles.
- **Role services** A role service is a software component that provides the functionality for a server role. Each role can have one or more related role services. Some server roles, such as Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP), have a single function, and installing the role installs this function. Other roles, such as Network Policy and Access Services, and Active Directory Certificate Services (AD CS), have multiple role services that you can install. With these server roles, you can choose which role services to install.
- **Features** A feature is a software component that provides additional functionality. Features, such as BitLocker Drive Encryption and Windows Server Backup, are



installed and removed separately from roles and role services. A computer can have zero or more features installed, depending on its configuration.

You configure roles, role services, and features by using Server Manager. Some roles, role services, and features are dependent on other roles, role services, and features. As you install roles, role services, and features, Server Manager prompts you to install other roles, role services, or features that are required. Similarly, if you try to remove a required component of an installed role, role service, or feature, Server Manager warns that you cannot remove the component unless you also remove dependent roles, role services, or features.

Because adding or removing roles, role services, and features can change hardware requirements, you should carefully plan any configuration changes and determine how they affect a server's overall performance. Although you typically want to combine complementary roles, doing so increases the workload on the server, so you need to optimize the server hardware accordingly. Unlike early releases of Windows, Windows Server 2016 does not install some important server features automatically. For example, you must add Windows Server Backup to use the built-in backup and restore features of the operating system.



# Managing Server Core

Server Core is designed to be managed locally using the command line or remotely using either command-line or graphical tools. You use Server Core when you don't need a desktop and want to reduce the overall footprint of the server. A server running in a core configuration doesn't have the Windows shell, Internet Explorer or local GUI management.

With a Server Core installation, you get a user interface that includes a limited desktop environment for local console management of the server. This minimal interface includes the following:

- Windows Logon screen for logging on and logging off
- Notepad (Notepad.exe) for editing files
- Registry Editor (Regedit.exe) for managing the registry
- Task Manager (Taskmgr.exe) for managing tasks and starting new tasks
- Command prompt (Cmd.exe) for administration using the command line
- Windows PowerShell prompt for administration using Windows PowerShell
- File Signature Verification tool (Sigverif.exe) for verifying digital signatures of system files
- System Information (Msinfo32.exe) for getting system information
- Windows Installer (Msiexec.exe) for managing Windows Installer
- Date And Time control panel applet (Timedate.cpl) for viewing or setting the date, time, and time zone
- Region And Language control panel applet (Intl.cpl) for viewing or setting regional and language options, including formats and the keyboard layout
- Server Configuration utility (Sconfig.cmd), which provides a text-based menu system for managing a server's configuration

When you start a server with a Server Core installation, you can use the Windows Logon screen to unlock the server and log on. The first time you press Ctrl+Alt+Delete you are prompted to enter and confirm a password for the Administrator account. Afterward, when you press Ctrl+Alt+Delete to unlock the server, you are prompted to enter the password for the Administrator account or press Esc to switch users or log-in methods.

In a domain, the standard restrictions apply for logging on to servers, and anyone with appropriate user rights and logon permissions can log on to the server. On servers that are not acting as domain controllers and for servers in workgroup environments, you


can use the NET USER command to add users, and the NET LOCALGROUP command to add users to local groups for the purposes of logging on locally.

After you log on to a Server Core installation, you have a limited desktop environment with an administrator command prompt. You can use the command prompt for administration of the server and to run PowerShell. One way to open a PowerShell prompt is to enter powershell in the command prompt window.

While you are logged in, you can press Ctrl+Alt+Delete to display the Logon UI. You navigate this text-based menu using the Up/Down arrow keys or Tab to move between options and Enter to select an option. The menu has the following options:

- Change a password
- Lock
- Sign out
- Task Manager
- Cancel

**NOTE** When you are logged on locally, you'll also have the option to switch users. If you are working remotely using Remote Desktop, you can't send the Ctrl+Alt+Delete keystroke combination to the remote computer. Here, use Ctrl+Alt+End instead.




At the command prompt, you have all the standard commands and command-line utilities available for managing the server. However, commands, utilities, and programs run only when all of their dependencies are available in the Server Core installation. If you accidentally close the default command prompt, you can open a new command prompt by following these steps:

1. Press Ctrl+Alt+Delete (or Ctrl+Alt+End when working remotely) and then select Task Manager.




2. In Task Manager, click More Details.
3. On the File menu, click Run New Task.
4. In the Create New Task dialog box, enter **cmd.exe** in the Open box, and then click OK.



You can use this technique to open additional Command Prompt windows, too. Although you can work with Notepad and Regedit by entering **notepad.exe** or **regedit.exe** instead of **cmd.exe**, you can also start Notepad and Regedit directly from a command prompt by entering **notepad.exe** or **regedit.exe** as appropriate.

The Server Configuration utility (Sconfig.cmd) provides a text-based menu system that makes it easy to do the following:

- Configure domain or workgroup membership
- Change a server's name
- Add a local Administrator account
- Configure remote management features
- Configure Windows Update settings
- Download and install Windows updates
- Enable or disable Remote Desktop
- Configure network settings for TCP/IP
- Configure the date and time
- Activate Windows Server
- Change CEIP participation options
- Log off, restart, or shut down



Although a Server Core installation supports a limited set of roles and role services, you can install most features. Server Core also supports the .NET Framework and Windows Remote Management (WinRM). This support allows you to perform local and remote administration by using PowerShell. Some of the common tasks you might want to perform when you are logged on include the following:

**Cscript Scregedit.wsf** Configure the operating system. Use the /cli parameter to list available configuration areas. **ipconfig /all** List information about the computer's IP address configuration.

**Netdom RenameComputer** Set the server's name.

**Netdom Join** Join the server to a domain.

**Netsh** Provide multiple contexts for managing the configuration of networking components. Enter **netsh interface ipv4** to configure IPv4 settings. Enter **netsh interface ipv6** to configure IPv6 settings.

**Pnputil.exe** Install or update hardware device drivers.

**Sc query type=driver** List installed device drivers.

**Serverweroptin.exe** Configure Windows Error Reporting.

**Slmgr /ato** Windows Software Licensing Management tool used to activate the operating system. Runs `Cscript slmgr.vbs -ato`.

**Slmgr /ipk** Install or replace the product key. Runs `Cscript slmgr.vbs -ipk`.

**SystemInfo** List the system configuration details.

**Wecutil.exe** Create and manage subscriptions to forwarded events.

**Wevtutil.exe** View and search event logs.

**Winrm quickconfig** Configure the server to accept WS-Management requests from other computers. Runs `Cscript winrm.vbs quickconfig`. Enter without the quickconfig parameter to see other options.

**Wmic datafile where name="FullPath" get version** List a file's version.

**Wmic nicconfig index=9 call enabledhcp** Set the computer to use dynamic IP addressing rather than static IP addressing.

**Wmic nicconfig index=9 call enablestatic("IPAddress"), ("SubnetMask")** Set a computer's static IP address and network mask.

**Wmic nicconfig index=9 call setgateways("GatewayIPAddress")** Set or change the default gateway.

**Wmic product get name /value** List installed Windows Installer (MSI) applications by name.

**Wmic product where name="Name" call uninstall** Uninstall an MSI application.

**Wmic qfe list** List installed updates and hotfixes.

**Wusa.exe PatchName.msu /quiet** Apply an update or hotfix to the operating system.

## Managing Nano Server

Nano Server provides a fully custom environment with just the components your server needs to support the applications and services you want to deploy. You use Nano Server when you don't need 32-bit support, the graphical desktop interface or other graphical components.

Nano Server is designed to be managed remotely using either command-line or graphical tools. As Nano Server doesn't have the Windows logon interface, this means you can't log in to servers using this configuration as you do other servers. With a Nano Server installation, the only interface available when you log on locally is the Recovery Console.

Recovery Console provides features similar to the Logon UI for Server Core with options similar to `Sconfig.cmd`. With a keyboard and monitor connected to a nano server, you can log in to the recovery console using an administrator account. The console runs automatically and has a basic text-based menu that is navigated using



the keyboard. You use the arrow keys or Tab to navigate between options and Enter to select an option. To go back a screen or a page, you press Esc.

Although the console is meant to help you diagnose and resolve networking issues that prevent you from connecting to a server, you can use the recovery features at any time to work with basic settings of the server, including

- [Network adapters](#)
- [TCP/IP settings](#)
- [Server naming](#)

You also can use the console to restart or shutdown the server.

As Nano Server doesn't have the remote desktop, you can't remotely manage a server using this configuration in the same way as you can a full-server or server core installation. Not supporting remote desktop, however, isn't the same as not being able to use remote management tools or other options.

Although you can't remote in to a nano server and nano server doesn't have the graphical interfaces, you can use the graphical administration tools to remotely manage a server using this configuration. Thus, the easiest way to manage a nano server remotely is to install the Remote Server Administration Tools (RSAT) on your management computer and then use the tools to manage the server.

You also can manage a nano server using Windows PowerShell Remoting. Start by opening an administrator command prompt. If you haven't added the server as a trusted host on your computer, do so by entering the following command:

```
Set-Item WSMan:\localhost\Client\TrustedHosts NanoServerIPAddress
```

Where NanoServerIPAddress is the IP address of the nano server. You can then establish a remote session with the server. The basic syntax is:

```
$Cred = Get-Credential  
$Session = New-PSSession -ComputerName IPAddress -Credential $Cred
```

To establish a connection to the nano server, you must pass in your user name and password. This example stores credentials in a Credential object and then uses GetCredential to prompt for the required credentials. You also could specify the credentials explicitly, as shown here:

```
$Session = New-PSSession -ComputerName IPAddress
```

After you establish a session with the nano server, you must import the server-side PowerShell session into your client-side session by running the following command:

```
Import-PSSession $Session
```





You can then work with the nano server remotely. When you are finished, you should disconnect the remote shell from the nano server.

**NOTE** Beginning with Windows PowerShell 3.0, sessions are persistent by default. When you disconnect from a session, any command or scripts that are running in the session continue running, and you can later reconnect to the session to pick up where you left off. You also can reconnect to a session if you were disconnected unintentionally, such as by a temporary network outage.

To disconnect a session without stopping commands or releasing resources, run the following command:

```
Disconnect-PSSession $Session
```

The `$Session` object was instantiated when you created the session. As long as you don't exit the PowerShell window in which this object was created, you can use this object to reconnect to the session by entering:

```
Connect-PSSession $Session
```

When you are completely finished with the session, you should remove it. Removing a session stops any commands or scripts that are running, ends the session, and releases the resources the session was using. Remove a session by running the following command:

```
Remove-PSSession $Session
```

## Performing a Clean Installation

Before you start an installation, you need to consider whether you want to manage the computer's drives and partitions during the setup process. If you want to use the advanced drive setup options that Setup provides for creating and formatting partitions, you need to start the computer by using the distribution media. If you don't start by using the distribution media, these options won't be available, and you'll be able to manage disk partitions at a command prompt only by using the DiskPart utility.

You can perform a clean installation of Windows Server 2016 by following these steps:

1. Use one of the following techniques to start the Setup program:
- To perform a new installation, start the computer with the Windows Server 2016 distribution media in the computer's disc drive, and then press any key when prompted to start Setup from your media. If you are not prompted to boot from the disc drive, you might need to access the boot manager or select advanced boot options and then boot from the media rather than from the hard disk, or you might



need to change the computer's firmware settings to allow starting and loading the operating system from media.

- To perform a clean installation over an existing installation, you can boot from the distribution media, or you can start the computer and log on by using an account with administrator privileges. When you insert the Windows Server 2016 distribution media into the computer's disc drive, Setup should start automatically. If Setup doesn't start automatically and you are trying to boot from media, you might need to select advanced boot options or change firmware settings to allow booting from media. If you are logged in to an existing installation, you can start use File Explorer to access the distribution media, and then double-click Setup.exe.

```
F2 = System Setup
F10 = Lifecycle Controller (Config iDRAC, Update FW, Install OS)
F11 = Boot Manager
F12 = PXE Boot

Initializing Serial ATA devices...
Port A: TOSHIBA DT01ACA050
Port B: ST31000340AS
Port C: ST3500630AS
Port E: HL-DT-ST DVD-ROM DU90N
```

**NOTE** If you are installing Windows Server 2016 as a guest operating system on a host server, you'll need to allocate resources beforehand, including those for memory, processors and disks. Although the operating system can operate in a minimal installation with only 512 MB RAM, you may need to allocate more RAM than this initially to complete the setup process.

2. If you started the computer by using the distribution media, you'll next need to choose your language, time and currency formats, and keyboard layout when prompted. Only one keyboard layout is available during installation. If your keyboard language and the language edition of Windows Server 2016 you are installing are different, you might see unexpected characters as you type. To avoid this, be sure that you select the correct keyboard language. When you are ready to continue with the installation, click Next.





3. If you are starting the installation from an existing operating system and are connected to a network or the Internet, choose whether to get updates during the installation. Click either **Download And Install Updates** or **Not Right Now**.




4. With volume and enterprise licensed editions of Windows Server 2016, you might not need to provide a product key during installation. With retail editions, however, you need to enter a product key when prompted. Click **Next** to continue.



**NOTE** Don't worry about using the correct letter case or entering dashes. Setup enters all letters you type in uppercase. When a dash is needed, Setup enters the dash automatically. If you started Setup from an existing installation, you can copy the key from a document and then paste it into the Product key text box. If you enter an invalid product key, Setup will continue to display a warning and you'll need to correct the issue. Sometimes it's easier to re-enter the product key than try to identify a mistyped value.




5. Choose whether to perform a standard (Server Core) installation or a full-server (with Desktop Experience) installation. Keep in mind the edition of the operating system listed is based on the product key you entered previously. If you entered a product key for the wrong edition, you can go back and enter a product key for the edition you want to use. Click Next.



6. The license terms for Windows Server 2016 have changed from previous releases of Windows. After you review the license terms, click Accept to continue.
7. If you started the installation while logged on to an existing installation, you'll next see the Choose What To Keep page. Generally, you won't be able to keep any existing files, apps or settings and the Nothing selection will be your only option. Before continuing make sure you really don't need any files or other data on the system drive. If necessary, you can cancel the installation and restart Setup once you've backed up important files and data. Click Next and then click Install to begin the installation process. Skip the remaining steps.





8. If you booted from distribution media, you'll need to continue through the remaining steps of this procedure to complete the installation. When prompted, select the type of installation you want Setup to perform. Because you are performing a clean installation to replace an existing installation completely or configure a new computer, select Custom: Install Windows Only (Advanced) as the installation type.
9. You'll next be able to select the disk or disk and partition on which you want to install the operating system. There are two versions of the related page, so you need to keep the following in mind:
  - When a computer has a single hard disk with a single partition encompassing the whole disk or a single area of unallocated space, the whole disk partition is selected by default, and you can click Next to choose this as the install location and continue. With a disk that is completely unallocated, you might want to create the necessary partition before installing the operating system, as discussed in the "Managing Disk Partitions During Installation" section later in this chapter.
  - When a computer has multiple disks or a single disk with multiple partitions, you need to select an existing partition to use for installing the operating system or create a partition. You can create and manage partitions, as discussed in the "Managing Disk Partitions During Installation" section later in this chapter.
  - If a disk has not been initialized for use or if the firmware of the computer does not support starting the operating system from the selected disk, you need to initialize it



by creating one or more partitions on the disk. You cannot select or format a hard disk partition that uses FAT or FAT32 or has other incompatible settings. To work around this issue, you might want to convert the partition to NTFS.

- When working with this page, you can access a command prompt to perform any necessary preinstallation tasks. See the “Using the Command Line During Installation” section later in this chapter.
- 10. If the partition you select contains a previous Windows installation, Setup provides a prompt stating that existing user and application settings will be moved to a folder named Windows.old and that you must copy these settings to the new installation to use them. Click OK.
- 11. Click Next. Setup starts the installation of the operating system. During this procedure, Setup copies the full disk image of Windows Server 2016 to the location you selected and then expands it. Afterward, Setup installs features based on the computer’s configuration and the hardware it detects. This process requires several automatic restarts.

When Setup finishes the installation, the operating system will be loaded, and you can perform initial configuration tasks such as setting the administrator password and server name. If you encounter any issues during the installation, you may need to go through the Setup process again. See “Performing Additional Tasks During Installation” for tips and guidance.

**REAL WORLD** Servers running core installations of Windows Server are configured to use DHCP by default. As long as the server has a network card and a connected network cable, a Server Core installation should be able to connect to your organization’s DHCP servers and obtain the correct network settings. You can configure the server by using Sconfig, which provides menu options for configuring domain/workgroup membership, the computer name, remote management, Windows Update, Remote Desktop, network settings, date and time, logoff, restart, and shutdown.

Alternatively, you can configure the server by using individual commands. If you want to use a static IP address, use Netsh to apply the settings you want. After networking is configured correctly, enter **Slmgr –ipk** to set the product key and **Slmgr –ato** to activate Windows. Enter **timedate.cpl** to set the server’s date and time. If you want to enable remote management by using the WS-Management protocol, enter **winrm quickconfig**.

Next, you'll probably want to set the name of the computer. To view the default computer name, enter **echo %computername%**. To rename the computer, use Netdom

RenameComputer with the following syntax: **netdom renamecomputer currentname /newname:newname**, where currentname is the current name of the computer and newname is the name you want to assign. An example is **netdom renamecomputer win-k4m6bnovlhe /newname:server18**. You'll need to restart the computer, and you can do this by entering **shutdown /r**.

When the computer restarts, you can join it to a domain by using Netdom Join. For the syntax, enter **netdom join /?**.

As part of a server's post-installation configuration, you'll also want to configure networking settings. Chapter 18 provides a detailed discussion of the related options, including details on how to team adapters for load balancing. Once you've finalized the TCP/IP configuration, you may need to join the server to a domain and Windows Server 2016 provides several ways to do this.

In PowerShell, you can add a server to a domain using Add-Computer. The basic syntax is:

```
$Cred = Get-Credential  
Add-Computer -Domain DomainName -Credential $Cred
```

where DomainName is the name of the domain to which the local computer should be added. When you are joining a computer to a domain, you must provide the credentials of a user account that has permission to join computers to the specified domain.

**TIP** You don't have to be logged on to the server to add it to the domain using PowerShell. When working remotely, use the **-ComputerName** parameter to specify the name or IP address of the server to join to the domain.

If you have any problem adding the server to a domain, you may want to:

- Check the TCP/IP settings and ensure the IP address, subnet mask, gateway and DNS server addresses are correct.
- Check the date and time of the server to ensure these settings are correct and as close to world time as possible.

Sometimes, it may seem like you are doing everything right and you still not be able to join a server to a domain. In this case, you can perform an offline domain join that doesn't require the server to communicate directly with a domain controller. Here, you:



1. Run the following command on a computer that is a member of the domain to create a provisioning file:

```
djoin /provision /domain DomainName /machine TargetName /savefile FileName
```

where DomainName is the name of the domain to which the server should be added, TargetName is the name of the server, and FileName is the name for the provisioning file. The following example creates a provisioning file called fs97join.txt to allow FileServer97 to join the ImaginedLands.local domain:

```
djoin /provision /domain imaginedlands.local /machine fileserver97  
/savefile fs97-join.txt
```

2. When you run the /provision command, a computer account is created in the domain for the server and then a provisioning file is generated. Copy the provisioning file to the target server. This text-based file contains encrypted strings that allow the server to join the domain without using TCP/IP networking for authentication.

3. Run the following command to provision the target server:

```
djoin /requestodj /localos /loadfile FilePath /windowspath WinPath
```

Here, you perform an offline domain request of the local operating system by loading a provisioning file. FileName specifies the path and name of the provisioning file and WinPath specifies the location of the Windows directory on the server. The following example loads the provisioning file created previously and use the %SystemRoot% environment variable to specify the location of the Windows directory:

```
djoin /requestodj /localos /loadfile f:\fs97-join.txt /windowspath %systemroot%
```

4. After you restart the target server, the server will be a joined to the previously specified domain and you'll be able to logon to the server using a domain account. As no domain user credentials are cached initially, the server will need an active network connection for logon to occur.

Provisioning can fail for a variety of reasons. If a server with the same name as the one you want to join to the domain was previously provisioned or added to the domain, provisioning will fail and you'll need to confirm whether the name is actually available for use. If no other domain joined computer is actually using the name, you can force DJoin to re-provision the name and create a provisioning file by adding the /Reuse parameter, as shown in the following example:

```
djoin /provision /domain tvpress.local /machine webserver23  
/savefile ws23-join.txt /reuse
```



TIP Copying the provisioning file to a core server using a USB drive? Use DiskPart to determine the drive letter. Enter **diskpart** to start the utility and then enter **list volume** to list all available volumes.

## Performing an Upgrade Installation

You can upgrade servers running Windows Server 2012 R2 to Windows Server 2016. During an upgrade, Setup attempts to keep Windows settings, user files and apps. If the server has apps installed that don't support Windows Server 2016, you should uninstall the apps before continuing. Because of the challenges presented with upgrading a server hosting services and application software, it often is more efficient to migrate services and applications a server is hosting to other servers and then perform a clean installation.

You can perform an upgrade installation of Windows Server 2016 by following these steps:


1. Start the computer, and log on by using an account with administrator privileges. When you insert the Windows Server 2016 distribution media into the computer's DVD-ROM drive, Setup should start automatically. If Setup doesn't start automatically, use File Explorer to access the distribution media and then double-click Setup.exe.
2. Because you are starting Setup from the current operating system, you are not prompted to choose your language, time and currency formats, or keyboard layout, and only the current operating system's keyboard layout is available during installation. If your keyboard language and the language of the edition of Windows Server 2016 you are installing are different, unexpected characters might be displayed as you type.
3. If the computer is connected to a network or the Internet, choose whether to get updates during the installation. Click either Download And Install Updates or Not Right Now.




4. With volume-licensed and enterprise-licensed editions of Windows Server 2016, you might not need to provide a product key during installation of the



operating system. With retail editions, however, you are prompted to enter a product key. If you enter an invalid product key, Setup will display a warning and you'll need to correct the issue before you can continue.



5. Choose whether to perform a standard (Server Core) installation or a full-server (with Desktop Experience) installation. Keep in mind the edition of the operating system listed is based on the product key you entered previously. If you entered a product key for the wrong edition, you can go back and enter a product key for the edition you want to use. Click Next.



6. The license terms for Windows Server 2016 have changed from previous releases of Windows. After you review the license terms, click Accept to continue.



7. Next, specify that you want to keep files and apps, which tells Setup that you want to perform an upgrade. This ensures that Setup will try to keep Windows settings, user files and apps.



8. If there are any issues that need your attention, you'll be asked to take corrective action, confirm that you understand the issues or both before you can continue. After taking a corrective action, click Refresh to have Setup perform the preinstallation checks again. If there are no other issues, Setup will continue. If you are asked for confirmation, clicking Confirm will also continue the Setup process, provided there are no other issues that need your attention.
9. Click Install. Setup will then start the installation. Because you are upgrading the operating system, you do not need to choose an installation location. During this process, Setup copies the full disk image of Windows Server 2016 to the system disk. Afterward, Setup installs features based on the computer's configuration and the hardware it detects.



When Setup finishes the installation, the operating system will be loaded, and you can perform initial configuration tasks such as setting the administrator password and server name. If you encounter any issues during the installation, you may need to go through the Setup process again. See “Performing Additional Tasks During Installation” for tips and guidance.



# Performing Additional Tasks During Installation

Sometimes you might forget to perform a preinstallation task prior to starting the installation. Rather than restarting the computer, you can access a command prompt from Setup or use advanced drive options to perform the necessary tasks.





# Using the Command Line During Installation

When you access a command prompt from Setup, you access the Windows Preinstallation Environment (Windows PE) used by Setup to install the operating system. During installation, on the Where Do You Want To Install Windows page, you can access a command prompt by pressing Shift+F10. The Windows PE gives you access to many of the same command-line tools that are available in a standard installation of Windows Server 2016. These tools include:

- **ARP** Displays and modifies the IP-to-physical address translation tables used by the Address Resolution Protocol (ARP).
- **ASSOC** Displays and modifies file extension associations.
- **ATTRIB** Displays and changes file attributes.
- **CALL** Calls a script or script label as a procedure.
- **CD/CHDIR** Displays the name of or changes the current directory.
- **CHKDSK** Checks a disk for errors and displays a report.
- **CHKNTFS** Displays the status of volumes. Sets or excludes volumes from automatic system checking when the computer is started.
- **CHOICE** Creates a list from which users can select one of several choices in a batch script.
- **CLS** Clears the console window.
- **CMD** Starts a new instance of the Windows command shell.
- **COLOR** Sets the colors of the command-shell window.
- **CONVERT** Converts FAT volumes to NTFS.
- **COPY** Copies or combines files.
- **DATE** Displays or sets the system date.
- **DEL** Deletes one or more files.
- **DIR** Displays a list of files and subdirectories within a directory.
- **DISKPART** Invokes a text-mode command interpreter so that you can manage disks, partitions, and volumes by using a separate command prompt and commands that are internal to DISKPART.
- **DISM** Services and manages Windows images.
- **DOSKEY** Edits command lines, recalls Windows commands, and creates macros.
- **ECHO** Displays messages or turns command echoing on or off.
- **ENDLOCAL** Ends localization of environment changes in a batch file.
- **ERASE** Deletes one or more files.
- **EXIT** Exits the command interpreter.



- **EXPAND** Uncompresses files.
- **FIND** Searches for a text string in files.
- **FOR** Runs a specified command for each file in a set of files.
- **FORMAT** Formats a drive or USB flash device.
- **FTP** Transfers files.
- **FTYPE** Displays or modifies file types used in file- name extension associations.
- **GOTO** Directs the Windows command interpreter to a labeled line in a script.
- **HOSTNAME** Prints the computer's name.
- **IF** Performs conditional processing in batch programs.
- **IPCONFIG** Displays TCP/IP configuration.
- **LABEL** Creates, changes, or deletes the volume label of a disk.
- **MD/MKDIR** Creates a directory or subdirectory.
- **MORE** Displays output one screen at a time.
- **MOUNTVOL** Manages a volume mount point.
- **MOVE** Moves files from one directory to another directory on the same drive.
- **NBTSTAT** Displays the status of NetBIOS.
- **NET ACCOUNTS** Manages user account and password policies.
- **NET COMPUTER** Adds or removes computers from a domain.
- **NET CONFIG SERVER** Displays or modifies the configuration of a server service.
- **NET CONFIG WORKSTATION** Displays or modifies the configuration of a workstation service.
- **NET CONTINUE** Resumes a paused service.
- **NET FILE** Displays or manages open files on a server.
- **NET GROUP** Displays or manages global groups.
- **NET LOCALGROUP** Displays or manages local group accounts.
- **NET NAME** Displays or modifies recipients for messenger service messages.
- **NET PAUSE** Suspends a service.
- **NET PRINT** Displays or manages print jobs and shared queues.
- **NET SEND** Sends a messenger service message.
- **NET SESSION** Lists or disconnects sessions.
- **NET SHARE** Displays or manages shared printers and directories.
- **NET START** Lists or starts network services.
- **NET STATISTICS** Displays workstation and server statistics.
- **NET STOP** Stops services.
- **NET TIME** Displays or synchronizes network time.





- NET USE** Displays or manages remote connections.
- NET USER** Displays or manages local user accounts.
- NET VIEW** Displays network resources or computers.
- NETSH** Invokes a separate command prompt that allows you to manage the configuration of various network services on local and remote computers.
- **NETSTAT** Displays the status of network connections.
- **PATH** Displays or sets a search path for executable files in the current command window.
- **PATHPING** Traces routes, and provides packet-loss information.
- **PAUSE** Suspends the processing of a script, and waits for keyboard input.
- **PING** Determines whether a network connection can be established.
- **POPD** Changes to the directory stored by PUSH.D.
- **PRINT** Prints a text file.
- **PROMPT** Modifies the command prompt.
- **PUSH.D** Saves the current directory and then changes to a new directory.
- **RD/RMDIR** Removes a directory.
- **RECOVER** Recovers readable information from a bad or defective disk.
- **REG ADD** Adds a new subkey or entry to the registry.
- **REG COMPARE** Compares registry subkeys or entries.
- **REG COPY** Copies a registry entry to a specified key path on a local or remote system.
- **REG DELETE** Deletes a subkey or entries from the registry.
- **REG QUERY** Lists the entries under a key and the names of subkeys (if any).
- **REG RESTORE** Writes saved subkeys and entries back to the registry.
- **REG SAVE** Saves a copy of specified subkeys, entries, and values to a file.
- **REGSVR32** Registers and unregisters DLLs.
- **REM** Adds comments to scripts.
- **REN** Renames a file.
- **ROUTE** Manages network routing tables.
- **SET** Displays or modifies Windows environment variables. Also used to evaluate numeric expressions at the command line.
- **SETLOCAL** Begins the localization of environment changes in a batch file.
- **SFC** Scans and verifies protected system files.
- **SHIFT** Shifts the position of replaceable parameters in scripts.
- **START** Starts a new command-shell window to run a specified program or command.
- 





- **SUBST** Maps a path to a drive letter.
- **TIME** Displays or sets the system time.
- TITLE** Sets the title for the command-shell window.
- TRACERT** Displays the path between computers.
- TYPE** Displays the contents of a text file.
- VER** Displays the Windows version.

- 
- 
- 
- 





- **VERIFY** Tells Windows whether to verify that your files are written correctly to a disk.
- **VOL** Displays a disk volume label and serial number.





## Forcing Disk Partition Removal During Installation

During installation, you might be unable to select the hard disk you want to use. This issue can arise if the hard-disk partition contains an invalid byte offset value. To resolve this issue, you need to remove the partitions on the hard disk (which destroys all associated data) and then create the necessary partition by using the advanced options in the Setup program. During installation, on the Where Do You Want To Install Windows page, you can remove unrecognized hard-disk partitions by following these steps:

1. Press Shift+F10 to open a command prompt.
2. At the command prompt, enter **diskpart** . This starts the DiskPart utility.
3. To view a list of disks on the computer, enter **list disk** .
4. Select a disk by entering **select disk DiskNumber** , where DiskNumber is the number of the disk you want to work with.
5. To permanently remove the partitions on the selected disk, enter **clean** .
6. When the cleaning process is finished, enter **exit** to exit the DiskPart utility.
7. Enter **exit** to exit the command prompt.

**CAUTION** Do not clean a disk if there are any files or other data on the disk that may be needed. Cleaning a disk removes all partitions on the selected disk, which in turn causes all data on the disk to be lost.

To get Setup to recognize the changes you've made, you'll need to go back to the previous window and then continue with the installation process. You'll then be able to select the disk you previously cleaned and use it for installation. With the disk selected, click the Disk Options link to display the partition configuration options and then click New. In the Size box, set the size of the partition in megabytes, and then click Apply. This creates a partition on the disk you cleaned. Make sure the partition is at least 32 GB (32768 MB).

## Loading Device Drivers During Installation

During installation, on the Where Do You Want To Install Windows page, you can use the Load Driver option to load the device drivers for disks and disk controllers. Typically, you use this option when a storage subsystem or disk drive you want to use for installing the operating system isn't available for selection.

To load the device drivers and make the disk available, follow these steps:





1. During installation, on the Where Do You Want To Install Windows page, click Load Driver.
  2. When prompted, insert the installation media into a DVD drive or USB flash drive, and then click OK. Setup then searches the computer's removable media drives for the device drivers.
- If Setup finds multiple device drivers, select the driver to install, and then click Next.
  - If Setup doesn't find the device driver, click Browse to use the Browse For Folder dialog box to select the device driver to load, click OK, and then click Next.

You can click the Rescan button to have Setup rescan the computer's removable media drives for the device drivers. If the driver is found, you'll be able to select and install it. If you are unable to install a device driver successfully, click the back arrow button in the upper-left corner of the Install Windows dialog box to go back to the previous page and then try to continue the installation. If this doesn't work, you'll need to determine which disk or storage subsystem drivers are needed and obtain them before you can continue.

## Managing Disk Partitions During Installation

When you are performing a clean installation and have started the computer from the distribution media, the Where Do You Want To Install Windows page is displayed during setup. This page has options for working with drives and partitions, including:

- **New** Creates a partition. You must then format the partition.
- **Format** Formats a new partition so that you can use it for installing the operating system.
- **Delete** Deletes a partition that is no longer wanted.
- **Extend** Extends a partition to increase its size.

The sections that follow discuss how to use each of these options. If these options aren't available, you can still work with the computer's disks. On the Where Do You Want To Install Windows page, press Shift+F10 to open a command prompt. At the command prompt, enter **diskpart** to start the DiskPart utility.





## Creating Disk Partitions During Installation

Creating a partition allows you to set the partition's size. Because you can create new partitions only in areas of unallocated space on a disk, you might need to delete existing partitions to be able to create a partition of the size you want. After you create a partition, you can format the partition so that you can use it to install a file system. If you don't format a partition, you can still use it for installing the operating system. In this case, Setup formats the partition when you continue installing the operating system.

You can create a new partition by following these steps:

1. During installation, on the [Where Do You Want To Install Windows](#) page, click the disk on which you want to create the partition, and then click [New](#).
2. In the [Size](#) box, set the size of the partition in megabytes, and then click [Apply](#) to have Setup create a partition on the selected disk.

After you create a partition, you need to format the partition to continue with the installation.







# Formatting Disk Partitions During Installation

Formatting a partition creates a file system on the partition. When formatting is complete, you have a formatted partition on which you can install the operating system. Keep in mind that formatting a partition destroys all data on the partition. You should format existing partitions (rather than ones you just created) only when you want to remove an existing partition and all its contents so that you can start the installation from a freshly formatted partition.

You can format a partition by following these steps:

1. During installation, on the Where Do You Want To Install Windows page, click the partition that you want to format.
2. Click Format. When prompted to confirm that you want to format the partition, click OK. Setup then formats the partition.





## Deleting Disk Partitions During Installation

Deleting a partition removes a partition you no longer want or need. When Setup finishes deleting the partition, the disk space previously allocated to the partition becomes unallocated space on the disk. Deleting the partition destroys all data on the partition. Typically, you need to delete a partition only when it is in the wrong format or when you want to combine areas of free space on a disk.

You can delete a partition by following these steps:

1. During installation, on the Where Do You Want To Install Windows page, click the partition you want to delete.
2. Click Delete. When prompted to confirm that you want to delete the partition, click OK. Setup then deletes the partition.





## Extending Disk Partitions During Installation

Windows Server 2016 requires at least 10 GB of disk space for installation, and at least 32 GB of available disk space is recommended. If an existing partition is too small, you won't be able to use it to install the operating system. To resolve this, you can extend a partition to increase its size by using areas of contiguous, unallocated space on the current disk. You can extend a partition with an existing file system only if it is formatted with NTFS 5.2 or later. New partitions created in Setup can be extended also, provided that the disk on which you create the partition has unallocated space.

You can extend a partition by following these steps:

1. During installation, on the Where Do You Want To Install Windows page, click the partition you want to extend.
2. Click Extend. In the Size box, set the size of the partition in megabytes, and then click Apply to extend the selected partition.
3. When prompted to confirm that you want to extend the partition, click OK. Setup then extends the partition.







# Chapter 18. Implementing TCP/IP Networking

You enable networked computers to communicate by using the basic networking protocols built into Windows Server 2016. The key protocol you use is TCP/IP, which is a suite of protocols and services used for communicating over a network and is the primary protocol used for internetwork communications. Compared to configuring other networking protocols, configuring TCP/IP communications is fairly complicated, but TCP/IP is the most versatile protocol available.

**NOTE** Group Policy settings can affect your ability to install and manage TCP/IP networking. The key policies you should examine are in User Configuration\Administrative Templates\Network\Network Connections and Computer Configuration\Administrative Templates\System\Group Policy. Group Policy is discussed in Chapter 29, “Using Group Policy for Administration.”





# Navigating Networking Options

TCP/IP is the primary networking protocol used by Windows computers. Modern computers use IP version 4 (IPv4) and IP version 6 (IPv6) addressing.




# Networking Tools & Essentials

Windows Server 2016 has an extensive set of tools for managing networking, including:

- **Network Connections** Lists the available network adapters and their status.
- **Network And Sharing Center** Provides a central console for viewing and managing a computer's networking and sharing configuration.
- **Network Diagnostics** Provides automated diagnostics to help diagnose and resolve networking problems.


You use Network Connections to view and manage a server's network adapters. To access Network Connections, follow these steps:

1. Open Settings by clicking Start and then selecting Settings.
2. On the Settings home page, click Network & Internet and then click Ethernet.
3. On the Ethernet page, click Change Adapter Options.




While you are working with Network Connections, you can:

- Double-click a network adapter to display its status.
- Right-click a network adapter and then select Properties to view the related Properties dialog box.
- Right-click a network adapter and then select Diagnose to begin troubleshooting networking problems.



Network And Sharing Center provides an overview of the network. The value below the network name shows the category of the current network. The three categories of networks are defined as follows:

- **Domain network** Designates a network in which computers are connected to the corporate domain to which they are joined
- **Private network** Designates a network in which computers are configured as members of a homegroup or workgroup and are not connected directly to the public Internet
- **Public network** Designates a network in a public place, such as a coffee shop or an airport, rather than an internal network



The Access Type box specifies whether and how the computer is connected to its current network. Values for this option are No Network Access, No Internet Access, or Internet.




If you click the name of a network connection, you can display the related status dialog box.

You can access Network And Sharing Center by following these steps:

1. Open Settings by clicking Start and then selecting Settings.
2. On the Settings home page, click Network & Internet and then click Ethernet.
3. On the Ethernet page, click Network And Sharing Center.

Clicking Change Adapter Settings displays Network Connections, which you can use to manage network connections. Clicking Change Advanced Sharing Settings provides options for configuring the computer's sharing and discovery settings for each network profile.


However, with Windows Server 2016, network discovery generally is turned off for all network categories by default, while file and print sharing are turned on for all network categories by default. Thus, the category is mostly applicable to firewall settings, which are used to block or allow applications according to the network category.



You can use Windows Firewall to manage which apps are blocked or allowed by following these steps:

1. Open Settings by clicking Start and then selecting Settings.
2. On the Settings home page, click Network & Internet and then click Ethernet.
3. On the Ethernet page, click Windows Firewall. In the Windows Firewall app, the active network is listed as Connected.
4. Click Allow An App Or Feature Through Windows Firewall. Then on the Allowed Apps page, click Change Settings.

5. Use the options provided to specify which apps and features are allowed. If an app or feature is allowed for a particular network category, select the related checkbox. Otherwise, clear the related checkbox.



When you connect a computer's network adapter to a network for the first time, Windows sets the network category based on the configuration of the computer. Typically, you will find that a network adapter is set as public before you join a computer to a domain.

If the category is set incorrectly, one way to fix the problem is to click Troubleshoot Problems while working with Network And Sharing Center and then follow the prompts to allow the network troubleshooter to detect the problem and help you resolve it. Another way to resolve the problem is to use Windows PowerShell to make the change.

For example, the default alias for the first network connection on a device typically is "Ethernet". If so, you can use the following command to change the network category of this connection to private:

```
Get-NetConnectionProfile -InterfaceAlias "Ethernet" | Set-NetConnectionProfile -NetworkCategory Private
```

From Network And Sharing Center, you can attempt to diagnose a networking problem.

To do this, click Troubleshoot Problems, and then click a troubleshooter to run, such as

Incoming Connections or Network Adapter, and then follow the prompts. Windows Network Diagnostics then attempts to identify the network problem and provide a possible solution.

# Network Interfaces & Teaming

A network connection is created automatically if a computer has a network adapter and is connected to a network. If a computer has multiple network adapters, one network connection is created for each adapter that is connected to a network. Each network connection can be configured in the following ways:

- **Manually** IP addresses that are assigned manually are called static IP addresses . Static IP addresses are fixed and don't change unless you change them. You usually assign static IP addresses to Windows servers, and when you do this, you need to configure additional information to help the server navigate the network.
- **Dynamically** A DHCP server (if one is installed on the network) assigns dynamic IP addresses at startup, and the addresses might change over time. Dynamic IP addressing is the default configuration.
- **Alternate addresses (IPv4 only)** When a computer is configured to use DHCP and no DHCP server is available, Windows Server 2016 assigns an alternate private IP address automatically. The alternate address is in the range 169.254.0.1 to 169.254.255.254 with a subnet mask of 255.255.0.0. You can also specify a userconfigured alternate address.

By default, servers use dynamic IPv4 addressing, if a DHCPv4 server is available, or autoconfiguration using alternate addressing otherwise. The alternate address is in the range 169.254.0.1 to 169.254.255.254 with a subnet mask of 255.255.0.0. You can also specify a user-configured alternate address.


Unlike IPv4, a DHCPv6 server is not required to dynamically configured IPv6. IPv6 can autoconfigure itself without DHCP or any other stateful configuration protocol. Computers autoconfigure link-local addresses for each network adapter configured to use IPv6 through router discovery, which also sets the addresses of routers and other essential configuration parameters automatically.

**TIP** The default link-local address is based on the FE80::/64 prefix and the 64bit identifier of the network adapter itself. This default address is used as long as it is unique and valid for the network. However, if for some reason the default address is not unique, you will need to manually specify another IPv6 address.


You can configure the TCP/IP settings for each individual network adapter available on a sever by following these steps:

1. **Open Settings.** On the Settings home page, click Network & Internet and then click Ethernet.

2. On the Ethernet page, click Change Adapter Options. In Network Connections, right-click the connection with which you want to work, and then click Properties. This displays a Properties dialog box for the connection.
3. In the Properties dialog box for the network connection, be sure that Internet Protocol Version 6 (TCP/IPv6), Internet Protocol Version 4 (TCP/IPv4), or both are selected.
4. As appropriate, follow the instructions in the next section for configuring static or dynamic IP addressing.




Although you generally must configure each network adapter available on a server separately, network interface teaming (NIC teaming) allows you to group adapters together for the purposes of load balancing. If a server has teamed adapters, you'll find that IPv4 and IPv6 are disabled on individual adapters and that the Microsoft Network Adapter Multiplexor Protocol is enabled instead.




The reason for this is that you configure TCP/IP addressing for teamed adapters by using the multiplexed interface rather than individual adapters. The IP addressing options available are the same as those for standard adapters. This means you use the IPv4 options in the adapter's properties dialog box to configure IPv4 and the related IPv6 options to configure IPv6.







Other options of the multiplexed adapter, however, such as whether load balancing is enabled, should be configured by using the NIC Teaming interface. In Server Manager, you can access the NIC Teaming interface by clicking the related Enabled or Disabled link on the Local Server page.



You also can access the NIC Teaming interface by entering the following in the Search box or at a prompt:

Lbfoadmin.exe /servers .



The `/servers` command passed to `Lbfoadmin` specifies the server or servers you want to work with. In the example, you use a period (.) to specify that you want to work with the local server. You also could pass in a comma-separated list of servers to work with, such as:

```
Lbfoadmin.exe /servers corpserver28, corpserver31
```





# Configuring TCP/IP Networking

As discussed previously, you configure TCP/IP settings using the Properties dialog box for the network adapter. If a server has multiple adapters and those adapters are teamed, you configure network teaming settings using the NIC Teaming app.







# Configuring Static IP Addresses

When you assign a static IP address, you need to tell the computer the IP address you want to use, the subnet mask for this IP address, and, if necessary, the default gateway to use for internetwork communications. An IP address is a numeric identifier for a computer. IP addressing schemes vary according to how your network is configured, but they're typically assigned based on a particular network segment.

IPv6 addresses and IPv4 addresses are very different. With IPv6, the first 64 bits represent the network ID and the remaining 64 bits represent the network interface. With IPv4, a variable number of the initial bits represent the network ID and the rest of the bits represent the host ID. For example, if you're working with IPv4 and a computer on the network segment 10.0.10.0 with a subnet mask of 255.255.255.0, the first three octets (8-bit groups) represent the network ID, and the address range you have available for computer hosts is 10.0.10.1 to 10.0.10.254. In this range, the address 10.0.10.255 is reserved for network broadcasts.

If you're on a private network that is indirectly connected to the Internet, you should use private IPv4 addresses. Table 18-1 summarizes private network IPv4 addresses.

**TABLE 18-1** Private IPv4 network addressing

| PRIVATE<br>NETWORK<br>ID | SUBNET<br>MASK | NETWORK<br>ADDRESS<br>RANGE     |
|--------------------------|----------------|---------------------------------|
| 10.0.0.0                 | 255.0.0.0      | 10.0.0.0–<br>10.255.255.255     |
| 172.16.0.0               | 255.240.0.0    | 172.16.0.0–<br>172.31.255.255   |
| 192.168.0.0              | 255.255.0.0    | 192.168.0.0–<br>192.168.255.255 |

All other IPv4 network addresses are public and must be leased or purchased. If the network is connected directly to the Internet and you've obtained a range of IPv4 addresses from your Internet service provider, you can use the IPv4 addresses you've been assigned.





## Using the Ping Command to Check an Address

Before you assign a static IP address, you should make sure that the address isn't already in use or reserved for use with DHCP. With the ping command, you can check whether an address is in use. Open a command prompt and enter **ping**, followed by the IP address you want to check.

To test the IPv4 address 10.0.10.12, you would use the following command:

```
ping 10.0.10.12
```

To test the IPv6 address FEC0::02BC:FF:BECB:FE4F:961D, you would use the following command:

```
ping FEC0::02BC:FF:BECB:FE4F:961D
```


If you receive a successful reply from the ping test, the IP address is in use and you should try another one. If the request times out for all four ping attempts, the IP address isn't active on the network at this time and probably isn't in use; however, a firewall could be blocking your ping request. Your company's network administrator would also be able to confirm whether an IP address is in use.

## Configuring a Static IPv4 or IPv6 Address

One network connection is available for each network adapter installed. These connections are created automatically. To configure static IP addresses for a particular connection, follow these steps:


1. In Network And Sharing Center, click Change Adapter Settings. In Network Connections, right-click the connection with which you want to work, and then click Properties.
  2. Double-click Internet Protocol Version 6 (TCP/IPv6) or Internet Protocol Version 4 (TCP/IPv4) as appropriate for the type of IP address you are configuring.
  3. For an IPv6 address, do the following:
    - Click Use The Following IPv6 Address, and then enter the IPv6 address in the IPv6 Address text box. The IPv6 address you assign to the computer must not be in use anywhere else on the network.
    - The Subnet Prefix Length option ensures that the computer communicates over the network properly. Windows Server 2016 should insert a default value for the subnet prefix into the Subnet Prefix Length text box. If the network doesn't use variablelength subnetting, the default value should suffice, but if it does use
- 

variable-length subnets, you need to change this value as appropriate for your network.



4. For an IPv4 address, do the following:

- Click Use The Following IP Address, and then enter the IPv4 address in the IP Address text box. The IPv4 address you assign to the computer must not be in use anywhere else on the network.
- The Subnet Mask option ensures that the computer communicates over the network properly. Windows Server 2016 should insert a default value for the subnet mask into the Subnet Mask text box. If the network doesn't use variable-length subnetting, the default value should suffice, but if it does use variable-length subnets, you need to change this value as appropriate for your network.



5. If the computer needs to access other TCP/IP networks, the Internet, or other subnets, you must specify a default gateway. Enter the IP address of the network's default router in the Default Gateway text box.
6. Domain Name System (DNS) is needed for domain name resolution. Enter a preferred address and an alternate DNS server address in the text boxes provided.
7. When you have finished, click OK twice. Repeat this process for other network adapters and IP protocols you want to configure.
8. With IPv4 addressing, configure WINS as necessary.


## Configuring Dynamic IP Addresses and Alternate IP Addressing

Although most servers have static IP addresses, you can configure servers to use dynamic addressing, alternate IP addressing, or both. You configure dynamic and alternate addressing by following these steps:

9. In Network And Sharing Center, click Change Adapter Settings. In Network Connections, one LAN connection is shown for each network adapter installed. These connections are created automatically. If you don't find a

LAN connection for an installed adapter, check the driver for the adapter. It might be installed incorrectly. Right-click the connection with which you want to work, and then click Properties.

10. Double-click Internet Protocol Version 6 (TCP/IPv6) or Internet Protocol Version 4 (TCP/IPv4) as appropriate for the type of IP address you are configuring.
11. Select Obtain An IPv6 Address Automatically or Obtain An IP Address Automatically as appropriate for the type of IP address you are configuring. You can select Obtain DNS Server Address Automatically, or you can select Use The Following DNS Server Addresses, and then enter a preferred and alternate DNS server address in the text boxes provided.



12. When you use dynamic IPv4 addressing, you can configure an automatic alternate address or manually configure the alternate address. To use an automatic configuration, on the Alternate Configuration tab, select Automatic Private IP Address. Click OK, click Close, and then skip the remaining step.
13. To use a manual configuration, on the Alternate Configuration tab, select User Configured, and then enter the IP address you want to use in the IP Address text box. The IP address you assign to the computer should be a private IP

address, and it must not be in use anywhere else when the settings are applied. Complete the alternate configuration by entering a subnet mask, default gateway, DNS server, and Windows Internet Name Service (WINS) settings. When you have finished, click OK, and then click Close.

The image shows a screenshot of the 'Internet Protocol Version 4 (TCP/IPv4) Properties' dialog box. The 'Alternate Configuration' tab is selected. The dialog box contains the following elements:

- General** and **Alternate Configuration** tabs.
- Text: 'If this computer is used on more than one network, enter the alternate IP settings below.'
- Two radio buttons:   
☐ Automatic private IP address   
☒ User configured
- Five input fields for user-configured settings, each with a placeholder 'x' for each octet:   
IP address:   
Subnet mask:   
Default gateway:   
Preferred DNS server:   
Alternate DNS server:



## Configuring Multiple Gateways


To provide fault tolerance in case of a router outage, you can choose to configure Windows Server 2016 computers so that they use multiple default gateways. When you assign multiple gateways, Windows Server 2016 uses the gateway metric to determine which gateway is used and at what time. The gateway metric indicates the routing cost of using a gateway. The gateway with the lowest routing cost, or metric, is used first. If the computer can't communicate with this gateway, Windows Server 2016 tries to use the gateway with the next lowest metric.

The best way to configure multiple gateways depends on the configuration of your network. If your organization's computers use DHCP, you probably want to configure the additional gateways through settings on the DHCP server. If computers use static IP addresses or you want to set gateways specifically, assign them by following these steps:

1. In Network And Sharing Center, click Change Adapter Settings. In Network Connections, right-click the connection with which you want to work, and then click Properties.
2. Double-click Internet Protocol Version 6 (TCP/IPv6) or Internet Protocol Version 4 (TCP/IPv4) as appropriate for the type of IP address you are configuring.
3. Click Advanced to open the Advanced TCP/IP Settings dialog box.







4. The Default Gateways panel shows the current gateways that have been manually configured (if any). You can enter additional default gateways as necessary. Click Add, and then enter the gateway address in the Gateway text box. By default, Windows Server 2016 automatically assigns a metric to the gateway. You can also assign the metric yourself. To do this, clear the Automatic Metric check box, enter a metric in the text box provided, and then click Add. Repeat this process for each gateway you want to add.
5. Click OK, and then click Close.

## Configuring Networking for Hyper-V

After you install Hyper-V and create an external virtual network, your server uses a virtual network adapter to connect to the physical network. When you work with the Network Connections page, you will find the original network adapter and a new virtual network adapter. The original network adapter will have nothing bound to it except the Microsoft Virtual Network Switch Protocol, and the virtual network adapter will have all the standard protocols and services bound to it. The virtual network adapter that appears under Network Connections will have the same name as the virtual network switch with which it is associated.





**NOTE** As part of the Hyper-V configuration, you can create an internal virtual network, which enables communications only between the server and hosted virtual machines. This configuration exposes a virtual network adapter to the parent server without the need to have a physical network adapter associated with it and isolates the virtual machine from the Internet and the rest of the LAN. HyperV binds the virtual network service to a physical network adapter only when an external virtual network is created. An external virtual network is required for communications on the LAN and the Internet.


Following this, when you install Hyper-V on a server and enable external virtual networking, you'll find that virtual network switching is being used. The server will have a network connection with the Hyper-V Extensible Virtual Switch Protocol enabled and all other networking components not enabled and an entry for a virtual connection with the key networking components enabled and the Hyper-V Extensible Virtual Switch Protocol disabled. This is the configuration you want to use to ensure proper communications for the server and any hosted virtual machines that use networking. If this configuration is changed, virtual machines won't be able to connect to the external network.

## Using NIC Teaming

NIC teaming allows you to optimize the way a server uses multiple network interfaces. When you add interfaces to a team, you group them together for load balancing, designating a standby interface or both. By load balancing interfaces, you can increase the available bandwidth for networking and reduce the likelihood that the capacity of a server's network card is a limiting factor with respect to overall performance. By designating a standby interface, you can ensure user can connect to a server in the event an adapter fails.

# Understanding NIC Teaming

By default, NIC teaming is disabled on Windows servers. A NIC team can be configured in a switch dependent or independent mode. A switch dependent mode requires configuration of a switch and the server to identify which links are part of the team. A switch independent mode only requires configuration of a server to identify the interfaces associated with the team.



**REAL WORLD** The benefit of using switch dependent teaming is that the overhead associated with load balancing is shifted to the switch, rather than requiring the server to manage the load balancing. The drawback is that you typically need enterprise-class switches and must either enable teaming on the switches or specifically configure the switches for teaming.

Up to 32 interfaces can be added to a team. Windows Server 2016 supports two switch dependent teaming modes:

- LACP, also referred to as dynamic teaming
- Static teaming, also referred to as generic teaming



Dynamic teaming, which is supported by most enterprise-class switches and based on IEEE 802.1ax, allows for the automatic creation of a team. Here, Windows Server dynamically identifies links between a server and a particular switch using the Link Aggregation Control Protocol (LACP) and adds them to the team automatically. Generally, enabling LACP on the switch and then creating a team that uses dynamic teaming is all you need to do to configure this mode.

Static teaming, which is also supported by most enterprise-class switches and based on IEEE 802.3ad, requires manual configuration of the server and switch to identify the links that are part of the team. Here, you configure the switch to use static teaming and then creating a team that uses static teaming for specific interfaces on the server.

Windows Server 2016 also supports a switch independent mode, which requires only that you creating a team on the server for specific interfaces. As switches don't participate in the team when using this mode, links connected to multiple switches can be used if necessary.

Regardless of which teaming mode a server uses, you have three choices for load balancing across the participating interfaces:

- [Dynamic](#)
- [Hyper-V port](#)
- [Address hash](#)

Each traffic distribution algorithm has its pros and cons. Dynamic load balancing is the default mode and it is arguably the most efficient when you want to try to distribute the load evenly across all available interfaces. With dynamic load balancing, traffic distribution also is similar regardless of whether you use a switch dependent or switch independent mode as well.

Hyper-V port traffic distribution is based on the MAC address of a particular virtual machine hosted on a server or the port to which the virtual machine is connected. If a server has many virtual machines, this approach can be useful to achieve balanced distribution over the interfaces used in the team. However, if a server has few virtual machines or several high bandwidth using virtual machines, traffic distribution may favor a select few interfaces rather than all interfaces and thus may not be balanced appropriately.

With Hyper-V port distribution, the available bandwidth of any particular MAC address or Hyper-V port is always limited to the bandwidth available on a single interface. Additionally, it's important to note that inbound and outbound traffic for the virtual machine occur over the same interface. Thus, you should use this distribution





mode when the number of virtual machines on the server is a multiple to the number of available interfaces and the sharing of inbound and outbound traffic for a particular virtual machine won't be a limiting factor.



Address hash distribution is based on a hash from the address components of a TCP or UDP packet stream, including the source and destination MAC addresses, the source and destination IP addresses, and the source and destination TCP or UDP ports. If a server has a heavy outbound workload, this approach can achieve good load balancing across all interfaces in a team. However, as an IP address can only be associated with a specific, designated MAC address for routing, all inbound traffic is received only on the primary interface of the team. Thus, inbound traffic is limited in terms of total available bandwidth and cannot exceed the bandwidth of the primary interface. With these limitations, you'll find that address hash distribution is best used when you have heavy outbound traffic and light inbound traffic.



# Configuring NIC Teaming

You enable teaming by creating a new team using either the NIC Teaming app or the cmdlets associated with the Netlbfo module in Windows Powershell. To list all of the cmdlets for the Netlbfo module, enter **Get-Command -Module Netlbfo** at a PowerShell prompt.

In Server Manager, access the NIC Teaming app by clicking the related Enabled or Disabled link on the Local Server page. Alternatively, while working with Windows Server 2016, enter **Lbfoadmin.exe /servers .** in the Search box or at a command prompt.



In most cases, you'll find it's easier to work directly with the server you want to configure for NIC teaming by logging on locally or accessing the server using Remote Desktop. That said, the NIC Teaming interface does support configuration of remote servers and virtual machines running on a host server. If the server or virtual machine you want to work with is listed on the Servers panel, click on it to select it. Otherwise, add the server by selecting the related Tasks option, clicking Add Servers and then using the dialog box provided to specify the name or IP address of the server or servers you want to work with. You can add a server by name on the Active Directory panel or by name or IP address on the DNS panel. For example, on the Active Directory panel, type the computer name or fully qualified domain name of a remote Windows server that is in a trusted domain and then click Find Now. In the Name list, double-click a server to add it to the Selected list.

Once you've selected the server you want to work with on the Servers panel, you can use the option on the Teams panel to create the NIC team. Select the related Tasks button and then select New Team. In the New Team dialog box, type a name of the team, such as Primary NIC or Primary NIC Team and select the interfaces that should be part of the team.

NIC Teaming

New team

Team name:

Primary NIC


Member adapters:

| In Team                             | Adapter | Speed  | State | Reason |
|-------------------------------------|---------|--------|-------|--------|
| <input checked="" type="checkbox"/> | NIC1    | 1 Gbps |       |        |
| <input checked="" type="checkbox"/> | NIC2    | 1 Gbps |       |        |

Additional properties

Next, expand the Additional Properties panel by clicking on the panel heading. You'll then be able to specify the teaming and load balancing modes. By default, all network interfaces are active. If the team should have a standby adapter for failover purposes,

click the Standby Adapter option and then select the standby adapter. Only one interface can be designated as the standby.



Additional properties

Teaming mode: Switch Independent

Load balancing mode: Dynamic

Standby adapter: None (all adapters Active)

Primary team interface: [Primary NIC; Default VLAN](#)

Close the New Team dialog box and create the team by clicking OK. If you are remotely accessing the server that you are configuring, your connection with the server will be lost and you will need to reestablish your connection.

When you create the team, the interfaces that are part of the team will be configured as discussed previously. This means the Microsoft Network Adapter Multiplexor Protocol will be enabled on the interfaces while all other protocols are disabled. A multiplexed interface will also be added and configured for the server. Generally, this interface will have the same TCP/IP settings as the server's primary adapter had previously. Open the Properties dialog box for the multiplexed adapter to view or change the TCP/IP settings.



# Chapter 19. Data Storage: The Essentials

A hard drive is the most common storage device used on network workstations and servers. Users depend on hard drives to store their documents, spreadsheets, and other types of data. Drives are organized into file systems that users can access either locally or remotely.

Local file systems are installed on a user's computer and can be accessed without remote network connections. The C drive, which is available on most workstations and servers, is an example of a local file system. You access the C drive by using the file path C:\.

On the other hand, you access remote file systems through a network connection to a remote resource. You can connect to a remote file system by using the Map Network Drive feature of File Explorer.

Wherever disk resources are located, your job as an IT professional is to manage them. The tools and techniques you use to add, prepare and work with drives are discussed in this chapter. Later chapters discuss techniques for partitioning drives, implementing storage management and more.

## Deploying File Services

A file server provides a central location for storing and sharing files across the network. When many users require access to the same files and application data, you should configure file servers in the domain. Although all servers are configured with basic file and storage services, you must still configure any additional services and features that may be needed using the File And Storage Services role.

The role services associated with the File And Storage Services include:

- **BranchCache For Network Files** Enables computers in a branch office to cache commonly used files from shared folders. It takes advantage of data deduplication techniques to optimize data transfers over the wide area networks (WAN) to branch offices.
- **Data Deduplication** Uses subfile variable-size chunking and compression to achieve greater storage efficiency. This works by segmenting files into 32-KB to 128KB chunks, identifying duplicate chunks, and replacing the duplicates with references to a single copy. Optimized files are stored as reparse points. After deduplication, files on the volume are no longer stored as data streams and instead are replaced with stubs that point to data blocks within a common chunk store.







- **DFS Namespaces** Enables you to group shared folders located on different servers into one or more logically structured namespaces. Each namespace appears as a single shared folder with a series of subfolders; however, the underlying structure of a namespace can come from shared folders on multiple servers in different sites.
- **DFS Replication** Enables you to synchronize folders on multiple servers across local or WAN connections by using a multimaster replication engine. The replication engine uses the Remote Differential Compression (RDC) protocol to synchronize only the portions of files that have changed since the last replication. You can use DFS Replication with DFS Namespaces or by itself. When a domain is running in a Windows 2008 domain functional level or higher, domain controllers use DFS Replication to provide more robust and granular replication of the SYSVOL directory.
- **File Server** Enables you to manage file shares that users can access over the network.
- **File Server Resource Manager (FSRM)** Installs a suite of tools that administrators can use to better manage data stored on servers. By using FSRM, administrators can generate storage reports, configure quotas, and define filescreening policies.
- **File Server VSS Agent Service** Enables VSS-aware backup utilities to create consistent shadow copies (snapshots) of applications that store data files on the file server.
- **iSCSI Target Server** Turns any Windows Server into a network-accessible block storage device, which can be used for testing of applications prior to deploying storage area network (SAN) storage. It supports shared storage on both Windows iSCSI initiators and those iSCSI initiators that are not based on Windows as well as network/diskless boot for diskless servers.
- **iSCSI Target Storage Provider** Supports managing iSCSI virtual disks and shadow copies (snapshots) from an iSCSI initiator.
- **Server for NFS** Provides a file-sharing solution for enterprises with a mixed Windows and UNIX environment. When you install Server for Network File System (NFS), users can transfer files between Windows Server and UNIX operating systems by using the NFS protocol.
- **Storage Services** Enables you to manage storage, including storage pools and storage spaces. Storage pools group disks so that you can create virtual disks from the available capacity. Each virtual disk you create is a storage space.
- **Work Folders** Enables users to synchronize their corporate data to their devices and vice versa. Those devices can be joined to the corporate domain or a workplace.



When you add any needed role services to a file server, you might also want to install the following optional features, available through the Add Roles And Features Wizard:

- **BitLocker Drive Encryption** Prevents unauthorized access to data by wrapping entire drives or only the used portion of volumes in tamper-proof encryption. Used with Trust Platform Module (TPM), BitLocker can validate the integrity of a server's boot manager and boot files at startup to guarantee that a hard disk has not been tampered with while the operating system was offline.
- **BitLocker Network Unlock** Enables a protected system volume to be automatically unlocked on startup when a computer is joined to a domain and using TPM.
- **Enhanced Storage** Supports additional functions made available by devices that support hardware encryption and enhanced storage. Enhanced storage devices support Institute of Electrical and Electronics Engineers (IEEE) standard 1667 to provide enhanced security, which can include authentication at the hardware level of the storage device.
- **Multipath I/O** Provides support for using multiple data paths between a file server and a storage device. Servers use multiple I/O paths for redundancy in case of the failure of a path and to improve transfer performance.
- **SMB Bandwidth Limit** Tracks SMB traffic by source category, either default, virtual machine or Live Migration, and allows you to limit the amount of traffic for each category by setting a bytes per second limitation.
- **Storage Replica** Enables block-level replication between servers, with or without clustering. When you install and configure storage replica, you can mirror the data in physical sites to safeguard data at the filesystem level using a synchronous or asynchronous technique.
- **Windows Server Backup** The standard backup utility included with Windows Server 2016.

Binaries needed to install roles and features are referred to as payloads . With Windows


Server 2016, payloads are stored in subfolders of the %SystemDrive%\Windows\WinSXS folder. If the binaries for the tools have been removed, you might need to install the tools by specifying a source.

**IMPORTANT** If payloads have been removed and you don't specify a source, payloads are restored via Windows Update by default. However, Group Policy can be used to control whether Windows Update is used to restore payloads and

to provide alternate source paths for restoring payloads. The policy with which you want to work is Specify Settings For Optional Component Installation And Component Repair, which is under Computer Configuration\Administrative Templates\System. This policy also is used for obtaining payloads needed to repair components.

You can configure the File And Storage Services role ON a server by following these steps:

1. In Server Manager, click Manage, and then click Add Roles And Features, or select Add Roles And Features in the Quick Start pane. This starts the Add Roles And Features Wizard. If the wizard displays the Before You Begin page, read the Welcome text, and then click Next.
2. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Click Next.




3. On the Server Selection page, you can choose to install roles and features on running servers or virtual hard disks. Either select a server from the server pool or select a server from the server pool on which to mount a virtual hard disk (VHD). If you are adding roles and features to a VHD, click Browse and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, click Next.




NOTE Only servers that are running Windows Server 2012 or later and that have been added for management in Server Manager are listed.

- On the Server Roles page, select File And Storage Services. Expand the related node, and select the additional file services to install. If additional features are required to install a role, you'll see an additional dialog box. Click Add Features to close the dialog box and add the required features to the server installation. When you are ready to continue, click Next.



- On the Features page, select the features you want to install. If additional functionality are required to install a feature you selected, you'll see an additional dialog box. Click Add Features to close the dialog box and add the required features to the server installation. When you are ready to continue, click Next. Depending on the added feature, there may be additional steps before you get to the Confirm page.



- 
6. Optionally, on the Confirm page, click the [Export Configuration Settings](#) link to generate an installation report that can be displayed in Internet Explorer.

**REAL WORLD** If the server on which you want to install roles or features doesn't have all the required binary source files, the server gets the files via Windows Update by default or from a location specified in Group Policy.

You can also specify an alternate path for the required source files. To do this, click the [Specify An Alternate Source Path](#) link, enter that alternate path in the box provided, and then click OK. For network shares, enter the UNC path to the share, such as \\CorpServer25\\WinServer2016\\. For mounted Windows images, enter the WIM path prefixed with WIM: and including the index of the image to use, such as WIM:\\CorpServer25\\WinServer2016\\install.wim:4.

7. After you review the installation options and save them as necessary, click [Install](#) to begin the installation process. The [Installation Progress](#) page tracks the progress of the installation. If you close the wizard, click the [Notifications](#) icon in Server Manager, and then click the link provided to reopen the wizard.
8. When Setup finishes installing the server with the roles and features you selected, the [Installation Progress](#) page will be updated to reflect this. Review the installation details to ensure that all phases of the installation were completed successfully.

Note any additional actions that might be required to complete the installation, such as restarting the server or performing additional installation tasks.

If any portion of the installation failed, note the reason for the failure. Review the Server Manager entries for installation problems, and take corrective actions as appropriate.

## Adding Drives

Windows Server 2016 supports both hard disk drives and solid state drives. Before you make a drive available to users, you need to configure it and consider how it will be used. With Windows Server 2016, you can configure drives in a variety of ways. The technique you choose depends primarily on the type of data with which you're working and the needs of your network environment. For general user data stored on workstations, you might want to configure individual drives as standalone storage devices. In that case, user data is stored on a workstation's drive, where it can be accessed and stored locally.



Although storing data on a single drive is convenient, it isn't the most reliable way to store data. To improve reliability and performance, you might want a set of drives to work together. Windows Server 2016 supports drive sets and arrays by using various storage technologies, including Storage Spaces.





# Physical Drives

Whether you use individual drives or drive sets, you need physical drives. Physical drives are the actual hardware devices that are used to store data. The amount of data a drive can store depends on its size and whether it uses compression. Windows Server 2016 supports both Standard Format and Advanced Format drives. Standard Format drives have 512 bytes per physical sector and are also referred to as 512 native sector or 512b drives. Advanced Format drives have 4096 bytes per physical sector and are available in both 4K native and 512 emulation versions. 4K native and 512e drives represent a significant shift for the drive industry, and it allows for large, multiterabyte drives.

Disks perform physical media updates in the granularity of their physical sector size. 512b disks work with data 512 bytes at a time; native 4K and 512e disks work with data 4096 bytes at a time. At an administrator prompt, you can use the command-line utility Fstutil to determine bytes per physical sector by entering the following:

```
Fstutil fsinfo ntfsinfo DriveDesignator
```

where DriveDesignator is the designator of the drive to check, such as:

```
Fstutil fsinfo sectorinfo c:
```



```
Administrator: Command Prompt
Microsoft Windows [Version 10.0.14300]
(c) 2016 Microsoft Corporation. All rights reserved.

C:\Windows\system32>Fstutil fsinfo sectorinfo c:
LogicalBytesPerSector :                512
PhysicalBytesPerSectorForAtomicity :    4096
PhysicalBytesPerSectorForPerformance :  4096
FileSystemEffectivePhysicalBytesPerSectorForAtomicity : 4096
Device Alignment :                    Aligned (0x000)
Partition alignment on device :       Aligned (0x000)
Performs Normal Seeks
Trim Not Supported
Not DAX capable

C:\Windows\system32>
```

Having a larger physical sector size is what allows drive capacities to jump well beyond previous physical capacity limits. When there is only a 512-byte write, drives must perform additional work to complete the sector write. For best performance, applications must be updated to read and write data properly in this new level of granularity (4096 bytes).

Windows Server 2016 supports many drive interface architectures, including

- [Small Computer System Interface \(SCSI\)](#)
- [Parallel ATA \(PATA\)](#), also known as IDE
- [Serial ATA \(SATA\)](#)



The terms SCSI, IDE, and SATA designate the interface type used by the drives. The interface is used to communicate with a drive controller. SCSI drives use SCSI controllers, IDE drives use IDE controllers, and so on.

SCSI is one of the most commonly used interfaces, and there are multiple bus designs for SCSI and multiple interface types. Parallel SCSI (also called SPI) has largely been replaced by Serial Attached SCSI (SAS). Internet SCSI (iSCSI) uses the SCSI architectural model, but it uses TCP/IP as the transport rather than the traditional physical implementation.

SATA was designed to replace IDE. SATA drives are increasingly popular as a lowcost alternative to SCSI. SATA II and SATA III, the most common SATA interfaces, are designed to operate at 3 gigabits per second and 6 gigabits per second, respectively. In addition, eSATA (also known as external SATA) is meant for externally connected drives.

Various improvements to SATA are available, including revision 3.3, which includes support for SATA Express. SATA Express combines two PCI Express 3.0 ports to achieve speeds of up to 12 gigabits per second and is designed to be used with high performance solid state drives. With Serial Attached Storage (SAS), which leverages SCSI commands, enterprise hard disk drives are also achieving speeds of up to 12 gigabits per second.

When setting up a new server, you should give considerable thought to the drive configuration. Start by choosing drives or storage systems that provide the appropriate level of performance. There really is a substantial difference in speed and performance among various drive specifications.

You should consider not only the capacity of the drive but also the following:

- **Rotational speed**    A measurement of how fast the disk spins
- **Average seek time**    A measurement of how long it takes to seek between disk tracks during sequential I/O operations

Generally speaking, when comparing drives that conform to the same specification, such as Ultra640 SCSI or SATA III, the higher the rotational speed (measured in thousands of rotations per minute) and the lower the average seek time (measured in milliseconds, or msec), the better. As an example, a drive with a rotational speed of 15,000 RPM gives you 45–50 percent more I/O per second than the average 10,000 RPM drive, all other things being equal. A drive with a seek time of 3.5 msec gives you a 25–30 percent response time improvement over a drive with a seek time of 4.7 msec.







Other factors to consider include the following:

- **Maximum sustained data transfer rate** A measurement of how much data the drive can continuously transfer
- **Mean time to failure (MTTF)** A measurement of how many hours of operation you can expect to get from the drive before it fails
- **Nonoperational temperatures** Measurements of the temperatures at which the drive fails

Most drives of comparable quality have similar transfer rates and MTTF. For example, if you compare enterprise SAS drives with 15,000 RPM rotational speed from different vendors, you will probably find similar transfer rates and MTTF. Transfer rates can be expressed in megabytes per second (MBps) or gigabits per second (Gbps). A rate of 1.5 Gbps is equivalent to a data rate of 187.5 MBps, and 3.0 Gbps is equivalent to 375 MBps. Sometimes you'll get a maximum external transfer rate (per the specification to which the drive complies) and an average sustained transfer rate. The average sustained transfer rate is the most important factor.

**NOTE** Don't confuse MBps and Mbps. MBps is megabytes per second. Mbps is megabits per second. Because there are 8 bits in a byte, a 100 MBps transfer rate is equivalent to an 800 Mbps transfer rate.

Power requirements and temperature are other important factors to consider when you're selecting a drive, but they're factors few administrators take into account. Typically, the faster a drive rotates, the more power it uses and the hotter it runs. This is not always the case, but it is certainly something you should consider when making your choice. For example, 15K drives tend to use more power and run hot, and you must be sure sufficient power is available while also carefully regulating temperature. Be sure to compare average operating power and temperature specifications. While a typical 10K drive may use 7.87 watts of power, a typical 15K drive may use 8.68 watts of power. Typical 15K drives can become nonoperational at temperatures of 70 degrees Centigrade or higher (as would most other drives).

Windows Server 2016 adds support for disk drives with hardware encryption (referred to as encrypted drives). Encrypted drives have built-in processors that shift the encryption-decryption activities from the operating system to hardware, freeing up operating system resources. Windows Server 2016 will use hardware encryption with BitLocker when available. Other security features available in Windows Server 2016 include Secure Boot and Network Unlock. Secure Boot provides boot integrity by validating Boot Configuration Data (BCD) settings according to the Trusted Platform





Module (TPM) validation profile settings. Network Unlock can be used to automatically unlock the operating system drive on domain-joined computers.

## Preparing a Physical Drive for Use

After you install a drive, you need to configure it for use. You configure the drive by partitioning it and creating file systems in the partitions as needed. A partition is a section of a physical drive that functions as if it were a separate unit. After you create a partition, you can create a file system in the partition.

## The MBR and GPT Partition Styles

Two partition styles are used for disks: master boot record (MBR) and GUID partition table (GPT). The MBR contains a partition table that describes where the partitions are located on the disk. With this partition style, the first sector on a drive contains the master boot record and a binary code file called the master boot code that's used to boot the system. This sector is unpartitioned and hidden from view to protect the system.

With the MBR partitioning style, disks traditionally support volumes of up to 4 terabytes (TB) and use one of two types of partitions: primary or extended. Each MBR drive can have up to four primary partitions or three primary partitions and one extended partition. Primary partitions are drive sections you can access directly for file storage. You make a primary partition accessible to users by creating a file system on it. Although you can access primary partitions directly, you can't access extended partitions directly. Instead, you can configure extended partitions with one or more logical drives that are used to store files. Being able to divide extended partitions into logical drives allows you to divide a physical drive into more than four sections.

GPT was originally developed for high-performance computers using Intel's Itanium processors. The key difference between the GPT partition style and the MBR partition style has to do with how partition data is stored. With GPT, critical partition data is stored in the individual partitions, and redundant primary and backup partition tables are used for improved structural integrity. Additionally, GPT disks support volumes of up to 18 exabytes (1 exabyte equals 1024x1024 terabytes) and as many as 128 partitions. Although the GPT and MBR partitioning styles have underlying differences, most disk-related tasks are performed in the same way.

## Legacy and Protective Mbrs

Most computers ship with Unified Extensible Firmware Interface (UEFI). Although UEFI is replacing BIOS and EFI as the top-level firmware interface, UEFI doesn't





replace all the functionality in either BIOS or EFI and typically is wrapped around BIOS or EFI. With respect to UEFI, GPT is the preferred partitioning scheme and a protective MBR may be located on any disk that uses the GPT disk layout. A legacy MBR and a protective MBR differ in many important ways.

A legacy MBR is located at the first logical block on a disk that is not using the GPT disk layout. The first 512 bytes on an MBR disk have the following layout:

- The MBR begins with a 424-byte boot code, which is used to select an MBR partition record and load the first logical block of that partition. The boot code on the MBR is not executed by UEFI.
- The boot code is followed by a 4-byte unique MBR disk signature, which can be used by the operating system to identify the disk and distinguish the disk from other disks on the system. The unique signature is written by the operating system and not used by UEFI.
- A 2-byte separator follows the disk signature. At byte offset 446, there is an array of four MBR partition records, with each record being 16 bytes in length. Block 510 contains 0x55 and block 511 contains 0xAA. Block 512 is reserved.

The four partition records each define the first and last logical blocks that a particular partition uses on a disk:

- Each 16-byte MBR partition record begins with a 1-byte boot indicator. For example, a value of 0x80 identifies a bootable legacy partition. Any other value indicates that this is not a bootable legacy partition. This value is not used by UEFI. • The boot indicator is followed by a 3-byte address identifying the start of the partition. At byte offset 4, there's a 1-byte value that indicates the operating system type, which is followed by a 3-byte value that identifies the end of the partition. These values are not used by UEFI.
- At byte offset 8, there is a 4-byte value indicating the first logical block of the partition, and this is followed by a 4-byte value indicating size of the partition in units of logical blocks. Both of these values are used by UEFI.

**NOTE** If an MBR partition has an operating system type value of 0xEF, firmware must add the UEFI system partition GUID to the handle for the MBR partition. This allows boot applications, operating system loaders, drivers, and other lower-level tools to locate the UEFI system partition, which must physically reside on the disk.

A protective MBR may be located at the first logical block on a disk that is using the GPT disk layout. The protective MBR precedes the GUID Partition Table Header and





is used to maintain compatibility with tools that do not understand GPT partition structures. The purpose of the protective MBR is to protect the GPT partitions from boot applications, operating system loaders, drivers, and other lower-level tools that don't understand the GPT partitioning scheme. The protective MBR does this by defining a fake partition covering the entire disk. When a disk has a protective MBR, the first 512 bytes on the disk have the following layout:

- The protective MBR begins with a 424-byte boot code, which is not executed by UEFI.
- The boot code is followed by a 4-byte disk signature, which is set to zero and not used by UEFI.
- A 2-byte separator follows the disk signature. This separator is set to zero and not used by UEFI.
- At byte offset 446, there is an array of four MBR partition records, with each record being 16-bytes in length. Only the first partition record—the protective partition record—is used. The other partition records are set to zero.
- Block 510 contains 0x55 and block 511 contains 0xAA. Block 512 is reserved.

The protective partition record reserves the entire space on the disk after the first 512 bytes for the GPT disk layout. The protective partition record begins with a 1-byte boot indicator that is set to 0x00, which indicates a non-bootable partition. The boot indicator is followed by a 3-byte address identifying the start of the partition at 0x000200, which is the first usable block on the disk.

At byte offset 4, there's a 1-byte value set to 0xEE to indicate the operating system type as GPT Protective. This is followed by a 3-byte value that identifies the last usable block on the disk, which is the end of the partition (or 0xFFFFFFFF if it is not possible to represent this value).

At byte offset 8, there is a 4-byte value set to 0x00000001, which identifies the logical block address of the GPT partition header. This is followed by a 4-byte value indicating size of the disk minus one block (or 0xFFFFFFFF if the size of the disk is too large to be represented).

## Disk Types and File Systems

In addition to a partition style, physical drives have a disk type, which is either basic or dynamic, as discussed later in the chapter under “Working with Disks.” After you set the partition style and disk type for a physical drive, you can format free areas of the drive to establish logical partitions. Formatting creates a file system on a partition. Windows Server 2016 supports the following file systems:



- FAT
- FAT32
- exFAT
- NTFS
- ReFS

With FAT, the number of bits used with the file allocation table determines the variant with which you are working and the maximum volume size. FAT16, also known simply as FAT, defines its file allocation tables using 16 bits. Volumes that are 4 gigabytes (GB) or less in size are formatted with FAT16.

FAT32 defines its file allocation tables using 32 bits, and you can create FAT32 volumes that are 32 GB or less by using the Windows format tools. Although Windows can mount larger FAT32 volumes created with third-party tools, you should use NTFS for volumes larger than 32 GB.

Extended FAT is an enhanced version of FAT. Technically, exFAT could have been called FAT64 (and is called that by some). Because exFAT defines its file allocation tables by using 64 bits, it can overcome the 4-GB file-size limit and the 32-GB volumesize limit of FAT32 file systems. The exFAT format supports allocation unit sizes of up to 128 KB for volumes up to 256 TB.

NTFS volumes have a very different structure and feature set than FAT volumes. The first area of the volume is the boot sector, which stores information about the disk layout, and a bootstrap program executes at startup and boots the operating system. Instead of a file allocation table, NTFS uses a relational database called the master file table (MFT) to store information about files.

The MFT stores a file record of each file and folder on the volume, pertinent volume information, and details about the MFT itself. NTFS gives you many advanced options, including support for the Encrypting File System, compression, and the option to configure file screening and storage reporting. File screening and storage reporting are available when you add the File Server Resource Manager role service to a server as part of the File Services role.

At an administrator prompt, you can use Fsutil to get information about an NTFS volume by entering the following:

```
Fsutil fsinfo ntfsinfo DriveDesignator
```

where DriveDesignator is the designator of the drive to check, such as:

```
Fsutil fsinfo ntfsinfo c:
```

```
Administrator: Command Prompt
C:\Windows\system32>fsutil fsinfo ntfsinfo c:
NTFS Volume Serial Number : 0x9a4ab50e4ab4e7e1
NTFS Version : 3.1
LFS Version : 2.0
Number Sectors : 0x00000000218d17ff
Total Clusters : 0x00000000431a2fff
Free Clusters : 0x000000003e70120
Total Reserved : 0x000000000027b95
Bytes Per Sector : 512
Bytes Per Physical Sector : 4096
Bytes Per Cluster : 4096
Bytes Per FileRecord Segment : 1024
Clusters Per FileRecord Segment : 8
Mft Valid Data Length : 0x0000000070c0000
Mft Start Lcn : 0x000000000000c0000
Mft2 Start Lcn : 0x0000000000000002
Mft Zone Start : 0x00000000000c70c0
Mft Zone End : 0x00000000000cc820
Max Device Trim Extent Count : 0
Max Device Trim Byte Count : 0x0
Max Volume Trim Extent Count : 62
Max Volume Trim Byte Count : 0x40000000
Resource Manager Identifier : 8887E6FA-1161-11E6-9B2D-FE362480C635
C:\Windows\system32>
```

Resilient File System (ReFS) can be thought of as the next generation of NTFS. As such, ReFS remains compatible with core NTFS features while cutting noncore features to focus relentlessly on reliability. This means disk quotas, Encrypting File System (EFS), compression, file screening, and storage reporting are not available but built-in reliability features have been added.

One of the biggest reliability features in ReFS is a data integrity scanner, also called a data scrubber . The scrubber provides proactive error identification, isolation, and correction. If the scrubber detects data corruption, a repair process is used to localize the area of corruption and perform automatic online correction. Through an automatic online salvage process, corrupted areas that cannot be repaired, such as those caused by bad blocks on the physical disk, are removed from the live volume so that they cannot adversely affect good data. Because of the automated scrubber and salvage processes, a Check Disk feature is not needed when you use ReFS (and there's no Check Disk utility for ReFS).

**NOTE** When you are working with File And Storage Services, you can group available physical disks into storage pools so that you can create virtual disks from available capacity. Each virtual disk you create is a storage space. Because only NTFS and ReFS support storage spaces, you'll want to keep that in mind when you are formatting volumes on file servers. For more information about storage spaces, see Chapter 22 "Using Storage Spaces".

Like NTFS, ReFS has been revised several times since it was introduced. Windows Server 2016 RTM supports ReFS 3.0, which has different features and enhancements than earlier releases of the file system.




**TIP** Although you can use the Convert utility to convert a FAT volume to NTFS, there currently is no utility available to convert NTFS volumes to ReFS or vice versa. If you upgrade from an earlier release of Windows Server or move drives to a server running Windows Server 2016, you may find that volumes that use older versions of ReFS are unavailable. This is because the different versions of ReFS are incompatible with each other. To recover the data, you'll need to move the drive to a computer running a version of Windows Server that supports the file system.

At an administrator prompt, you can use the command-line utility Fsutil to get information about an ReFS volume by entering the following:

Fsutil fsinfo refsinfo DriveDesignator

where DriveDesignator is the designator of the drive to check, such as:

Fsutil fsinfo refsinfo d:



```
Administrator: Command Prompt
C:\Windows\system32>Fsutil fsinfo refsinfo g:
REFS Volume Serial Number : 0xa0162ca3162c7d02
REFS Version : 3.0
Number Sectors : 0x0000000018680000
Total Clusters : 0x00000000030d0000
Free Clusters : 0x0000000002fc67d1
Total Reserved : 0x00000000004a50c
Bytes Per Sector : 512
Bytes Per Physical Sector : 4096
Bytes Per Cluster : 4096
Checksum Type: CHECKSUM_TYPE_NONE


C:\Windows\system32>
```

## Using Disk Management

You use Disk Management to configure drives. Disk Management is included as part of the Computer Management console. You can also add it to custom MMCs.

In Computer Management, you can access Disk Management by expanding the Storage node, and then selecting Disk Management. Alternatively, you can type **diskmgmt.msc** in the Search box and then press Enter.

Disk Management has three views: Disk List, Graphical View, and Volume List. With remote systems, you're limited in the tasks you can perform with Disk Management. Remote management tasks you can perform include viewing drive details, changing drive letters and paths, and converting disk types. With removable media drives, you can also eject media remotely. To perform more advanced manipulation of remote drives, you can use the DiskPart command-line utility.



**NOTE** You should be aware that if you create a partition but don't format it, the partition is labeled as Free Space. In addition, if you haven't assigned a portion of the disk to a partition, this section of the disk is labeled Unallocated.

In the default configuration, the Volume List view is in the upper-right corner, and the Graphical View is in the lower-right corner. This is the default configuration. You can change the view for the top or bottom pane as follows:


- To change the top view, select **View**, choose **Top**, and then select the view you want to use.
- To change the bottom view, select **View**, choose **Bottom**, and then select the view you want to use.
- To hide the bottom view, select **View**, choose **Bottom**, and then select **Hidden**.

Windows Server 2016 supports four types of disk configurations:

- **Basic** The standard fixed disk type. Basic disks are divided into partitions and are the original disk type for early Windows operating systems.
- **Dynamic** An enhanced fixed disk type that you can update without having to restart the operating system (in most cases). Dynamic disks are divided into volumes.
- **Removable** The standard disk type associated with removable storage devices.
- **Virtual** The virtual hard disk (VHD) disk type associated with virtualization. Computers can use VHDs just like they use regular fixed disks and can even be configured to boot from a VHD.



These disk configurations can be used with legacy storage approaches as well as standards-based storage. From the Disk Management window, you can get more detailed information on a drive section by right-clicking it, and then selecting Properties. When you do this, you get a dialog box. The example that follows shows the dialog boxes for two fixed disks. The one on the left uses NTFS, and the one on the right uses ReFS. Both disks have additional tabs based on the server configuration.



If you've configured remote management through Server Manager and MMCs, you can use Disk Management to configure and work with disks on remote computers. Keep in mind, however, that your options are slightly different from when you are working with the disks on a local computer. Tasks you can perform include the following:

- Viewing limited disk properties, but not volume properties. When you are viewing disk properties, you'll see only the General and Volumes tabs. You won't be able to see volume properties.
- Changing drive letters and mount paths.
- Formatting, shrinking, and extending volumes. With mirrored, spanned, and striped volumes, you are able to add and configure related options.
- Deleting volumes (except for system and boot volumes).
- Creating, attaching, and detaching VHDs. When you create and attach VHDs, you need to enter the full file path and won't be able to browse for the .vhd file.



Some tasks you perform with disks and volumes depend on the Plug and Play and Remote Registry services.





## Using Removable Storage Devices

Removable storage devices can be formatted with ReFS, NTFS, FAT, FAT32, or exFAT. You connect external storage devices to a computer rather than installing them inside the computer. This makes external storage devices easier and faster to install than most fixed disk drives. Most external storage devices have either a USB, FireWire or eSATA interface.

Working with removable disks is similar to working with fixed disks. You can do the following:

- Right-click a removable disk, and then select **Open** or **Explore** to examine the disk's contents in File Explorer.
- Right-click a removable disk, and then select **Format** to format a removable disk as discussed in “Formatting Partitions” later in this chapter. Removable disks generally are formatted with a single partition.
- Right-click a removable disk, and then select **Properties** to view or set properties. On the **General** tab of the Properties dialog box, you can set the volume label.

When you work with removable disks, you can customize disk and folder views. To do this, right-click the disk or folder, select **Properties**, and then click the **Customize** tab. You can then specify the default folder type to control the default details displayed. For example, you can set the default folder type as **Documents** or **Pictures And Videos**. You can also set folder pictures and folder icons.

With exFAT, FAT, or FAT32, folders and files stored on a removable disk do not have any security permissions or features other than the basic read-only or hidden attribute flags that you can set. With NTFS, removable disks support network file and folder sharing. You configure sharing on removable disks in the same way you configure standard file sharing. You can assign share permissions, configure caching options for offline file use, and limit the number of simultaneous users. You can share an entire removable disk as well as individual folders stored on the removable disk. You can also create multiple share instances.

## Installing and Checking for a New Drive

Hot swapping is a feature that allows you to remove internal devices without shutting off the computer. Typically, hot-swappable internal drives are installed and removed from the front of the computer. If your computer supports hot swapping of internal drives, you can install drives without having to shut down. After you have installed a new drive, open **Disk Management**, and then choose **Rescan Disks** from the **Action**





menu. New disks that are found are added with the appropriate disk type. If a disk that you've added isn't found, restart the computer.

If the computer doesn't support hot swapping of internal drives, you must turn the computer off and then install the new drives. Then you can scan for new disks as described previously. If you are working with new disks that have not been initialized—meaning they don't have disk signatures—Disk Management will start the Initialize Disk dialog box as soon it starts up and detects the new disks. You can then initialize the disks by following these steps:

1. Each disk you install needs to be initialized. In the Initialize Disk dialog box, select the disk or disks you installed.
2. Disks can use either the MBR or GPT partition style. Select the partition style you want to use for the disk or disks you are initializing.
3. Click OK. If you elected to initialize disks, Windows writes a disk signature to the disks and initializes the disks with the basic disk type.

If you don't want to use the Initialize Disk dialog box, you can close it and use Disk Management instead to view and work with the new disk. In the Disk List view, the new disk is marked with a red downward-pointing arrow icon, the disk's type is listed as Unknown, and the disk's status is listed as Not Initialized. Right-click the disk's icon and select Online. Right-click the disk's icon again, and select Initialize Disk. You can then initialize the disk as discussed previously.

**NOTE** At an administrator PowerShell prompt, you can use `Get-Disk` to list available disks and `Initialize-Disk` to initialize new disks.

## Understanding Drive Status

Knowing the status of a drive is useful when you install new drives or troubleshoot drive problems. Disk Management shows the drive status in Graphical View and Volume List view. Here are the most common status values and the techniques you can use to resolve error conditions:

- **Online** The normal disk status. It means the disk is accessible and doesn't have problems. Both dynamic disks and basic disks display this status. The drive doesn't have any known problems. You don't need to take any corrective action.
- **Online (Errors)** I/O errors have been detected on a dynamic disk. You can try to correct temporary errors by right-clicking the disk and selecting `Reactivate Disk`. If this doesn't work, the disk might have physical damage or you might need to run a thorough check of the disk.



- **Offline** The disk isn't accessible and might be corrupted or temporarily unavailable. If the disk name changes to Missing, the disk can no longer be located or identified on the system. Check for problems with the drive, its controller, and cables. Make sure that the drive has power and is connected properly. Use the Reactivate Disk command to bring the disk back online (if possible).
- **Foreign** The disk has been moved to your computer but hasn't been imported for use. A failed drive brought back online might sometimes be listed as Foreign. Rightclick the disk, and then click Import Foreign Disks to add the disk to the system.
- **Unreadable** The disk isn't accessible currently, which can occur when disks are being rescanned. Both dynamic and basic disks display this status. With FireWire and USB card readers, you might get this status if the card is unformatted or improperly formatted. You might also get this status after the card is removed from the reader. Otherwise, if the drives aren't being scanned, the drive might be corrupted or have I/O errors. Right-click the disk, and then click Rescan Disk (on the Action menu) to try to correct the problem. You might also want to restart the system.
- **Unrecognized** The disk is of an unknown type and can't be used on the system. A drive from a system that is not based on Windows might display this status. If the disk is from another operating system, don't do anything. Normally, you can't use the drive on the computer without initializing and formatting it, so try a different drive.
- **Not Initialized** The disk doesn't have a valid signature. A drive from a system not based on Windows might display this status. If the disk is from another operating system, don't do anything. You can't use the drive on the computer, so try a different drive. To prepare the disk for use on Windows Server 2016, right-click the disk, and then click Initialize Disk.
- **No Media** No media has been inserted into the DVD or removable drive, or the media has been removed. Only DVD and removable disk types display this status. Insert a DVD or a removable disk to bring the disk online. With FireWire and USB card readers, this status is usually (but not always) displayed when the card is removed.





# Working with Disks

Windows Server 2016 supports both legacy storage approaches and standards-based storage approaches. Legacy storage uses software-based RAID. Standards-based storage uses Storage Spaces.





# Disk Management Essentials

Basic, dynamic, and virtual disk configurations can be used with both legacy storage approaches and standards-based storage. Normally, Windows Server 2016 disk partitions are initialized as basic disks. The exception is when you want to use software-based RAID instead of standards-based storage.

With software-based RAID, you can't create new fault-tolerant drive sets by using the basic disk type. You need to convert to dynamic disks and then create volumes that use striping, mirroring, or striping with parity (referred to as RAID 0, 1, and 5, respectively). The fault-tolerant features and the ability to modify disks without having to restart the computer are the key capabilities that distinguish dynamic disks from basic disks. Other features available on a disk depend on the disk formatting.

You can use both basic and dynamic disks on the same computer; however, volume sets must use the same disk type and partitioning style. For example, if you want to mirror drives C and D using software RAID, both drives must have the dynamic disk type and use the same partitioning style, which can be either MBR or GPT. Note that Disk Management allows you to start many disk configuration tasks regardless of whether the disks with which you are working use the dynamic disk type. The catch is that during the configuration process, Disk Management will convert the disks to the dynamic disk type. To learn how to convert a disk from basic to dynamic, see “Changing Drive Types” on the next page.

You can perform different disk configuration tasks with basic and dynamic disks. With basic disks, you can do the following:

- [Format partitions, and mark them as active](#)
- [Create and delete primary and extended partitions](#)
- [Create and delete logical drives within extended partitions](#)
- [Convert from a basic disk to a dynamic disk](#)

With dynamic disks, you can do the following:

- [Create and delete simple, striped, spanned, mirrored, and RAID-5 volumes](#)
- [Remove a mirror from a mirrored volume](#)
- [Extend simple or spanned volumes](#)
- [Split a volume into two volumes](#)
- [Repair mirrored or RAID-5 volumes](#)
- [Reactivate a missing or offline disk](#)
- [Revert to a basic disk from a dynamic disk \(requires deleting volumes and restoring from backup\)](#)



With either disk type, you can do the following:

- View properties of disks, partitions, and volumes
- Make drive-letter assignments
- Configure security and drive sharing
- Use Storage Spaces to implement standards-based storage







# Special Considerations for Basic and Dynamic Disks

Whether you're working with basic or dynamic disks, you need to keep in mind five special types of drive sections:

- **Active** The active partition or volume is the drive section for system caching and startup. Some devices with removable storage might be listed as having an active partition.
- **Boot** The boot partition or volume contains the operating system and its support files. The system and boot partition or volume can be the same.
- **Crash dump** The partition to which the computer attempts to write dump files in the event of a system crash. By default, dump files are written to the %SystemRoot% folder, but they can be located on any partition or volume.
- **Page file** A partition containing a paging file used by the operating system. Because a computer can page memory to multiple disks, according to the way virtual memory is configured, a computer can have multiple page file partitions or volumes.
- **System** The system partition or volume contains the hardware-specific files needed to load the operating system. The system partition or volume can't be part of a striped or spanned volume.

| Volume          | Layout | Type  | File System | Status                                                                   | Capacity  | Free Space | % Free |
|-----------------|--------|-------|-------------|--------------------------------------------------------------------------|-----------|------------|--------|
| ---             | Simple | Basic |             | Healthy (Recovery Partition)                                             | 450 MB    | 450 MB     | 100 %  |
| (C:)            | Simple | Basic | NTFS        | Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition) | 160.22 GB | 118.73 GB  | 74 %   |
| New Volume (D:) | Simple | Basic | ReFS        | Healthy (Primary Partition)                                              | 72.19 GB  | 71.77 GB   | 99 %   |

**REAL WORLD** GPT is becoming more common. With Windows Server 2016, a disk with the GPT partition style will have a recovery partition and an EFI system partition. The only Disk Management view that shows the partition style is the Disk list view. While you are working with the Graphical view, you can determine the partition style by right-clicking a disk (not a volume on the disk), and selecting Properties. In the Properties dialog box, the partition style is listed on the Volumes tab.

**NOTE** You can mark a partition as active by using Disk Management. In Disk Management, right-click the primary partition you want to mark as active, and then click Mark Partition As Active. You can't mark dynamic disk volumes as active. When you convert a basic disk containing the active partition to a dynamic disk, this partition becomes a simple volume that's active automatically.





## Changing Drive Types

You can use dynamic disks with any current version of Windows and many other operating systems, including most UNIX variants. However, keep in mind that you need to create a separate volume for any operating system not based on Windows.

You can't use dynamic disks on portable computers. When you are working with nonportable computers and servers, you only can use dynamic disks with drives connected to internal controllers (as well as some eSATA controllers). Although you can't use dynamic disks with portable or removable drives on these computers, you can connect such a drive to an internal controller or a recognized eSATA controller, and then use Disk Management to import the drive.

Windows Server 2016 provides the tools you need to convert a basic disk to a dynamic disk and to change a dynamic disk back to a basic disk. When you convert to a dynamic disk, partitions are changed to volumes of the appropriate type automatically. You can't change these volumes back to partitions. Instead, you must delete the volumes on the dynamic disk, and then change the disk back to a basic disk. Deleting the volumes destroys all the information on the disk.

## Converting a Basic Disk to a Dynamic Disk

Before you convert a basic disk to a dynamic disk, you should make sure that you don't need to boot the computer to an operating system that doesn't support dynamic disks. With MBR disks, you should also make sure that the disk has 1 MB of free space at the end of the disk. Although Disk Management reserves this free space when creating partitions and volumes, disk management tools on other operating systems might not. Without the free space at the end of the disk, the conversion will fail.

With GPT disks, you must have contiguous, recognized data partitions. If the GPT disk contains partitions that Windows doesn't recognize, such as those created by another operating system, you can't convert to a dynamic disk.

With either type of disk, the following holds true:

- There must be at least 1 MB of free space at the end of the disk. Disk Management reserves this free space automatically, but other disk management tools might not.
- You can't use dynamic disks on portable computers or with removable media. You can configure these drives only as basic drives with primary partitions.
- You shouldn't convert a disk if it contains multiple installations of the Windows operating system. If you do, you might be able to start the computer only by using the installation which did the conversion.




To convert a basic disk to a dynamic disk, follow these steps:


1. In Disk Management, right-click a basic disk that you want to convert, either in the Disk List view or in the left pane of the Graphical View. Then click Convert To Dynamic Disk.

**NOTE** Make sure you right-click a disk and not a volume on a disk. If you rightclick a volume, you won't have the same options.

2. In the Convert To Dynamic Disk dialog box, select the check boxes for the disks you want to convert.



3. Click OK to continue. This displays the Disks To Convert dialog box, which shows the disks you're converting.




The buttons and columns in this dialog box contain the following information:

- **Name** Shows the disk number.
- **Disk Contents** Shows the type and status of partitions, such as boot, active, or in use.



- **Will Convert** Specifies whether the drive will be converted. If the drive doesn't meet the criteria, it won't be converted, and you might need to take corrective action, as described previously.
  - **Details** Shows the volumes on the selected drive.
  - **Convert** Starts the conversion.
4. To begin the conversion, click Convert. Disk Management warns you that after the conversion is complete, you won't be able to start previous versions of Windows from volumes on the selected disks. Click Yes to continue.



5. Disk Management restarts the computer if a selected drive contains the boot partition, system partition, or a partition in use.

## Changing a Dynamic Disk Back to a Basic Disk

Before you can change a dynamic disk back to a basic disk, you must delete all dynamic volumes on the disk. After you do this, right-click the disk and select Convert To Basic Disk to change the dynamic disk to a basic disk. You can then create new partitions and logical drives on the disk.

## Reactivating Dynamic Disks

If the status of a dynamic disk is Online (Errors) or Offline, you can often reactivate the disk to correct the problem. You reactivate a disk by following these steps:

1. In Disk Management, right-click the dynamic disk you want to reactivate, and then click Reactivate Disk. Confirm the action when prompted.
2. If the drive status doesn't change, you might need to reboot the computer. If this still doesn't resolve the problem, check for problems with the drive, its controller, and the cables. Also make sure that the drive has power and is connected properly.



## Rescanning Disks

Rescanning all drives on a system updates the drive configuration information on the computer. Rescanning can sometimes resolve a problem with drives that show a status of Unreadable. You rescan disks on a computer by choosing Rescan Disks from the Action menu in Disk Management.





## Moving a Dynamic Disk to a New System

An important advantage of dynamic disks over basic disks is that you can easily move dynamic disks from one computer to another. For example, if after setting up a computer you decide that you don't really need an additional drive, you can move it to another computer where it can be better used.

Windows Server 2016 greatly simplifies the task of moving drives to a new system. Before moving disks, you should follow these steps:

1. Open Disk Management on the system where the dynamic disks are currently installed. Check the status of the disks, and ensure that they're marked as Healthy. If the status isn't Healthy, you should repair partitions and volumes before you move the disks.

**NOTE** Drives with BitLocker Drive Encryption cannot be moved by using this technique. BitLocker Drive Encryption wraps drives in a protected seal so that any offline tampering is detected and results in the disk being unavailable until an administrator unlocks it.

2. Check the drive subsystems on the original computer and the computer to which you want to transfer the disk. Both computers should have identical drive subsystems. If they don't, the Plug and Play ID on the system drive from the original computer won't match what the destination computer is expecting. As a result, the destination computer won't be able to load the right drivers, and the boot attempt might fail.
3. Check whether any dynamic disks you want to move are part of a spanned, extended, or striped set. If they are, you should make a note of which disks are part of which set and plan on moving all disks in a set together. If you are moving only part of a disk set, you should be aware of the consequences. For spanned, extended, or striped volumes, moving only part of the set will make the related volumes unusable on the current computer and on the computer to which you are planning to move the disks.

When you are ready to move the disks, follow these steps:

1. On the original computer, start Computer Management. Then, in the left pane, select Device Manager. In the Device list, expand Disk Drives. This shows a list of the physical disk drives on the computer. Right-click each disk you want to move, and then click Uninstall. If you are unsure which disks to uninstall, rightclick each disk and click Properties. In the Properties dialog box, click the Volumes tab and then select Populate to show the volumes on the selected disk.





2. Next, on the original computer, select the Disk Management node in Computer Management. If the disk or disks you want to move are still listed, right-click each disk, and then click Remove Disk.
3. After you perform these procedures, you can move the dynamic disks. If the disks are hot-swappable disks and this feature is supported on both computers, remove the disks from the original computer, and then install them on the destination computer. Otherwise, turn off both computers, remove the drives from the original computer, and then install them on the destination computer. When you have finished, restart the computers.
4. On the destination computer, access Disk Management, and then choose Rescan Disks from the Action menu. When Disk Management finishes scanning the disks, right-click any disk marked Foreign, and then click Import. You should now be able to access the disks and their volumes on the destination computer.

**NOTE** In most cases, the volumes on the dynamic disks should retain the drive letters they had on the original computer. However, if a drive letter is already used on the destination computer, a volume receives the next available drive letter. If a dynamic volume previously did not have a drive letter, it does not receive a drive letter when moved to the destination computer. Additionally, if automounting is disabled, the volumes aren't automatically mounted, and you must manually mount volumes and assign drive letters.





# Managing Virtual Hard Disks

Virtual hard disks (VHDs) provide alternate storage locations that function similar to traditional disks while requiring some additional setup prior to use. Before you can use a VHD, you must create, attach and initialize it using Disk Management or another management tool. You must also partition the virtual disk and then format this partition.





# Creating Virtual Hard Disks

In Disk Management, you can create a VHD by choosing Create VHD from the Action menu. In the Create And Attach Virtual Hard Disk dialog box, specify whether you want to use the standard VHD format or the newer VHDX format. As the VHDX format has many enhancements over the VHD format, including resiliency for power failure events and support for very large virtual disks, you'll typically want to use this format. The only time you wouldn't want to use this format is when you may move the virtual disk to an early Windows or Windows Server operating system.

Create and Attach Virtual Hard Disk

Specify the virtual hard disk location on the machine.

Location:

G:\data-storage4.vhdx Browse...

Virtual hard disk size: 500 GB

Virtual hard disk format

☐ VHD  
Supports virtual disks up to 2040 GB in size.

☒ VHDX  
Supports virtual disks larger than 2040 GB in size (Supported maximum of 64 TB) and is resilient to power failure events. This format is not supported in operating systems earlier than Windows 8 or Windows Server 2012.


Virtual hard disk type

☐ Fixed size  
The virtual hard disk file is allocated to its maximum size when the virtual hard disk is created.


☒ Dynamically expanding (Recommended)  
The virtual hard disk file grows to its maximum size as data is written to the virtual hard disk.

OK Cancel



Next, click Browse. Use the Browse Virtual Disk Files dialog box to select the location where you want to create the .vhd or .vhdx file for the VHD. After you enter a file name for the virtual disk, such as data-storage, click in the Save As Type list and select Virtual Disk Files (\*.vhd) or Virtual Disk Files (\*.vhdx) as appropriate for the virtual disk format you are using. Click Save to close the Browse... dialog box.



After you specify the storage location, size the virtual disk using the Virtual Hard Disk Size options. You can enter the size of the disk in megabytes, gigabytes, or terabytes. Complete the configuration by specifying whether the size of the VHD dynamically expands to its fixed maximum size as data is saved to it or instead uses a fixed amount of space regardless of the amount of data stored on it. When you click OK, Disk Management creates the VHD.



The VHD is attached automatically and added as a new disk. To initialize the disk for use, right-click the disk entry in Graphical View, and then click Initialize Disk. In the Initialize Disk dialog box, the disk is selected for initialization. Specify the disk type as MBR or GPT, and then click OK.



After initializing the disk, you can create a volume on the disk using the unpartitioned space. Right-click the unpartitioned space on the disk and select New Simple Volume. Use the New Simple Volume wizard to create the volume, as discussed in the section of this chapter entitled “Creating Volumes and Volume Sets.” After you create the volume, the VHD is available for use.

## Working with Virtual Hard Disks

After you've created, attached, initialized, and formatted a VHD, you can work with a virtual disk in much the same way as you work with other disks. You can write data to and read data from a VHD. You can boot the computer from a VHD.

You are able to take a VHD offline or put a VHD online by right-clicking the disk entry in Graphical View and selecting Offline or Online, respectively. If you no longer want to use a VHD, you can detach it by right-clicking the disk entry in Graphical View, selecting Detach VHD, and then clicking OK in the Detach Virtual Hard Disk dialog box.



You can use VHDs created with other programs as well. If you created a VHD using another program or have a detached VHD you want to attach, you can work with the VHD by completing the following steps:

1. In Disk Management, click the Attach VHD option on the Action menu.
2. In the Attach Virtual Hard Disk dialog box, click Browse. Use the Browse Virtual Disk Files dialog box to select the .vhd or .vhdx file for the virtual disk, and then click Open.
3. If you want to attach the VHD in read-only mode, select Read-Only. Click OK to attach the VHD.



## Chapter 20. Partitioning and Optimizing Drives

Whether you are working with hard disk drives (HDDs) or solid state drives (SSDs), you can use disks with either the basic or dynamic disk type. Basic disks use



partitions, while dynamic disks use volumes, and both of which can be compressed or encrypted. Compression allows you to store more information on a drive, but shouldn't be used with some types of data, such as professionally-produced audio. Encryption protects sensitive data by making it unreadable if files are copied to another computer without first being decrypted. In addition to standard encryption, Windows Server 2016 also supports BitLocker Drive Encryption, a drive encryption technology discussed in Chapter 21.





# Using Basic Disks and Partitions

When you install a new computer or update an existing computer, you often need to partition the drives on the computer. You partition drives by using Disk Management.



# Partitioning Essentials

In Windows Server 2016, a physical drive using the MBR partition style can have up to four primary partitions and one extended partition. This allows you to configure MBR drives in one of two ways: by using one to four primary partitions, or by using one to three primary partitions and one extended partition. A primary partition can fill an entire disk, or you can size it as appropriate for the workstation or server you're configuring. Within an extended partition, you can create one or more logical drives. A logical drive is simply a section of a partition with its own file system. Generally, you use logical drives to divide a large drive into manageable sections. With this in mind, you might want to divide a 600-GB extended partition into three logical drives of 200 GB each. Physical disks with the GPT partition style can have up to 128 partitions.

After you partition a drive, you format the partitions. This is high-level formatting that creates the file system structure rather than low-level formatting that sets up the drive for initial use. You're probably very familiar with the C drive used by Windows Server 2016. Well, the C drive is simply the designator for a disk partition. If you partition a disk into multiple sections, each section can have its own drive letter. You use the drive letters to access file systems in various partitions on a physical drive. Unlike MS-DOS, which assigns drive letters automatically starting with the letter C, Windows Server 2016 lets you specify drive letters. Generally, the drive letters C through Z are available for your use.

**NOTE** The drive letter A used to be assigned to a system's floppy disk drive. If the system had a second floppy disk drive, the letter B was assigned to it, so you could use only the letters C through Z. Don't forget that DVD drives and other types of media drives need drive letters as well. The total number of drive letters you can use at one time is 24. If you need additional volumes, you can create them by using drive paths.

By using drive letters, you can have only 24 active volumes. To get around this limitation, you can mount disks to drive paths. A drive path is set as a folder location on another drive. For example, you might mount additional drives as E:\Data1, E:\Data2, and E:\Data3. You can use drive paths with basic and dynamic disks. The only restriction for drive paths is that you mount them on empty folders that are on NTFS drives.

To help you differentiate between primary partitions and extended partitions with logical drives, Disk Management color codes the partitions. For example, primary partitions might be color coded with a dark-blue band and logical drives in extended



partitions might be color coded with a light-blue band. The key for the color scheme is shown at the bottom of the Disk Management window. You can change the colors in the Settings dialog box by choosing Settings from the View menu.






# Creating Partitions and Simple Volumes

Windows Server 2016 simplifies the Disk Management user interface by using one set of dialog boxes and wizards for both partitions and volumes. The first three volumes on a basic drive are created automatically as primary partitions. If you try to create a fourth volume on a basic drive, the remaining free space on the drive is converted automatically to an extended partition with a logical drive of the size you designate by using the new volume feature in the extended partition. Any subsequent volumes are created in the extended partitions as logical drives automatically.

In Disk Management, you create partitions, logical drives, and simple volumes by following these steps:

1. In Disk Management's Graphical View, right-click an unallocated or free area, and then click New Simple Volume. This starts the New Simple Volume Wizard. Read the Welcome page, and then click Next.




The screenshot shows the 'New Simple Volume Wizard' window, specifically the 'Specify Volume Size' step. The window title is 'New Simple Volume Wizard' with a close button (X) in the top right corner. Below the title bar, the text 'Specify Volume Size' is followed by the instruction 'Choose a volume size that is between the maximum and minimum sizes.' The main area contains three rows of information: 'Maximum disk space in MB:' with the value '199999', 'Minimum disk space in MB:' with the value '8', and 'Simple volume size in MB:' with a spin box currently set to '199999'. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

2. The Specify Volume Size page in the New Simple Volume Wizard specifies the minimum and maximum size for the volume in megabytes and lets you size the volume within these limits. Size the partition in megabytes in the Simple Volume Size In MB box, and then click Next.
3. On the Assign Drive Letter Or Path page, specify whether you want to assign a drive letter or path, and then click Next. The following options are available:

- **Assign The Following Drive Letter** Choose this option to assign a drive letter. Then select an available drive letter in the list provided. By default, Windows Server 2016 selects the lowest available drive letter and excludes reserved drive letters as well as those assigned to local disks or network drives.
- **Mount In The Following Empty NTFS Folder** Choose this option to mount the partition in an empty NTFS folder. You must then type the path to an existing folder or click Browse to search for or create a folder to use.
- **Do Not Assign A Drive Letter Or Drive Path** Choose this option if you want to create the partition without assigning a drive letter or path. If you later want the partition to be available for storage, you can assign a drive letter or path at that time.

NOTE You don't have to assign volumes a drive letter or a path. A volume with no designators is considered to be unmounted and is for the most part unusable. An unmounted volume can be mounted by assigning a drive letter or a path at a later date. See "Assigning Drive Letters and Paths" in Chapter 24.




4. On the Format Partition page, determine whether and how the volume should be formatted. If you want to format the volume, select Format This Volume With The Following Settings, and then configure the following options:

- **File System** Sets the file system type as FAT, FAT32, exFAT, NTFS, or ReFS. The file system types available depend on the size of the volume you are

formatting. If you use FAT32, you can later convert to NTFS with the Convert utility. You can't, however, convert NTFS partitions to FAT32.

- **Allocation Unit Size** Sets the cluster size for the file system. This is the basic unit in which disk space is allocated. The default allocation unit size is based on the size of the volume and is set dynamically prior to formatting by default. To override this feature, you can set the allocation unit size to a specific value. If you use many small files, you might want to use a smaller cluster size, such as 512 or 1024 bytes. With these settings, small files use less disk space. Note that ReFS volumes have a fixed allocation unit size.




- **Volume Label** Sets a text label for the partition. This label is the partition's volume name and is set to New Volume by default. You can change the volume label at any time by right-clicking the volume in File Explorer, clicking Properties, and typing a new value in the Label box provided on the General tab.
- **Perform A Quick Format** Tells Windows Server 2016 to format without checking the partition for errors. With large partitions, this option can save you a few minutes. However, it's usually better to check for errors, which enables Disk Management to mark bad sectors on the disk and lock them out.
- **Enable File And Folder Compression** Turns on compression for the disk. Built-in compression is available only for NTFS (and is not supported for FAT, FAT32, exFAT, or ReFS). Under NTFS, compression is transparent to users and

compressed files can be accessed just like regular files. If you select this option, files and directories on this drive are compressed automatically. For more information on compressing drives, files, and directories, see “Compressing Drives and Data” later in this chapter.

## Formatting Partitions

Formatting creates a file system on a partition and permanently deletes any existing data. This is high-level formatting that creates the file system structure rather than low-level formatting that initializes a drive for use. To format a partition, right-click the partition, and then click Format. This opens the Format dialog box.



You use the formatting options as follows:

- **Volume Label** Specifies a text label for the partition. This label is the partition’s volume name.
- **File System** Specifies the file system type as FAT, FAT32, exFAT, NTFS, or ReFS. The file system types available depend on the size of the volume you are formatting.
- **Allocation Unit Size** Specifies the cluster size for the file system. This is the basic unit in which disk space is allocated. The default allocation unit size is based on the size of the volume and is set dynamically prior to formatting. To override this feature, you can set the allocation unit size to a specific value. If you use lots of small files, you might want to use a smaller cluster size, such as 512 or 1024 bytes. With these settings, small files use less disk space.
- **Perform A Quick Format** Tells Windows Server 2016 to format without checking the partition for errors. With large partitions, this option can save you a few minutes. However, it’s more prudent to check for errors, which allows Disk Management to mark bad sectors on the disk and lock them out.

When you’re ready to proceed, click OK. Because formatting a partition destroys any existing data, Disk Management gives you one last chance to cancel the procedure.



Click OK to start formatting the partition. Disk Management changes the drive's status to reflect the formatting and the percentage of completion. When formatting is complete, the drive status changes to reflect this.

## Using Volumes and Volume Sets

You create volumesets and RAID arrays on dynamic drives. With a volume set, you can create a single volume that spans multiple drives. Users can access this volume as if it were a single drive, regardless of how many drives the volume is spread over. A volume that's on a single drive is referred to as a simple volume . A volume that spans multiple drives is referred to as a spanned volume .

With a RAID array, you can protect important business data and sometimes improve the performance of drives. RAID can be implemented by using the built-in features of the operating system (a software approach) or by using hardware. Windows Server 2016 supports three levels of software RAID: 0, 1, and 5. RAID arrays are implemented as mirrored, striped, and striped with parity volumes.

You create and manage volumes in much the same way in which you create and manage partitions. A volume is a drive section you can use to store data directly.

**NOTE** With spanned and striped volumes on basic disks, you can delete a volume but you can't create or extend volumes. With mirrored volumes on basic disks, you can delete, repair, and resync the mirror. You can also break the mirror. For striped with parity volumes (RAID-5) on basic disks, you can delete or repair the volume, but you can't create new volumes.


## Understanding Volume Basics

Disk Management color codes volumes by type, much like it does partitions. Volumes also have the following properties:

- **Layout** Volume layouts include simple, spanned, mirrored, striped, and striped with parity.
- **Type** Volumes always have the type dynamic . Partitions always have the type basic .
- **File System** Like partitions, each volume can have a different file system type, such as FAT or NTFS file system. Note that FAT16 is available only when the partition or volume is 2 GB or less in size.
- **Status** The state of the drive. In Graphical View, the state is shown as Healthy, Failed Redundancy, and so on. The next section, "Understanding volume sets," discusses volume sets and the various states you might see.



- 
- **Free Space** The total amount of available space on the volume.
- **% Free** The percentage of free space out of the total storage size of the volume.



The screenshot shows the 'Computer Management' window with the 'Disk Management' tab selected. It displays a list of volumes and a graphical representation of the disks below.

| Volume          | Layout | Type    | File System | Status                                                   | Capacity   | Free Space | % Free |
|-----------------|--------|---------|-------------|----------------------------------------------------------|------------|------------|--------|
| (C:)            | Simple | Basic   | NTFS        | Healthy (OEM Partition)                                  | 39 MB      | 39 MB      | 100 %  |
| History (E:)    | Simple | Basic   | NTFS        | Healthy (Boot, Page File, Crash Dump, Primary Partition) | 268.41 GB  | 249.13 GB  | 93 %   |
| New Volume (F:) | Simple | Dynamic | NTFS        | Healthy                                                  | 465.76 GB  | 41.76 GB   | 9 %    |
| OS              | Simple | Basic   | FAT32       | Healthy (System, Active, Primary Partition)              | 8191.87 GB | 8190.69 GB | 100 %  |

| Disk   | Type    | Capacity   | State  |
|--------|---------|------------|--------|
| Disk 0 | Basic   | 465.76 GB  | Online |
| Disk 2 | Dynamic | 465.76 GB  | Online |
| Disk 3 | Dynamic | 8191.88 GB | Online |

Legend: ■ Unallocated ■ Primary partition ■ Simple volume

**Capacity** The total storage size of the drive.

An important advantage of dynamic volumes over basic volumes is that dynamic volumes enable you to make changes to volumes and drives without having to restart the system (in most cases). Volumes also let you take advantage of the fault-tolerance enhancements of Windows Server 2016. You can install other operating systems and dual boot a Windows Server 2016 system by creating a separate volume for the other operating system. For example, you could install Windows Server 2016 on volume C and Windows 10 on volume D.

With volumes, you can do the following:

- Assign drive letters and drive paths as discussed in “Assigning Drive Letters and Paths” later in Chapter 24
- Create any number of volumes on a disk as long as you have free space
- Create volumes that span two or more disks and, if necessary, configure fault tolerance
- Extend volumes to increase the volumes’ capacity
- Designate active, system, and boot volumes as described in “Special considerations for basic and dynamic disks” in Chapter 19, “Data Storage: The Essentials”





# Understanding Volume Sets

With volume sets, you can create volumes that span several drives by using free space on different drives to create what users see as a single volume. Files are stored on the volume set segment by segment, with the first segment of free space being used to store files before other segments. When the first segment fills up, the second segment is used, and so on.

You can create a volume set using free space on up to 32 hard disk drives. The key advantage to volume sets is that they let you tap into unused free space and create a usable file system. The key disadvantage is that if any hard disk drive in the volume set fails, the volume set can no longer be used, which means that essentially all the data on the volume set is lost.

Understanding the volume status is useful when you install new volumes or are trying to troubleshoot problems. Disk Management shows the drive status in Graphical View and Volume List view. Status values for dynamic volumes and the resolution for error states follow:

- **Data Incomplete** Spanned volumes on a foreign disk are incomplete. You must have forgotten to add the other disks from the spanned volume set. Add the disks that contain the rest of the spanned volumes, and then import all the disks at one time.
- **Data Not Redundant** Fault-tolerant volumes on a foreign disk are incomplete. You must have forgotten to add the other disks from a mirror or RAID-5 set. Add the remaining disks, and then import all the disks at one time.
- **Failed** An error disk status. The disk is inaccessible or damaged. Ensure that the related dynamic disk is online. As necessary, right-click the volume, and then click Reactivate Volume. For a basic disk, you might need to check the disk for a faulty connection.
- **Failed Redundancy** An error disk status. One of the disks in a mirror or RAID-5 set is offline. Ensure that the related dynamic disk is online. If necessary, reactivate the volume. Next, you might need to replace a failed mirror or repair a failed RAID-5 volume.
- **Formatting** A temporary status that indicates the volume is being formatted. The progress of the formatting is indicated as the percent complete unless you choose the Perform A Quick Format option.
- **Healthy** The normal volume status. The volume doesn't have any known problems. You don't need to take any corrective action.





- **Healthy (At Risk)** Windows had problems reading from or writing to the physical disk on which the dynamic volume is located. This status appears when Windows encounters errors. Right-click the volume, and then click Reactivate Volume. If the disk continues to have this status or has this status periodically, the disk might be failing, and you should back up all data on the disk.
- **Healthy (Unknown Partition)** Windows does not recognize the partition. This can occur because the partition is from a different operating system or is a manufacturer-created partition used to store system files. No corrective action is necessary.
- **Initializing** A temporary status that indicates the disk is being initialized. The drive status should change after a few seconds.
- **Regenerating** A temporary status that indicates that data and parity for a RAID-5 volume are being regenerated. Progress is indicated as the percent complete. The volume should return to Healthy status.
- **Resynching** A temporary status that indicates that a mirror set is being resynchronized. Progress is indicated as the percent complete. The volume should return to Healthy status.
- **Stale Data** Data on foreign disks that are fault tolerant are out of sync. Rescan the disks or restart the computer, and then check the status. A new status should be displayed, such as Failed Redundancy.
- **Unknown** The volume cannot be accessed. It might have a corrupted boot sector. The volume might have a boot sector virus. Check it with an up-to-date antivirus program. Rescan the disks or restart the computer, and then check the status.





# Creating Volumes and Volume Sets

You can format simple volumes as exFAT, FAT, FAT32, or NTFS. To make management easier, you should format volumes that span multiple disks as NTFS, which enables you to expand the volume set if necessary. If you find you need more space on a volume, you can extend simple and spanned volumes by selecting an area of free space and adding it to the volume. You can extend a simple volume within the same disk, and you can also extend a simple volume onto other disks. When you do this, you create a spanned volume, which you must format as NTFS.

You create volumes and volume sets by following these steps:

1. In Disk Management's Graphical View, right-click an unallocated area, and then click New Spanned Volume or New Striped Volume as appropriate. Read the Welcome page, and then click Next.
2. Select the disks that you want to be part of the volume, and then size the volume segments on those disks.

The screenshot shows the 'New Spanned Volume' wizard in Windows Disk Management. The title bar says 'New Spanned Volume' with a close button. The main heading is 'Select Disks' with a subtitle 'You can select the disks and set the disk size for this volume.' Below this, it says 'Select the disk you want to use, and then click Add.' There are two lists: 'Available:' on the left and 'Selected:' on the right. The 'Available:' list is empty. The 'Selected:' list contains 'Disk 2 476937 MB' and 'Disk 3 1240000 MB', with 'Disk 3' highlighted. Between the lists are three buttons: 'Add >', '< Remove', and '< Remove All'. Below the lists, there are three input fields: 'Total volume size in megabytes (MB):' with the value '1716937', 'Maximum available space in MB:' with the value '8388478', and 'Select the amount of space in MB:' with a value of '1240000' and a spinner control. At the bottom are three buttons: '< Back', 'Next >' (which is highlighted with a blue border), and 'Cancel'.

3. Available disks are shown in the Available list. If necessary, select a disk in this list, and then click Add to add the disk to the Selected list. If you make a mistake, you can remove disks from the Selected list by selecting the disk, and then clicking Remove.



**CAUTION** The disk wizards in Windows Server 2016 show both basic and dynamic disks with available disk space. If you add space from a basic disk, the wizard converts the disk to a dynamic disk before creating the volume set. Before clicking Yes to continue, be sure you really want to do this because it can affect how the disk is used by the operating system.

4. Select a disk in the Selected list, and then specify the size of the volume on the disk in the Select The Amount Of Space In MB box. The Maximum Available Space In MB box shows you the largest area of free space available on the disk. The Total Volume Size In Megabytes box shows you the total disk space selected for use with the volume. Click Next.

**TIP** Although you can size a volume set any way you want, consider how you'll use volume sets on the system. Simple and spanned volumes aren't fault tolerant; rather than creating one monstrous volume with all the available free space, you might want to create several smaller volumes to help ensure that losing one volume doesn't mean losing all your data.

5. Specify whether you want to assign a drive letter or path to the volume, and then click Next. You use the available options as follows:
    - **Assign The Following Drive Letter** To assign a drive letter, choose this option, and then select an available drive letter in the list provided.
    - **Mount In The Following Empty NTFS Folder** To assign a drive path, choose this option, and then type the path to an existing folder on an NTFS drive, or click Browse to search for or create a folder.
    - **Do Not Assign A Drive Letter Or Drive Path** To create the volume without assigning a drive letter or path, choose this option. You can assign a drive letter or path later if necessary.
  6. Specify whether the volume should be formatted. If you elect to format the volume, set the following formatting options:
    - **File System** Specifies the file system type, such as NTFS or ReFS.
    - **Allocation Unit Size** Specifies the cluster size for the file system. This is the basic unit in which disk space is allocated. The default allocation unit size is based on the volume's size and is set dynamically prior to formatting. Although you can't change the default size if you select ReFS, you can set the allocation unit size to a specific value with other formats. If you use a lot of small files, you might want to use a smaller cluster size, such as 512 or 1024 bytes. With these settings, small files use less disk space.
- 



- **Volume Label** Specifies a text label for the partition. This label is the partition's volume name.
  - **Perform A Quick Format** Tells Windows to format without checking the partition for errors. With large partitions, this option can save you a few minutes. However, it's more prudent to check for errors, which allows Disk Management to mark bad sectors on the disk and lock them out.
  - **Enable File And Folder Compression** Turns on compression for the disk. Compression is transparent to users, and compressed files can be accessed just like regular files. If you select this option, files and directories on this drive are compressed automatically. For more information on compressing drives, files, and directories, see "Compressing Drives and Data" in Chapter 20. (NTFS only)
7. Click Next, and then click Finish.





## Deleting Volumes and Volume Sets

You use the same technique to delete all volumes, whether they're simple, spanned, mirrored, striped, or RAID-5 (striped with parity). Deleting a volume set removes the associated file system, and all associated data is lost. Before you delete a volume set, you should back up any files and directories the volume set contains.

You can't delete a volume that contains the system, boot, or active paging files for Windows Server 2016.

To delete volumes, follow these steps:

1. In Disk Management, right-click any volume in the set, and then click Delete Volume. You can't delete a portion of a spanned volume without deleting the entire volume.
2. Click Yes to confirm that you want to delete the volume.





# Managing Volumes

You manage volumes much like you manage partitions. Follow the techniques outlined in Chapter 24.





## Compressing Drives and Data

When you format a drive for NTFS, Windows Server 2016 allows you to turn on the built-in compression feature. With compression, all files and directories stored on a drive are automatically compressed when they're created. Because this compression is transparent to users, compressed data can be accessed just like regular data. The difference is that you can store more information on a compressed drive than you can on an uncompressed drive.

**IMPORTANT** File Explorer shows the names of compressed resources in blue. It's also important to point out that ReFS does not support NTFS compression.


**REAL WORLD** Although compression is certainly a useful feature when you want to save disk space, you can't encrypt compressed data. Compression and encryption are mutually exclusive alternatives for NTFS volumes, which means you have the choice of using compression or using encryption. You can't use both techniques. For more information on encryption, see "Encrypting Drives and Data" later in this chapter. If you try to compress encrypted data, Windows Server 2016 automatically decrypts the data and then compresses it. Likewise, if you try to encrypt compressed data, Windows Server 2016 uncompresses the data, and then encrypts it.



## Compressing Drives

To compress a drive and all its contents, follow these steps:

1. In File Explorer or Disk Management, right-click the drive you want to compress, and then click Properties.
2. On the General tab, select Compress Drive To Save Disk Space, and then click OK.
3. In the Confirm Attribute Changes dialog box, select whether to apply the changes to subfolders and files, and then click OK.



## Compressing Directories and Files

If you decide not to compress a drive, Windows Server 2016 lets you selectively compress directories and files. To compress a file or directory, follow these steps:

1. In File Explorer, right-click the file or directory you want to compress, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced. In the Advanced Attributes dialog box, select the Compress Contents To Save Disk Space check box. Click OK twice.



For an individual file, Windows Server marks the file as compressed, and then compresses it. For a directory, Windows Server marks the directory as compressed and then compresses all the files in it. If the directory contains subfolders, Windows Server displays a dialog box that allows you to compress all the subfolders associated with the directory. Simply select **Apply Changes To This Folder, Subfolders, And Files**, and then click **OK**. After you compress a directory, any new files added or copied to the directory are compressed automatically.


**NOTE** If you move an uncompressed file from a different drive, the file is compressed. However, if you move an uncompressed file to a compressed folder on the same NTFS drive, the file isn't compressed. Note also that you can't encrypt compressed files.

## Expanding Compressed Drives

File Explorer shows the names of compressed files and folders in blue. You can remove compression from a drive by following these steps:

1. In File Explorer or Disk Management, right-click the drive that contains the data you want to expand, and then click **Properties**.
2. Clear the **Compress Drive To Save Disk Space** check box, and then click **OK**.
3. In the **Confirm Attribute Changes** dialog box, select whether to apply the change to subfolders and files, and then click **OK**.





**TIP** Windows always checks the available disk space before expanding compressed data. You should, too. If less free space is available than used space, you might not be able to complete the expansion. For example, if a compressed drive uses 150 GB of space and has 70 GB of free space available, you won't have enough free space to expand the data. Generally, you need about 1.5 to 2 times as much free space as you have compressed data.

## Expanding Compressed Directories and Files

If you decide that you want to expand a compressed file or directory, follow these steps:

1. Right-click the file or directory in File Explorer, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced. Clear the Compress Contents To Save Disk Space check box. Click OK twice.

With files, Windows Server removes compression and expands the file. With directories, Windows Server expands all the files within the directory. If the directory contains subfolders, you also have the opportunity to remove compression from the subfolders. To do this, select Apply Changes To This Folder, Subfolders, And Files when prompted, and then click OK.

**TIP** Windows Server also provides command-line utilities for compressing and uncompressing data. The compression utility is called Compact (Compact.exe). The uncompression utility is called Expand (Expand.exe).



# Encrypting Drives and Data

NTFS has many advantages over other file systems you can use with Windows Server.

One advantage is the capability to automatically encrypt and decrypt data by using the Encrypting File System (EFS). When you encrypt data, you add an extra layer of protection to sensitive data, and this extra layer acts as a security blanket blocking all other users from reading the contents of the encrypted files. Indeed, one of the great benefits of encryption is that only the designated user can access the data. This benefit is also a disadvantage in that the user must remove encryption before authorized users can access the data.

**NOTE** As discussed previously, you can't compress encrypted files. The encryption and compression features of NTFS are mutually exclusive. You can use one feature or the other, but not both. Note also that ReFS doesn't support this type of encryption.





# Understanding Encryption and the Encrypting File System

File encryption is supported on a per-folder or per-file basis. Any file placed in a folder marked for encryption is automatically encrypted. Files in encrypted format can be read only by the person who encrypted the file. Before other users can read an encrypted file, the user must decrypt the file or grant special access to the file by adding a user's encryption key to the file.

Every encrypted file has the unique encryption key of the user who created the file or currently has ownership of the file. An encrypted file can be copied, moved, backed up, restored, or renamed just like any other file, and in most cases these actions don't affect the encryption of the data. (For details, see "Working With Encrypted Files and Folders" later in this chapter.) The user who encrypts a file always has access to the file, provided that the user's public-key certificate is available on the computer that she is using. For this user, the encryption and decryption process is handled automatically and is transparent.

EFS is the process that handles encryption and decryption. The default setup for EFS allows users to encrypt files without needing special permission. Files are encrypted by using a public/private key that EFS automatically generates on a per-user basis.

Encryption certificates are stored as part of the data in user profiles. If a user works with multiple computers and wants to use encryption, an administrator needs to configure a roaming profile for that user. A roaming profile ensures that the user's profile data and public-key certificates are accessible from other computers. Without this, users won't be able to access their encrypted files on another computer.

**SECURITY ALERT** An alternative to a roaming profile is to copy the user's encryption certificate to the computers that the user uses. You can do this by using the certificate backup and restore process. Simply back up the certificate on the user's original computer, and then restore the certificate on each of the other computers the user logs on to.

EFS has a built-in data recovery system to guard against data loss. This recovery system ensures that encrypted data can be recovered if a user's public-key certificate is lost or deleted. The most common scenario for this is when a user leaves the company and the associated user account is deleted. A manager might have been able to log on to the user's account, check files, and save important files to other folders, but if the user account has been deleted, encrypted files will be accessible only if the encryption





is removed or if the files are moved to an exFAT, FAT, or FAT32 volume (where encryption isn't supported).

To access encrypted files after the user account has been deleted, you need to use a recovery agent. Recovery agents have access to the file encryption key necessary to unlock data in encrypted files. To protect sensitive data, however, recovery agents don't have access to a user's private key or any private key information.

Windows Server won't encrypt files without designated EFS recovery agents. Therefore, recovery agents are designated automatically, and the necessary recovery certificates are generated automatically as well. This ensures that encrypted files can always be recovered.

EFS recovery agents are configured at two levels:

- **Domain** The recovery agent for a domain is configured automatically when the first Windows Server domain controller is installed. By default, the recovery agent is the domain administrator. Through Group Policy, domain administrators can designate additional recovery agents. Domain administrators can also delegate recovery agent privileges to designated security administrators.
- **Local computer** When a computer is part of a workgroup or in a standalone configuration, the recovery agent is the administrator of the local computer by default. Additional recovery agents can be designated. Further, if you want local recovery agents in a domain environment rather than domain-level recovery agents, you must delete the recovery policy from Group Policy for the domain.

You can delete recovery agents if you don't want them to be used. However, if you delete all recovery agents, EFS will no longer encrypt files. One or more recovery agents must be configured for EFS to function.


## Encrypting Directories and Files

With NTFS volumes, Windows Server lets you select files and folders for encryption. When a file is encrypted, the file data is converted to an encrypted format that can be read only by the person who encrypted the file. Users can encrypt files only if they have the proper access permissions. When you encrypt folders, the folder is marked as encrypted, but only the files within it are actually encrypted. All files that are created in or added to a folder marked as encrypted are encrypted automatically. Note that File Explorer shows names of encrypted resources in green.

To encrypt a file or directory, follow these steps:



1. In File Explorer, right-click the file or directory you want to encrypt, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced, and then select the Encrypt Contents To Secure Data check box. Click OK twice.



**NOTE** You can't encrypt compressed files, system files, or read-only files. If you try to encrypt compressed files, the files are automatically uncompressed and then encrypted. If you try to encrypt system files, you get an error.

For an individual file, Windows Server marks the file as encrypted, and then encrypts it. For a directory, Windows Server marks the directory as encrypted, and then encrypts all the files in it. If the directory contains subfolders, Windows Server displays a dialog box that allows you to encrypt all the subfolders associated with the directory. Simply select Apply Changes To This Folder, Subfolders, And Files, and then click OK.

**NOTE** On NTFS volumes, files remain encrypted even when they're moved, copied, or renamed. If you copy or move an encrypted file to an exFAT, FAT, or FAT32 volume, the file is automatically decrypted before being copied or moved. Thus, you must have proper permissions to copy or move the file.

You can grant special access to an encrypted file or folder by right-clicking the file or folder in File Explorer, and then selecting Properties. On the General tab of the Properties dialog box, click Advanced. In the Advanced Attributes dialog box, click Details. In the Encryption Details For dialog box, users who have access to the



encrypted file are listed by name. To allow another user access to the file, click Add. If a user certificate is available for the user, select the user's name in the list provided, and then click OK. Otherwise, click Find User to locate the certificate for the user.



## Working with Encrypted Files and Folders

Previously, I said you can copy, move, and rename encrypted files and folders just like any other files. This is true, but I qualified this by saying “in most cases.” When you work with encrypted files, you’ll have few problems as long as you work with NTFS volumes on the same computer. When you work with other file systems or other computers, you might run into problems. Two of the most common scenarios are the following:

- **Copying between volumes on the same computer** When you copy or move an encrypted file or folder from one NTFS volume to another NTFS volume on the same computer, the files remain encrypted. However, if you copy or move encrypted files to a FAT volume, the files are decrypted before transfer and then transferred as standard files, and therefore end up in their destination as unencrypted files. FAT doesn’t support encryption.
- **Copying between volumes on a different computer** When you copy or move an encrypted file or folder from one NTFS volume to another NTFS volume on a different computer, the files remain encrypted as long as the destination computer allows you to encrypt files and the remote computer is trusted for delegation. Otherwise, the files are decrypted and then transferred as standard files. The same is true when you copy or move encrypted files to a FAT volume on another computer. FAT doesn’t support encryption.

After you transfer a sensitive file that has been encrypted, you might want to confirm that the encryption is still applied. Right-click the file, and then select Properties. On the General tab of the Properties dialog box, click Advanced. The Encrypt Contents To Secure Data option should be selected.

## Configuring Recovery Policies

Recovery policies are configured automatically for domain controllers and workstations. By default, domain administrators are the designated recovery agents for domains, and the local administrator is the designated recovery agent for a standalone workstation.

Group Policy Management Console (GPMC) is a feature you can add to any installation of Windows Server 2016 by using the Add Roles And Features Wizard. The GPMC is also available on Windows desktops when you install the Remote Server





Administration Tools (RSAT). After you add the GPMC to a computer, it is available on the Tools menu in Server Manager. Through the Group Policy console, you can view, assign, and delete recovery agents by following these steps:

1. With the GPMC, you can edit a Group Policy Object (GPO) by right-clicking the GPO, and then selecting Edit on the shortcut menu. The GPMC then opens the Group Policy Management Editor, which you use to manage policy settings.
2. Open the Encrypted Data Recovery Agents node in Group Policy. To do this, access the Administrative Templates policies for Computer Configuration under Windows Settings, Security Settings, Public Key Policies, and then select Encrypting File System.
3. The pane at the right lists the recovery certificates currently assigned. Recovery certificates are listed according to who issued them, who they are issued to, expiration date, purpose, and more.
4. To designate an additional recovery agent, right-click Encrypting File System, and then click Add Data Recovery Agent. This starts the Add Recovery Agent Wizard, which you can use to select a previously generated certificate that has been assigned to a user and mark it as a designated recovery certificate. Click Next.
5. On the Select Recovery Agents page, you can select certificates published in Active Directory or use certificate files. If you want to use a published certificate, click Browse Directory and then—in the Find Users, Contacts, And Groups dialog box—select the user with which you want to work. You'll then be able to use the published certificate of that user. If you want to use a certificate file, click Browse Folders. In the Open dialog box, use the options provided to select and open the certificate file you want to use.

**SECURITY ALERT** Before you designate additional recovery agents, you should consider setting up a root certificate authority (CA) in the domain. Then you can use the Certificates snap-in to generate a personal certificate that uses the EFS Recovery Agent template. The root CA must then approve the certificate request so that the certificate can be used.

6. To delete a recovery agent, select the recovery agent's certificate in the right pane, and then press Delete. When prompted to confirm the action, click Yes to permanently and irrevocably delete the certificate. If the recovery policy is empty (meaning that it has no other designated recovery agents), EFS will be
- 





turned off so that files can no longer be encrypted; existing EFS-encrypted resources won't have a recovery agent.





## Decrypting Files and Directories

File Explorer shows names of encrypted resources in green. If you want to decrypt a file or directory, follow these steps:

1. In File Explorer, right-click the file or directory, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced. Clear the Encrypt Contents To Secure Data check box. Click OK twice.

With files, Windows Server decrypts the file and restores it to its original format. With directories, Windows Server decrypts all the files within the directory. If the directory contains subfolders, you also have the option to remove encryption from the subfolders. To do this, select Apply Changes To This Folder, Subfolders, And Files when prompted, and then click OK.

**TIP** Windows Server also provides a command-line utility called Cipher (Cipher.exe) for encrypting and decrypting your data. Entering **cipher** at a command prompt without additional parameters shows you the encryption status of all folders in the current directory.





# Chapter 21. Using TPM and BitLocker Drive Encryption

Windows Server 2016 has many built-in features designed to protect servers from attacks by individuals accessing a server over a network or from the Internet. But what about when individuals have direct physical access to a server or your data? In these instances, Windows security safeguards don't apply. If someone can start a server—even if it's to another operating system he's installed—he could gain access to any data stored on the server, perhaps even your organization's most sensitive data. In addition, with the increased use of USB flash drives, users often take their data with them, and if they lose the USB flash drive, the data typically has no protection, meaning that anyone who finds the flash drive could read and access the data.

To protect servers and data in these instances, Windows Server 2016 includes Measured Boot, BitLocker Drive Encryption, BitLocker To Go, and the Trusted Platform Module (TPM) Services architecture. Together these features help protect servers and your data whether stored on fixed or removable drives. BitLocker Drive Encryption is a full-volume encryption technology. BitLocker To Go is a virtual-volume encryption technology for USB flash drives. TPM is a feature you can use with BitLocker Drive Encryption to enhance security.

## Creating Trusted Platforms

A server running Windows Server 2016 must be equipped with a compatible TPM in addition to compatible firmware to take advantage of TPM Services. Windows Server 2016 supports TPM version 2.0 or later and requires Trusted Computing Group (TCG)–compliant firmware. Firmware that is TCG-compliant supports the Static Root of Trust Measurement as defined by the TCG. For some configurations of TPM and BitLocker Drive Encryption, you also need to be sure that the firmware supports reading USB flash drives at startup.





## Navigating TPM Essentials

As discussed in the previous chapter, Windows Server 2016 includes the Encrypting File System (EFS) for encrypting files and folders. By using EFS, users can protect sensitive data so that it can be accessed only with their public key certificate.

Encryption certificates are stored as part of the data in a user's profile. So long as users have access to their profiles and the encryption keys they contain, they can access their encrypted files.

Although EFS offers excellent data protection, it doesn't safeguard a server from attack by someone who has direct physical access. In a situation in which a user loses a server, a server is stolen, or an attacker is logging on to a server, EFS might not protect the data because the attacker might be able to gain access to the server before it boots. He could then access the server from another operating system and change the server's configuration. He might then be able to hack into a logon account on the original operating system and log on as the user, or configure the server so that he can log on as a local administrator. Either way, the attacker could eventually gain full access to a server and its data.

To seal a server from physical attack and wrap it in an additional layer of protection, Windows Server 2016 includes the TPM Services architecture. TPM Services protect a server by using a dedicated hardware component called a TPM. A TPM is a microchip that is usually installed on the motherboard of a server, where it communicates with the rest of the system by using a hardware bus. Servers running Windows Server 2016 can use a TPM to provide enhanced protection for data, to ensure early validation of the boot file's integrity, and to guarantee that a disk has not been tampered with while the operating system was offline.

A TPM has the ability to create cryptographic keys and encrypt them so that they can be decrypted only by the TPM. This process, referred to as wrapping or binding, protects the key from disclosure. A TPM has a master "wrapping" key called the Storage Root Key (SRK). The SRK is stored within the TPM to ensure that the private portion of the key is secure.

Servers that have a TPM can create a key that has been not only wrapped but sealed. The process of sealing the key ensures that the key is tied to specific platform measurements and can be unwrapped only when those platform measurements have the same values they had when the key was created. This is what gives TPM-equipped servers increased resistance to attack.

Because TPM stores private portions of key pairs separately from memory controlled by the operating system, keys can be sealed to the TPM to provide absolute assurances





about the state of a system and its trustworthiness. TPM keys are unsealed only when the integrity of the system is intact. Further, because the TPM uses its own internal firmware and logic circuits for processing instructions, it does not rely on the operating system and is not subject to external software vulnerabilities.

The TPM can also be used to seal and unseal data that is generated outside the TPM, and this is where the true power of the TPM lies. In Windows Server 2016, the feature that accesses the TPM and uses it to seal a server is called BitLocker Drive Encryption. Although BitLocker Drive Encryption can be used in both TPM and non-TPM configurations, the most secure method is to use TPM.

When you use BitLocker Drive Encryption and a TPM to seal the boot manager and boot files of a server, the boot manager and boot files can be unsealed only if they are unchanged since they were last sealed. This means that you can use the TPM to validate a server's boot files in the pre-operating system environment. When you seal a hard disk by using TPM, the hard disk can be unsealed only if the data on the disk is unchanged since it was last sealed. This guarantees that a disk has not been tampered with while the operating system was offline.

When you use BitLocker Drive Encryption but do not use a TPM to seal the boot manager and boot files of a server, TPM cannot be used to validate a server's boot files in the pre-operating system environment. This means that in this instance, there is no way to guarantee the integrity of the boot manager and boot files of a server.

## TPM Management and Policies

Windows Server 2016 provides several tools for working with a TPM, including the following:


- **Trusted Platform Module Management** A console for configuring and managing a TPM. You can access this tool by typing **tpm.msc** in the Apps Search box, and then pressing Enter.
- **Manage The TPM Security Hardware** A wizard for creating the required TPM owner password. You can access this tool by typing **tpmunit** in the Apps Search box, and then pressing Enter.

**REAL WORLD** Access to the Trusted Platform Module Management console can be restricted in Group Policy. If you are unable to open the console, check if a Group Policy Object (GPO) being processed includes Management Console restrictions under Windows Components\Microsoft Management Console.



To perform TPM management tasks on a local server, you must be a member of the local server's Administrators group or be logged on as the local server administrator. When you are working with Trusted Platform Module Management, you can determine the exact state of the TPM. If you try to start Trusted Platform Module Management without turning on TPM, you'll get an error stating this. You'll also get an error if you try to run the Initialize The TPM Security Hardware Wizard without turning on TPM.

You will be able to perform management tasks with the TPM tools only when you've turned on TPM in firmware. When you are working with the Trusted Platform Module Management console, you should note the TPM status and the TPM manufacturer information. The TPM manufacturer information shows the version supported. If you want to use TPM with BitLocker Drive Encryption, TPM version 2.0 or later is required.




**NOTE** Although earlier releases of Windows showed the exact TPM state, Windows Server 2016 typically shows either a status of “The TPM is ready for use” or “The TPM is not ready for use.” If the TPM is ready for use, the TPM is on and ownership has been taken.

By default, Windows Server 2016 stores full TPM owner authorization information in the registry of the local server. This significant change enables administrators of the local server to perform TPM management tasks without having to provide the TPM owner password.

The Configure The Level Of TPM Owner Authorization Information Available To The Operating System policy controls the level of authorization information stored in the registry. This policy is found in the Administrative Templates policies for Computer Configuration under System\Trusted Platform Module Services. This policy has three enabled settings:

- **Full** The full TPM owner authorization, the TPM administrative delegation blob, and the TPM user delegation blob are stored in the registry. This setting enables a TPM to be used without requiring remote or external storage of the TPM owner authorization. Note that TPM-based applications designed for earlier versions of Windows or that rely on TPM anti-hammering logic might not support full TPM owner authorization in the registry.
- **Delegated** Only the TPM administrative delegation blob and the TPM user delegation blob are stored in the registry. This level is appropriate for TPM-based applications that rely on TPM anti-hammering logic. When you use this setting, Microsoft recommends that you remotely or externally store the TPM owner authorization.
- **None** No TPM owner authorization information is stored in the registry. Use this setting for compatibility with earlier releases of Windows and for applications that require external or remote storage of the TPM owner authorization. When using this setting, remote or external storage of the TPM owner authorization is required, just as it was in earlier releases of Windows.


**CAUTION** If you change the policy setting from Full to Delegated or vice versa, the full TPM owner authorization value is regenerated and any copies of the original TPM value will be invalid.





When this policy is set to Delegated or None, you'll be prompted for the TPM owner password before you are able to perform most TPM administration tasks.

With earlier releases of Windows, Microsoft recommended remotely storing the TPM owner authorization in Active Directory Domain Services (AD DS) for domain-joined servers, which could be accomplished by enabling the Turn On TPM Backup To Active Directory Domain Services policy, extending schema for the directory, and setting appropriate access controls.



Enabling backup to Active Directory changes the default way TPM owner information is stored. Specifically, when Turn On TPM Backup To Active Directory Domain

Services is enabled and Configure The Level Of TPM Owner Authorization Information Available To The Operating System is disabled or not configured, only the TPM administrative delegation blob and the TPM user delegation blob are stored in the registry. Here, to store the full TPM owner information, you must use the enabled setting of Full (or disable Active Directory backup of the TPM owner authorization).

Related policies under System\Trusted Platform Module Services include:

- [Ignore The Default List Of Blocked TPM Commands](#)
- [Ignore The Local List Of Blocked TPM Commands](#)
- [Standard User Lockout Duration](#)
- [Standard User Individual Lockout Threshold](#)
- [Standard User Total Lockout Threshold](#)

These policies control the way command block lists are used and when lockout is triggered after multiple failed authorization attempts. An administrator can fully reset all lockout-related parameters in the Trusted Platform Module Management console. On the Action menu, click Reset TPM Lockout. When the full TPM owner authorization is stored in the registry, you don't need to provide the TPM owner password. Otherwise, follow the prompts to provide the owner password or select the file containing the TPM owner password.




# Enabling TPM in Firmware

The TPM Services architecture in Windows Server 2016 provides the basic features required to configure and deploy TPM-equipped servers. This architecture can be extended with a feature called BitLocker Drive Encryption, which is discussed in the section “BitLocker Drive Encryption: the Essentials” later in this chapter.

Before you can use TPM, you must enable TPM in firmware. In many cases, servers that have TPM ship with it already enabled. Sometimes, however, you’ll find TPM is not enabled and you’ll need to enable it in firmware. With one of my servers, I needed to do the following:

1. Start the server, and then press F2 startup to access the firmware. In the firmware, I accessed the System BIOS screen and then the System Security screen.
2. On the System Security screen, TPM Security was listed as an option and I set it to On using the mouse. When I enabled TPM, the storage of TPM hierarchies was also enabled, which typically is what you want.



3. To save the changes to the setting and exit the firmware, I clicked Exit. When prompted to confirm whether I wanted to save the changes, I clicked Yes, and the server then rebooted.

With a different server, I needed to do the following:

1. Start the server, and then press F2 during startup to access the firmware. In the firmware, I accessed the Security menu and then the TPM Security screen.
2. On the TPM Security screen, I needed to select the TPM Security check box, and click Apply.
3. A prompt reminded me that I needed to turn off and then restart the server for TPM security to be fully enabled.
4. When I exited firmware, the server rebooted.



Next, you need to initialize and prepare the TPM for first use in software. As part of this process, you take ownership of the TPM, which sets the owner password on the TPM.

After TPM is enabled, you can manage the TPM configuration.



# Initializing and Preparing a TPM for First Use


Initializing a TPM configures it for use on a server so that you can use the TPM to secure volumes on the server's hard drives. The initialization process involves turning on the TPM and then setting ownership of the TPM. By setting ownership of the TPM, you assign a password that helps ensure that only the authorized TPM owner can access and manage the TPM. The TPM password is required to turn off the TPM if you no longer want to use it, and to clear the TPM before the server is recycled. In an Active Directory domain, you can configure Group Policy to save TPM passwords.

By using an administrator account, you can initialize the TPM and create the owner password by completing the following steps:

1. Start the Trusted Platform Module Management console. On the Action menu, click Prepare The TPM to start the Manage The TPM Security Hardware Wizard (tpmunit).

**NOTE** If the Initialize The TPM Security Hardware Wizard detects firmware that does not meet Windows requirements for a TPM or a TPM is not found, you will not be able to continue and should check that the TPM has been turned on in firmware.

2. When the wizard finishes its initial tasks, the TPM owner password is set and you are notified that the TPM is ready for use.



- 
3. Before clicking Close, you should save the TPM owner password. Click Remember My TPM Owner Password. In the Save As dialog box, select a location to save the password backup file, and then click Save.
  4. In the TPM Management console, the status should be listed as “The TPM is ready for use.”

NOTE By default, the password backup file is saved as ServerName .tpm. Ideally, you should save the TPM ownership password to removable media, such as a USB flash drive, and store the media in a secure location. In a domain where the TPM Backup To Active Directory Domain Services policy is applied, you won't have the option to save the TPM password. Here, the password is saved to Active Directory automatically.

# Clearing the TPM

Clearing the TPM erases information stored on the TPM and cancels the related ownership of the TPM. You should clear the TPM when a TPM-equipped server is to be recycled. Clearing the TPM invalidates any stored keys, and data encrypted by these keys can no longer be accessed.

After clearing the TPM, you should take ownership of the TPM. This will write new information to the TPM. You might then want to turn off the TPM so it isn't available for use.

By using an administrator account, you can clear the TPM, take ownership, and then turn off TPM by completing the following steps:

1. Start the Trusted Platform Module Management console. On the Action menu, click Clear TPM. This starts the Manage the TPM Security Hardware Wizard.

**CAUTION** Clearing the TPM resets it to factory defaults. As a result, you lose all keys and data protected by those keys. You do not need the TPM owner password to clear the TPM.

2. Read the warning on the Clear The TPM Security Hardware page, and then click Restart. Click Cancel to exit without clearing the TPM.



3. Typically, hardware designed for Windows Server 2016 can automatically complete the re-initialization process. On other hardware, you'll need physical access to the server to respond to the manufacturer's firmware confirmation



prompt. For example, you may need to press F12 during startup to clear, enable, and activate the TPM, or press ESC to reject the change request and continue loading the operating system.



# Changing the TPM Owner Password

You can change the TPM password at any time. The key reason to do this is if you suspect that the TPM owner password has been compromised. Your company's security policy also might require TPM owner password changes in certain situations.

To change the TPM owner password, complete the following steps:

1. Start the Trusted Platform Module Management console. On the Action menu, click Change Owner Password. This starts the Manage The TPM Security Hardware Wizard.
2. When the full TPM owner authorization is stored in the registry, you don't need to provide the TPM owner password. Otherwise, follow the prompts to provide the owner password or select the file containing the TPM owner password.
3. On the Create The TPM Owner Password page, you can elect to create the password automatically or manually.



4. If you want the wizard to create the password for you, select Automatically Create The Password (Recommended). The new TPM owner password is displayed. Click Change Password.
5. If you want to create the password, select Manually Create The Password. Enter and confirm a password of at least eight characters, and then click Change Password.
6. Before clicking Close, you might want to save the TPM owner password. Click Remember My TPM Owner Password. In the Save As dialog box, select a location to save the password backup file, and then click Save.



# BitLocker Drive Encryption: the Essentials

BitLocker is built into all editions of Windows Server 2016 and added as a feature for all editions of Windows Server. Although BitLocker Drive Encryption and BitLocker To Go are often referred to simply as BitLocker, they are separate but similar features. BitLocker Drive Encryption is designed to protect the data on the internal hard drives of lost, stolen, or inappropriately decommissioned servers and is a volume-level encryption technology. BitLocker To Go is designed to protect the data on removable data drives, such as external hard drives and USB flash drives, and is a virtual-volume encryption technology. Standard BitLocker encrypts by wrapping the entire volume or only the used portion of the volume in protected encryption. BitLocker To Go, on the other hand, creates a virtual volume on a USB flash drive. This virtual volume is encrypted by using an encryption key stored on the USB flash drive.







# Understanding BitLocker Drive Encryption

On a server without BitLocker Drive Encryption, a user with direct physical access to the server has a variety of ways he could gain full control and then access the server's data, whether that data is encrypted with EFS or not. For example, a user could use a boot disk to boot the server and reset the administrator password. A user could also install and then boot to a different operating system and then use this operating system to unlock the other installation.

BitLocker Drive Encryption prevents all access to a server's drives except by authorized personnel by wrapping entire drives or only the used portion of drives in tamper-proof encryption. If an unauthorized user tries to access a BitLocker-encrypted drive, the encryption prevents the user from viewing or manipulating the protected data in any way. This dramatically reduces the risk of an unauthorized person gaining access to confidential data through offline attacks.

**IMPORTANT** BitLocker Drive Encryption is meant to be used when a server is not in a physically secure location and requires additional protection. However, the encryption processes can reduce disk throughput, which in turn can affect a server's overall performance.

BitLocker Drive Encryption can use a TPM to validate the integrity of a server's boot manager and boot files at startup and to guarantee that a server's hard disk has not been tampered with while the operating system was offline. BitLocker Drive Encryption also stores measurements of core operating system files in the TPM.

Every time the server is started, Windows validates the boot files, the operating system files, and any encrypted volumes to ensure that they have not been modified while the operating system was offline. If the files have been modified, Windows alerts the user and refuses to release the key required to access Windows. The server then goes into Recovery mode, prompting the user to provide a recovery key before it permits access to the boot volume. The Recovery mode is also used if a BitLocker-encrypted disk drive is transferred to another system.

BitLocker Drive Encryption can be used in both TPM and non-TPM servers. If a server has a TPM, BitLocker Drive Encryption uses the TPM to provide enhanced protection for your data and to ensure early boot file integrity. These features together help prevent unauthorized viewing and accessing of data by encrypting the entire Windows volume and by safeguarding the boot files from tampering. If a server doesn't have a TPM or its TPM isn't compatible with Windows, BitLocker Drive Encryption can be used to encrypt entire volumes, and in this way protect the volumes





from tampering. This configuration, however, doesn't offer the added security of early boot file integrity validation.

On servers with a compatible TPM that is initialized, BitLocker Drive Encryption typically uses one of the following TPM modes:

- **TPM-Only** In this mode, only TPM is used for validation. When the server boots, TPM is used to validate the boot files, the operating system files, and any encrypted volumes. Because the user doesn't need to provide an additional startup key, this mode is transparent to the user, and the user logon experience is unchanged. However, if the TPM is missing or the integrity of files or volumes has changed, BitLocker enters Recovery mode and requires a recovery key or password to regain access to the boot volume.
- **TPM and PIN** In this mode, both TPM and a user-entered numeric key are used for validation. When the server boots, TPM is used to validate the boot files, the operating system files, and any encrypted volumes. The user must enter a PIN when prompted to continue startup. If the user doesn't have the PIN or is unable to provide the correct PIN, BitLocker enters Recovery mode instead of booting to the operating system. As before, BitLocker also enters Recovery mode if the TPM is missing or the integrity of boot files or encrypted volumes has changed.
- **TPM and Startup Key** In this mode, both TPM and a startup key are used for validation. When the server boots, TPM is used to validate the boot files, the operating system files, and any encrypted volumes. The user must have a USB flash drive with a startup key to log on to the server. If the user doesn't have the startup key or is unable to provide the correct startup key, BitLocker enters Recovery mode. As before, BitLocker also enters Recovery mode if the TPM is missing or the integrity of boot files or encrypted volumes has changed.
- **TPM and Smart Card Certificate** In this mode, both TPM and a smart card certificate are used for validation. When the server boots, TPM is used to validate the boot files, the operating system files, and any encrypted volumes. The user must have a smart card with a valid certificate to log on to the server. If the user doesn't have a smart card with a valid certificate and is unable to provide one, BitLocker enters Recovery mode. As before, BitLocker also enters Recovery mode if the TPM is missing or the integrity of boot files or encrypted volumes has changed.

Network unlock enables the system volume on a server with TPM to be automatically unlocked on startup, if the server is joined and connected to a domain. When not joined and connected to a domain, other means of validation can be used, such as a startup PIN.



On servers without a TPM or on servers that have incompatible TPMs, Windows Server 2016 can be configured to use an unlock password for the operating system drive. To configure this, you must enable the Configure Use Of Passwords For Operating System Drives policy in the Administrative Templates policies for Computer Configuration under Windows Components\BitLocker Drive Encryption\Operating System Drives. The unlock password can be configured with minimum length and complexity requirements. The default minimum password length is 8 characters, meaning the password must be at least 8 characters. Complexity requirements can be:

- Always validated using the Require Password Complexity setting.
- Validated if possible using the Allow Password Complexity setting.
- Not validated using the Do Not Allow Password Complexity setting.

The screenshot shows the 'Configure use of passwords for operating system drives' Group Policy setting. At the top, there are 'Previous Setting' and 'Next Setting' buttons. Below them are three radio buttons: 'Not Configured', 'Enabled' (which is selected), and 'Disabled'. To the right of the radio buttons is a 'Comment:' text box. Below the radio buttons is a 'Supported on:' dropdown menu showing 'At least Windows Server 2012 or Windows 8'. Under the 'Options:' section, there is a 'Configure password complexity for operating system drives:' section. It includes a dropdown menu set to 'Require password complexity', a 'Minimum password length for operating system drive:' spinner box set to '8', and a note: 'Note: You must enable the "Password must meet complexity requirements" policy setting for the password complexity setting to take effect.' There is also an unchecked checkbox labeled 'Require ASCII-only passwords for removable OS drives'. To the right of the options is a 'Help:' section with a scrollable text area containing detailed information about the policy setting and its requirements. At the bottom of the help section is another note: 'Note: These settings are enforced when turning on BitLocker, not when unlocking a volume. BitLocker will allow unlocking a drive with any of the protectors available on the drive.'

The unlock password is validated when you enable BitLocker Drive Encryption and set the password, and whenever the password is changed by a user. With required complexity, you can only set a password (and enable encryption) when the server can connect to a domain controller and validate the complexity of the password. With allowed complexity, the server will attempt to validate the complexity of the password when you set it but will allow you to continue and enable encryption if no domain controllers are available.

On servers without a TPM or on servers that have incompatible TPMs, BitLocker Drive

Encryption also can use Startup Key Only or Smart Card Certificate Only mode.

Startup Key Only mode requires a USB flash drive containing a startup key. The user

inserts the USB flash drive in the server before turning it on. The key stored on the flash drive unlocks the server.

Smart Card Certificate Only mode requires a smart card with a valid certificate. The user validates the smart card certificate after turning on the server. The certificate unlocks the server.

It's also important to point out that standard users can reset the BitLocker PIN and password on operating system drives, fixed data drives, and removable data drives. If you don't want standard users to be able to perform these tasks, enable the Disallow Standard Users From Changing The PIN Or Password policy. This Computer Configuration policy is found under Windows Components\BitLocker Drive Encryption\Operating System Drives.



Several important changes have been made to BitLocker Drive Encryption since the technology was first implemented. Specifically, you can do the following:

- **Encrypt FAT volumes and NTFS volumes.** When you encrypt FAT volumes, you have the option of specifying whether encrypted volumes can be unlocked and viewed. This option is configured through Group Policy and is enabled when you turn on BitLocker. In the Administrative Templates policies for Computer Configuration under Windows Components\BitLocker Drive Encryption, there are separate policies for earlier versions of Windows that allow FAT-formatted fixed drives and FATformatted removable drives to be unlocked and viewed.
- **Allow a data-recovery agent to be used with BitLocker Drive Encryption.** This option is configured through Group Policy. The data-recovery agent allows an encrypted volume to be unlocked and recovered by using a recovery agent's personal certificate or a 48-digit recovery password. You can optionally save the recovery information in Active Directory. In the Administrative Templates policies for Computer Configuration, there are separate policies for operating system volumes, other fixed drives, and removable drives.
- **Deny write access to removable data drives not protected with BitLocker.** This option is configured through Group Policy. If you enable this option, users have



read-only access to unencrypted removable data drives and read/write access to encrypted removable data drives.

In a domain, domain administrators are the default data-recovery agents. A homegroup or workgroup has no default data-recovery agent, but you can designate one. Any user you want to designate as a data-recovery agent needs a personal encryption certificate. You can generate a certificate by using the Cipher utility and then use the certificate to assign the data-recovery agent in Local Security Policy under Public Key Policies\BitLocker Drive Encryption.

Earlier releases of Windows support AES encryption with a diffuser. Windows 8 and Windows Server 2016 move away from this to support standard AES with 128-bit encryption by default or 256-bit encryption (if you enable the Choose Drive Encryption Method And Cipher Strength policy to set the cipher strength to 256-bit encryption). The cipher strength must be set prior to turning on BitLocker. Changing the cipher strength has no effect if the drive is already encrypted or encryption is in progress.





# Hardware Encryption, Secure Boot, and Network Unlock

Most BitLocker-related settings are controlled with the Administrative Templates policies for Computer Configuration under Windows Components\BitLocker Drive Encryption. Windows Server 2016 supports disk drives with hardware encryption (referred to as encrypted hard drives ). Encryption in hardware is faster and moves the processing burden from the server's processor to the hardware processor on the hard disk. By default, if a server has hardware encryption, Windows Server 2016 will use it with BitLocker.

In Group Policy, you can precisely control whether to permit software-based encryption when hardware encryption is not available and whether to restrict encryption to those algorithms and cipher strengths supported by hardware. Do this by enabling the Configure Use Of Hardware-Based Encryption For Fixed Data Drives policy and configuring the related options. When the policy is enabled, you must specifically allow software-based encryption when hardware-based encryption isn't available.

The screenshot shows the Group Policy setting 'Configure use of hardware-based encryption for operating system drives'. The policy is set to 'Enabled'. The 'Supported on' field indicates 'At least Windows Server 2012, Windows 8 or Windows RT'. Under 'Options', the checkbox 'Use BitLocker software-based encryption when hardware encryption is not available' is checked. The checkbox 'Restrict encryption algorithms and cipher suites allowed for hardware-based encryption' is unchecked. Below this, a text box contains the IP address '2.16.840.1.101.3.4.1.2;2.16.840.1.101.3.4.1.'. The 'Help' section on the right explains that this policy allows managing BitLocker's use of hardware-based encryption and specifies which encryption algorithms it can use with hardware-based encryption. It also notes that enabling this policy allows specifying additional options to control whether BitLocker software-based encryption is used instead of hardware-based encryption on computers that do not support hardware-based encryption and whether to restrict the encryption algorithms and cipher suites used with hardware-based encryption.

**MORE INFO** The Choose Drive Encryption Method And Cipher Strength policy doesn't apply to hardware-based encryption. You can set the encryption methods you want for hardware-based policy under Fixed Data Drives, in the Configure Use Of Hardware-Based Encryption For Fixed Data Drives policy. With hardware-based encryption, the encryption algorithm is set when the drive is partitioned.

Next, you might want to configure policy to control the permitted encryption types. Windows Server 2016 allows users to encrypt full volumes or used space only.

Encrypting full volumes takes longer, but it's more secure because the entire volume is protected. Encrypting used space protects only the portion of the drive used to store data. By default, either option can be used. If you want to allow only one type or the other, enable and configure related Enforce Drive Encryption Type policy for BitLocker. There are separate Enforce Drive Encryption Type policies for the operating system, fixed data, and removable data drives.

Enforce drive encryption type on operating system drives

Previous Setting Next Setting

☐ Not Configured ☒ Enabled ☐ Disabled

Comment:

Supported on: At least Windows Server 2012 or Windows 8

Options:

Select the encryption type:

Full encryption

Help:

This policy setting allows you to configure the encryption type used by BitLocker Drive Encryption. This policy setting is applied when you turn on BitLocker. Changing the encryption type has no effect if the drive is already encrypted or if encryption is in progress. Choose full encryption to require that the entire drive be encrypted when BitLocker is turned on. Choose used space only encryption to require that only the portion of the drive used to store data is encrypted when BitLocker is turned on.

**REAL WORLD** In high-security environments, you will want to encrypt entire volumes. At the time of this writing, and unless fixed with a future update or service pack, deleted files appear as free space when you encrypt used space only. As a result, until the files are wiped or overwritten, information in the files could be recovered with certain tools.

Operating system drives are handled as special cases. Windows Server 2016 enables you to pre-provision BitLocker so that you can turn on encryption prior to installation. Windows Server 2016 also can be configured to do the following:

- **Require additional authentication at startup.** If you enable and configure the related policy, **Require Additional Authentication At Startup**, user input is required, even if the platform lacks a pre-boot input capability. To enable a USB keyboard to be used on such a platform in the pre-boot environment, you should set the **Enable Use Of BitLocker Authentication Requiring Preboot Keyboard Input On Slates** policy to **Enabled**.
- **Allow secure boot for integrity validation.** Secure boot is used by default to verify boot configuration data (BCD) settings according to the TPM validation profile settings (also referred to as **Secure Boot** policy). When you use secure boot, the settings of the **Use Enhanced Boot Configuration Data Validation Profile** policy are

ignored (unless you specifically disable secure boot support by setting Allow Secure Boot For Integrity Validation to Disabled).

You set TPM validation profile settings by platform. For BIOS-based firmware, you use the Configure TPM Platform Validation Profile For BIOS-Based Firmware Configurations policy. For Unified Extensible Firmware Interface (UEFI)–based firmware, you use the Configure TPM Platform Validation Profile For Native UEFI Firmware Configurations policy. When you enable these policies, you specify exactly which platform configuration registers to validate during boot.

For BIOS-based firmware, Microsoft recommends validating Platform Configuration Registers (PCRs) 0, 2, 4, 8, 9, 10, and 11. For UEFI firmware, Microsoft recommends validating PCRs 0, 2, 4, 7, and 11. In both instances, PCR 11 validation is required for BitLocker protection to be enforced. PCR 7 validation is required to support secure boot with UEFI (and you'll need to enable this by selecting the related option).

The screenshot shows the 'Configure TPM platform validation profile for BIOS-based firmware configurations' Group Policy window. At the top, there are 'Previous Setting' and 'Next Setting' buttons. Below them, the 'Not Configured', 'Enabled' (selected), and 'Disabled' radio buttons are visible. A 'Comment:' text box is on the right. The 'Supported on:' dropdown menu is set to 'At least Windows Server 2012 or Windows 8'. Under the 'Options:' section, a list of Platform Configuration Registers (PCRs) is shown with checkboxes: PCR 0 (checked), PCR 1 (unchecked), PCR 2 (checked), PCR 3 (unchecked), PCR 4 (checked), PCR 5 (unchecked), PCR 6 (unchecked), PCR 7 (unchecked), and PCR 8 (checked). A 'Help:' section on the right provides detailed information about the policy, including its purpose, applicability to BIOS and UEFI systems, and instructions for enabling it before BitLocker.

Configure TPM platform validation profile for BIOS-based firmware configurations

Previous Setting Next Setting

☐ Not Configured ☒ Enabled ☐ Disabled

Comment:

Supported on: At least Windows Server 2012 or Windows 8

Options:

PCR 0: Core Root of Trust of Measurement (CRTM), BIOS, and Platform Extensions ☒

PCR 1: Platform and Motherboard Configuration and Data ☐

PCR 2: Option ROM Code ☒

PCR 3: Option ROM Configuration and Data ☐

PCR 4: Master Boot Record (MBR) Code ☒

PCR 5: Master Boot Record (MBR) Partition Table ☐

PCR 6: State Transition and Wake Events ☐

PCR 7: Computer Manufacturer-Specific ☐

PCR 8: NTFS Boot Sector ☒

Help:

This policy setting allows you to configure how the computer's Trusted Platform Module (TPM) security hardware secures the BitLocker encryption key. This policy setting does not apply if the computer does not have a compatible TPM or if BitLocker has already been turned on with TPM protection.

Important: This group policy only applies to computers with BIOS configurations or to computers with UEFI firmware with a Compatibility Service Module (CSM) enabled. Computers using a native UEFI firmware configuration store different values into the Platform Configuration Registers (PCRs). Use the "Configure TPM platform validation profile for native UEFI firmware configurations" group policy setting to configure the TPM PCR profile for computers using native UEFI firmware.

If you enable this policy setting before turning on BitLocker, you can configure the boot components that the TPM will validate before unlocking access to the BitLocker-encrypted operating system drive. If any of these components change while BitLocker

**IMPORTANT** Windows Server 2016 is designed to be run on servers with UEFI. UEFI doesn't replace all the functionality in either BIOS or EFI and can, in fact, be wrapped around BIOS or EFI. When a server has UEFI and is running Windows Server 2016, UEFI is the first link in the chain of trust for secure boot. UEFI 2.3.1 and later are able to run internal integrity checks that verify the firmware's digital signature before running it. If the firmware's digital signature has been modified or replaced (as an example, by a firmware rootkit), the firmware will not load.





With Secure Boot, firmware also verifies the digital signature on the Windows boot loader as part of initialization. If a rootkit is installed and the Windows boot loader has been modified, the server will be prevented from starting. After the Secure Boot, the boot loader verifies the digital signature of the operating system kernel as part of Trusted Boot. Also as part of Trusted Boot, the kernel in turn verifies all remaining boot components, including boot drivers and startup files. Finally, Measured Boot allows third-party software running on a remote server to verify the security of every startup component as well.

When you protect a server with BitLocker, you can require additional authentication at startup. Typically, this means a user is required to have a startup key on a USB flash drive, a startup PIN, or both. The Network Unlock feature provides this additional layer of protection without requiring the startup key, startup PIN, or both by automatically unlocking the operating system drive when a server is started, if the following conditions are met:

- The BitLocker-protected server has an enabled TPM.
- The server is on a trusted, wired network.
- The server is joined to and connected to a domain.
- A Network Unlock server with an appropriate Network Unlock certificate is available.

Because the server must be joined to and connected to the domain for network unlock to work, user authentication is still required when a server is not connected to the domain. When connected to the domain, the client server connects to a Network Unlock server to unlock the system drive. Typically, the Network Unlock server is a domain controller configured to use and distribute Network Unlock certificates to clients. The Network Unlock certificates in turn are used to create the network unlock keys.

You can configure a domain controller to distribute this certificate to clients. To do this, create an X.509 certificate for the server, for example by using Certmsg.mc, then use the BitLocker Driver Encryption Network Unlock Certificate setting to add this certificate to a GPO applied to the domain controller. You'll find this Computer Configuration setting under Windows Settings\Security Settings\Public Key Policies.

Finally, Windows Server 2016 also enables you to provision BitLocker during operating system deployment. You can do this from the Windows PreInstallation Environment (Windows PE). It's important to point out that Windows PowerShell includes a DISM module. As this module doesn't support wildcards when searching





for feature names, you can use the `Get-WindowsOptionalFeature` cmdlet to list feature names, as shown in this example:

```
get-windowsoptionalfeature -online | ft
```

To completely install BitLocker and related management tools, use the following command:

```
enable-windowsoptionalfeature -online -featurename bitlocker,bitlocker-utilities, bitlocker-networkunlock -all
```





# Deploying BitLocker Drive Encryption

Deploying BitLocker Drive Encryption in an enterprise changes the way administrators and users work with servers. A server with BitLocker Drive Encryption usually requires user intervention to boot to the operating system—a user must enter a PIN, insert a USB flash drive containing a startup key, or use a smart card with a valid certificate. Because of this requirement, after you deploy BitLocker Drive Encryption, you can no longer be assured that you can perform remote administration that requires a server to be restarted without having physical access to the server—someone might need to be available to enter the required PIN, insert the USB flash drive with the startup key, or use a smart card with a valid certificate.

To work around this issue, you can configure network unlock on your trusted, wired networks. Before you use BitLocker Drive Encryption, you should perform a thorough evaluation of your organization's servers. You need to develop plans and procedures for the following:

- Evaluating the various BitLocker authentication methods and applying them as appropriate
- Determining whether servers support TPM, and thus whether you must use TPM or non-TPM BitLocker configurations
- Storing, using, and periodically changing encryption keys, recovery passwords, and other validation mechanisms used with BitLocker

You need to develop procedures for tasks such as these:

- Performing daily operations with BitLocker-encrypted drives
- Providing administrative support for BitLocker-encrypted drives
- Recovering servers with BitLocker-encrypted drives

These procedures need to take into account the way BitLocker encryption works and the requirements to have PINs, startup keys, smart cards, and recovery keys available whenever you work with BitLocker-encrypted servers. After you evaluate your organization's servers and develop basic plans and procedures, you need to develop a configuration plan for implementing BitLocker Drive Encryption.

Several implementations of BitLocker Drive Encryption are available. Although servers running Windows Server 2016 can work with any of the available versions, earlier versions of Windows can't necessarily work with the latest version of BitLocker. For example, you might need to configure Group Policy to allow access from earlier versions of Windows.



To turn on BitLocker Drive Encryption on the drive containing the Windows operating system, the drive must have at least two partitions:



- The first partition is for BitLocker Drive Encryption. This partition, designated as the active partition, holds the files required to start the operating system and is not encrypted.
- The second is the primary partition for the operating system and your data. This partition is encrypted when you turn on BitLocker.

When you install Windows Server 2016, an additional partition is created automatically during setup. By default, this additional partition is used by the Windows Recovery

Environment (Windows RE). However, if you enable BitLocker on the system volume, BitLocker usually moves Windows RE to the system volume and then uses the additional partition for BitLocker. Because of this, during BitLocker setup, you'll typically see a notification that a new system drive will be created from free space on the original system drive and also be notified that the new drive will not have a drive letter.



Using BitLocker on a hard disk is easy. On a server with a compatible TPM, you must initialize the TPM as discussed in the “Initializing and Preparing a TPM for First Use” section earlier in this chapter, and then you need to enable BitLocker. On a server without a compatible TPM, you need to enable BitLocker only on your hard disk.

You can use local Group Policy and Active Directory–based Group Policy to help manage and maintain TPM and BitLocker configurations. Group Policy settings for TPM Services are found in Administrative Templates policies for Computer Configuration under System\Trusted Platform Module Services. Group Policy settings for BitLocker are found in Administrative Templates policies for Computer Configuration under Windows Components\BitLocker Drive Encryption. There are separate subfolders for fixed data drives, operating system drives, and removable data drives.

Policies you might want to configure include the following:

### **Trusted Platform Module Services policies**



- Configure The Level of TPM Owner Authorization Information Available To The Operating System
- Configure The List Of Blocked TPM Commands
- Ignore The Default List Of Blocked TPM Commands
  - Ignore The Local List Of Blocked TPM Commands
- Standard User Individual Lockout Threshold
- Standard User Lockout Duration
- Standard User Total Lockout Threshold
- Turn On TPM Backup To Active Directory Domain Services

### **BitLocker Drive Encryption policies**

- Choose Default Folder For Recovery Password
- Choose Drive Encryption Method And Cipher Strength
- Prevent Memory Overwrite On Restart
- Provide The Unique Identifiers For Your Organization
- Validate Smart Card Certificate Usage Rule Compliance

### **Fixed Drive policies**

- Allow Access To BitLocker-Protected Fixed Data Drives From Earlier Versions Of Windows
- Choose How BitLocker-Protected Fixed Drives Can Be Recovered
- Configure Use Of Hardware-Based Encryption For Fixed Data Drives
- Configure Use Of Passwords For Fixed Data Drives
- Configure Use Of Smart Cards On Fixed Data Drives
- Deny Write Access To Fixed Drives Not Protected By BitLocker
- Enforce Drive Encryption Type On Fixed Data Drives

### **Operating System Drive policies**

- Allow Enhanced PINs For Startup
- Allow Network Unlock At Startup
- Allow Secure Boot For Integrity Validation
- Choose How BitLocker-Protected Operating System Drives Can Be Recovered
- Configure Minimum PIN Length For Startup
- Configure TPM Platform Validation Profile For BIOS-Based Firmware Configurations
- Configure TPM Platform Validation Profile For Native UEFI Firmware Configurations





- •Configure TPM Platform Validation Profile (Windows Vista, Windows 7, Windows Server 2008, Windows Server 2008 R2)
- Configure Use Of Passwords For Operating System Drives
- Disallow Standard Users From Changing The PIN Or Password
- Enforce Drive Encryption Type On Operating System Drives
- Enable User Of BitLocker Authentication Requiring Preboot Keyboard Input On Slates
  - Require Additional Authentication At Startup
- Reset Platform Validation Data After BitLocker Recovery

### **Removable Data Drive policies**

- Allow Access To BitLocker-Protected Removable Data Drives From Earlier Versions Of Windows
- Choose How BitLocker-Protected Removable Drives Can Be Recovered
- Configure Use Of Hardware-Based Encryption For Removable Data Drives
- Configure Use Of Passwords For Removable Data Drives
- Configure Use Of Smart Cards On Removable Data Drives
- Control Use Of BitLocker On Removable Drives
- Deny Write Access To Removable Drives Not Protected By BitLocker
- Enforce Drive Encryption Type On Removable Data Drives

Active Directory includes TPM and BitLocker recovery extensions for Server objects. For TPM, the extensions define a single property of the Server object, called ms-TPMOwnerInformation. When the TPM is initialized or when the owner password is changed, the hash of the TPM ownership password can be stored as a value of the msTPM-OwnerInformation attribute on the related Server object. For BitLocker, these extensions define Recovery objects as child objects of Server objects and are used to store recovery passwords and associate them with specific BitLocker-encrypted volumes.

By default, Windows Server 2016 stores the full TPM owner authorization, the TPM administrative delegation blob, and the TPM user delegation in the registry. Because of this change, you no longer have to save this information separately to Active Directory for backup and recovery purposes. For more information, see the “TPM Management and Policies” section earlier in this chapter.

To ensure that BitLocker recovery information is always available, you can configure Group Policy to save recovery information in Active Directory as follows:





- With Choose How BitLocker-Protected Fixed Drives Can Be Recovered, enable the policy and accept the default options to allow data recovery agents and save the recovery information in Active Directory.
- With Choose How BitLocker-Protected Operating System Drives Can Be Recovered, enable the policy and accept the default options to allow data recovery agents and save the recovery information in Active Directory.
- With Choose How BitLocker-Protected Removable Drives Can Be Recovered, enable the policy and accept the default options to allow data recovery agents, and then save the recovery information in Active Directory.

REAL WORLD For Federal Information Processing Standard (FIPS) compliance, you







cannot create or save BitLocker recovery passwords. Instead, you need to configure Windows to create recovery keys. The FIPS setting is located in the Security Policy Editor at Local Policies\Security Options\System Cryptography: Use FIPS Compliant Algorithms For Encryption, Hashing, And Signing.

To configure BitLocker to use recovery keys, enable the security option System Cryptography: Use FIPS Compliant Algorithms For Encryption, Hashing, And Signing in local Group Policy or Active Directory–based Group Policy as appropriate. With this setting enabled, users can only generate recovery keys.





# Managing BitLocker Drive Encryption

You can configure and enable BitLocker Drive Encryption on both system volumes and data volumes. When you encrypt system volumes, you must unlock the server at startup, typically by using a TPM and network unlock when connected to the domain in addition to a TPM, a startup key, a startup PIN, or any required or optional combination of these. To enforce the strictest and highest security possible, use all three authentication methods.

In the current implementation of BitLocker, you do not have to encrypt a server's system volume prior to encrypting a server's data volumes. When you use encrypted data volumes, the operating system mounts BitLocker data volumes as it would any other volume, but it requires either a password or a smart card with a valid certificate to unlock the drive.

The encryption key for a protected data volume is created and stored independently from the system volume and all other protected data volumes. To allow the operating system to mount encrypted volumes, the key chain protecting the data volume is stored in an encrypted state on the operating system volume. If the operating system enters Recovery mode, the data volumes are not unlocked until the operating system is out of Recovery mode.

Setting up BitLocker Drive Encryption requires these steps:

1. Partitioning a server's hard disks appropriately and installing the operating system (if you are configuring a new server). Windows Setup partitions the drives for you automatically. However, the volume where BitLocker data is stored must always be the active, system volume.
2. Initializing and configuring a server's TPM (if applicable).
3. Turning on the BitLocker Drive Encryption feature (as necessary).
4. Checking firmware to ensure that the server is set to start first from the disk containing the active, system partition and the boot partition, not from USB or CD/DVD drives (applicable only when you encrypt system volumes).
5. Turning on and configuring BitLocker Drive Encryption.

**NOTE** When you are using a Microsoft account on a non-domain-joined server, you have an additional save option. You can save the recovery key to the Windows Live SkyDrive. The user's SkyDrive account will then contain a BitLocker folder with a separate file for each saved recovery key.

After you've turned on and configured BitLocker encryption, you can use several techniques to maintain the environment and perform recovery.

## Preparing for BitLocker Drive Encryption

As discussed previously, BitLocker Drive Encryption can be used in a TPM or a nonTPM configuration. Either configuration requires some preliminary work before you can turn on and configure BitLocker Drive Encryption.

You can determine the readiness status of a server by accessing the BitLocker Drive Encryption console. In Control Panel, click System And Security, and then click BitLocker Drive Encryption. If the system isn't properly configured, you'll receive an error message. Note the following:

- If you get an error message related to TPM on a server with a compatible TPM, refer to the “Enabling TPM in Firmware” section earlier in this chapter to learn more about TPM states and enabling TPM in firmware.
- If you get an error message related to TPM on a server with an incompatible TPM or no TPM, you need to change the server's Group Policy settings so that you can turn on BitLocker Drive Encryption without a TPM.




You can configure policy settings for BitLocker encryption in local Group Policy or in Active Directory–based Group Policy. In local policy, you apply the settings to the server's local GPO. For domain policy, you apply the settings to a Group Policy object processed by the server. While you are working with domain policy, you can also specify requirements for servers with a TPM.

To configure the way BitLocker can be used with or without a TPM, follow these steps:

1. Open the appropriate GPO for editing in the Group Policy Management Editor.
2. In the Administrative Templates policies for Computer Configuration under Windows Components\BitLocker Drive Encryption\Operating System Drives, double-click the Require Additional Authentication At Startup setting.

**IMPORTANT** There are several versions of this policy and they are specific to the operating system. Configure the version or versions of this policy that are appropriate for your working environment and the servers to which the policy will be applied. The options for each related policy are slightly different because the TPM features supported are slightly different for each.

3. In the Require Additional Authentication At Startup dialog box define the policy setting by selecting Enabled.



4. Do one of the following:
  - If you want to allow BitLocker to be used without a compatible TPM, select the Allow BitLocker Without A Compatible TPM check box. This changes the policy setting so that you can use BitLocker encryption with a password or startup key on a server without a TPM.
  - If you want to require BitLocker to be used with a TPM, clear the Allow BitLocker Without A Compatible TPM check box. This changes the policy setting so that you

can use BitLocker encryption on a server with a TPM by using a startup PIN, a startup key, or both.

5. On a server with a compatible TPM, several authentication methods can be used at startup to provide added protection for encrypted data. These authentication methods can be allowed or required. The methods available depend on the operating system specific version of the policy with which you are working. Use

the following guidelines to help you configure the options:

- If you want to allow TPM to be used at startup, set Configure TPM Startup to Allow TPM and use Do Not Allow for all other settings.
- If you want to require TPM to be used at startup, set Configure TPM Startup to Require TPM and use Do Not Allow for all other settings.
- If you want to use TPM only with a startup key, set Configure TPM Startup to Allow or Require TPM as appropriate and then set Configure TPM Startup Key to Allow or Require Startup Key with TPM as appropriate. Use Do Not Allow for all other settings.
- If you want to use TPM only with a startup PIN, set Configure TPM Startup to Allow or Require TPM as appropriate and then set Configure TPM Startup PIN to Allow or Require Startup PIN with TPM as appropriate. Use Do Not Allow for all other settings.
- If you want to use TPM only with a startup key and PIN, set Configure TPM Startup to Allow or Require TPM and then set Configure TPM Startup Key And Pin to Allow or Require Startup Key and PIN with TPM. Use Do Not Allow for all other settings. • If you want to allow TPM with any other authentication method, set allow options to allow.

6. Click OK to save your settings. This policy is enforced the next time Group Policy is applied.

7. Close the Group Policy Management Editor. To apply Group Policy immediately to the server you are logged on to, enter **gpupdate.exe /force** in the Apps Search box, and then press Enter.

Servers that have a startup key or a startup PIN also have a recovery password or certificate. The recovery password or certificate is required in the event any of the following occur:

- Changes are made to the system startup information
- The encrypted drive must be moved to another server
-



## The user is unable to provide the appropriate startup key or PIN

The recovery password or certificate should be managed and stored separately from the startup key or startup PIN. Although users are given the startup key or startup PIN, administrators should be the only ones with the recovery password or certificate. As an administrator, you need the recovery password or certificate to unlock the encrypted data on the volume if BitLocker enters a locked state. Generally, unless you use a common data-recovery agent, the recovery password or certificate is unique to this particular BitLocker encryption. This means you cannot use it to recover encrypted data from any other BitLocker-encrypted volume—even from other BitLocker-encrypted volumes on the same server. To increase security, you should store startup keys and recovery data apart from the server.

When BitLocker is installed, the BitLocker Drive Encryption console is available in Control Panel. Your configuration options for BitLocker depend on whether the server has a TPM and on how you've configured Group Policy.





## Enabling BitLocker on Nonsystem Volumes

Encrypting a nonsystem volume protects the data stored on the volume. Any volume formatted with FAT, FAT32, exFAT, NTFS or ReFS can be encrypted with BitLocker. The length of time it takes to encrypt a drive depends on the amount of data to encrypt, the processing power of the server, and the level of activity on the server.

Before you enable BitLocker, you should configure the appropriate Fixed Data Drive policies and settings in Group Policy, and then wait for Group Policy to be refreshed. If you don't do this and you enable BitLocker, you might need to turn BitLocker off, and then turn BitLocker back on because certain state and management flags are set when you turn on BitLocker. To ensure that you can recover an encrypted volume, you should allow data-recovery agents and store recovery information in Active Directory.

To enable BitLocker encryption on a nonsystem volume, follow these steps:

1. In File Explorer, right-click the data volume, and then click **Turn On BitLocker**. BitLocker then verifies that your server meets its requirements and then initializes the drive.


**NOTE** If BitLocker is already enabled, the **Manage BitLocker** option is displayed instead of **Turn On BitLocker**.

2. On the **Choose How You Want To Unlock This Drive** page, choose one or more of the following options, and then click **Next**:
  - **Use A Password To Unlock The Drive** Select this option if you want the user to be prompted for a password to unlock the drive. Passwords allow a drive to be unlocked in any location and to be shared with other people.
  - **Use My Smart Card To Unlock The Drive** Select this option if you want the user to use a smart card and enter the smart card PIN to unlock the drive. Because this feature requires a smart card reader, it's typically used to unlock a drive in the workplace and not for drives that might be used outside the workplace.

**NOTE** When you click **Next**, the wizard generates a recovery key. You can use the key to unlock the drive if BitLocker detects a condition that prevents it from unlocking the drive during boot. Note that you should save the key on removable media or on a network share. You can't store the key on the encrypted volume or the root directory of a fixed drive.







← BitLocker Drive Encryption (E:)

Choose how you want to unlock this drive

☒ Use a password to unlock the drive

Passwords should contain uppercase and lowercase letters, numbers, spaces, and symbols.

Enter your password

Reenter your password

☐ Use my smart card to unlock the drive

You'll need to insert your smart card. The smart card PIN will be required when you unlock the drive.

3. On the How Do You Want To Back Up Your Recovery Key? page, choose a save location for the recovery key, preferably a USB flash drive, other removable media or a network location.



← BitLocker Drive Encryption (E:)


How do you want to back up your recovery key?

If you forget your password or lose your smart card, you can use your recovery key to access your drive.

→ Save to a file

→ Print the recovery key

4. You can now print the recovery key if you want to, such as for storage in your office safe. When you have finished, click Next.
5. If allowed in Group Policy, you can elect to encrypt used disk space only or the entire drive, and then click Next. Encrypting the used disk space only is faster than encrypting an entire volume. It's also the recommended option for newer servers and drives (except in high-security environments).



← BitLocker Drive Encryption (E:)

Choose how much of your drive to encrypt

If you're setting up BitLocker on a new drive or a new PC, you only need to encrypt the part of the drive that's currently being used. BitLocker encrypts new data automatically as you add it.

If you're enabling BitLocker on a PC or drive that's already in use, consider encrypting the entire drive. Encrypting the entire drive ensures that all data is protected—even data that you deleted but that might still contain retrievable info.

☐ Encrypt used disk space only (faster and best for new PCs and drives)

☒ Encrypt entire drive (slower but best for PCs and drives already in use)


6. Choose the encryption mode. Windows Server 2016 and Windows 10 support standard AES encryption as well as enhanced XTS-AES encryption in 128-bit



and 256-bit configurations. As earlier releases of Windows do not support the new XTS-AES option, you should only use this option if you the encrypted drive won't be used with earlier releases of Windows or Windows Server.



7. On the Are You Ready To Encrypt This Drive? page, click Start Encrypting. How long the encryption process takes depends on the amount of data being encrypted and other factors.



As the encryption process can be paused and resumed, you can shut down the server before the drive is completely encrypted and the encryption of the drive will resume when you restart the server. The encryption state is maintained in the event of a power loss as well.

## Enabling BitLocker on USB Flash Drives


Encrypting USB flash drives protects the data stored on the volume. Any USB flash drive formatted with FAT, FAT32, exFAT, or NTFS can be encrypted with BitLocker. The length of time it takes to encrypt a drive depends on the size of the drive, the processing power of the server, and the level of activity on the server.

Before you enable BitLocker, you should configure the appropriate Removable Data Drives policies and settings in Group Policy and then wait for Group Policy to be refreshed. If you don't do this and you enable BitLocker, you might need to turn BitLocker off and then turn BitLocker back on because certain state and management flags are set when you turn on BitLocker.

To be sure that you can recover an encrypted volume, you should allow data-recovery agents and store recovery information in Active Directory. If you use a flash drive with earlier versions of Windows, the Allow Access To BitLocker-Protected Removable Data Drives From Earlier Versions Of Windows policy can ensure that you have access to the USB flash drive on other operating systems and servers. Unlocked drives are read-only.

To enable BitLocker encryption on a USB flash drive, follow these steps:

1. Insert the USB flash drive. In File Explorer, right-click the USB flash drive, and then click Turn On BitLocker. BitLocker then verifies that your server meets its requirements and then initializes the drive.
2. On the Choose How You Want To Unlock This Drive page, choose one or more of the following options, and then click Next:
  - **Use A Password To Unlock This Drive** Select this option if you want the user to be prompted for a password to unlock the drive. Passwords allow a drive to be unlocked in any location and to be shared with other people.
  - **Use My Smart Card To Unlock The Drive** Select this option if you want the user to use a smart card and enter the smart card PIN to unlock the drive. Because this feature requires a smart card reader, it's typically used to unlock a drive in the workplace and not for drives that might be used outside the workplace.



← BitLocker Drive Encryption (G:)

Choose how you want to unlock this drive

☒ Use a password to unlock the drive  
Passwords should contain uppercase and lowercase letters, numbers, spaces, and symbols.

Enter your password

Reenter your password

☐ Use my smart card to unlock the drive  
You'll need to insert your smart card. The smart card PIN will be required when you unlock the drive.

3. On the How Do You Want To Back Up Your Recovery Key? page, click Save The Recovery Key To A File.

4. In the Save BitLocker Recovery Key As dialog box, choose a save location, and then click Save. As a best practice, you should save the recovery key to a network location or USB drive rather than the server to which the key belongs.
5. You can now print the recovery key if you want to, such as for storage in your office safe. When you have finished, click Next.
6. If allowed in Group Policy, you can elect to encrypt used disk space only or the entire drive, and then click Next. Encrypting the used disk space only is faster than encrypting an entire volume. It's also the recommended option for newer servers and drives (except in high-security environments).
7. Choose the encryption mode. Windows Server 2016 and Windows 10 support standard AES encryption as well as enhanced XTS-AES encryption in 128-bit and 256-bit configurations. As earlier releases of Windows do not support the new XTS-AES option, you should only use this option if you the encrypted drive won't be used with earlier releases of Windows or Windows Server.
8. On the Are You Ready To Encrypt This Drive? page, click Start Encrypting. Be sure to pause encryption before removing the drive and then resume to complete the encryption. Do not otherwise remove the USB flash drive until the encryption process is complete. How long the encryption process takes depends on the amount of data to encrypt and other factors.




The encryption process does the following:

1. It adds an Autorun.inf file, the BitLocker To Go reader, and a Read Me.txt file to the USB flash drive.
2. It creates a virtual volume with the encrypted contents of the drive.


3. It encrypts the virtual volume to protect it. USB flash drive encryption takes approximately 6 to 10 minutes per gigabyte to complete. The encryption process can be paused and resumed, as long as you don't remove the drive.

When you insert an encrypted drive into a USB slot on a server running Windows Server 2016, a notification on the secure desktop is displayed. If the notification disappears before you can click it, right-click the drive in File Explorer and then select Unlock Drive.



**Unlock drive H:**  
This drive is BitLocker-protected.

Click the notification to display the BitLocker dialog box. This dialog box is also displayed on the secure desktop.



BitLocker (H:)

Enter password to unlock this drive.

**Fewer options**

Enter recovery key

☐ Automatically unlock on this PC

**Unlock**

When you are prompted, enter the password. Optionally, click More Options to expand the dialog box so that you select Automatically Unlock On This Server to save the password in an encrypted file on the server's system volume. Finally, click Unlock to unlock the drive so that you can use it.

If you forget or lose the password for the drive but have the recovery key, click More Options, and then click Enter Recovery Key. Enter the 48-digit recovery key, and then click Unlock. This key is stored in the XML-formatted recovery key file as plain text.


# Enabling BitLocker on System Volumes

Before you can encrypt a system volume, you should remove all bootable media from a server's CD/DVD drives, in addition to all USB flash drives. You should also back up critical files and data. You can then enable BitLocker encryption on the system volume by completing the following steps:


1. In File Explorer, right-click the system volume, and then click Turn On BitLocker. Windows checks the server and the drive to ensure that BitLocker can be enabled. If so, you'll see a prompt warning you that turning on BitLocker deletes any previously created restore points for the system volume. Click Yes to continue.



2. BitLocker Drive Encryption will then initialize the system volume, perform preliminary checks of the volume and also ensure the server meets requirements for using the feature. The system volume must be formatted with NTFS. When the preliminary checks are completed, click Next.



3. BitLocker warns you that a new system drive will be created. Click Next. As part of preparing the drive, BitLocker must create a new, unencrypted partition for the files required to start the operating system. This establishes the new system drive, which does not have a drive letter. The original partition containing the other operating system files and your data will then be encrypted.



**NOTE** If BitLocker is already enabled, the Manage BitLocker option is displayed instead of Turn On BitLocker. As part of the setup, Windows prepares the required BitLocker partition, if necessary. If Windows RE is in this partition, Windows moves Windows RE to the system volume and then uses this additional partition for BitLocker.

**NOTE** If the server doesn't have a TPM, the Allow BitLocker Without A Compatible TPM option must be enabled for operating system volumes in the Require Additional Authentication At Startup policy.

4. Once the first phase of drive preparation is complete, you'll be prompted to restart the server. If there is a problem moving the Windows RE to the system volume, you'll be notified of this and should consider creating a recovery drive containing a backup of the system files once the server restarts.
5. BitLocker setup resumes automatically. When you are ready to continue, click Next. You can now configure BitLocker startup preferences. Continue as discussed in the separate procedures that follow. If the server doesn't have a TPM, your options will be different. You'll be able to create a password to unlock the drive, or you can insert a USB flash drive and store the startup key on the flash drive.





When a server has a TPM, you can use BitLocker to provide basic integrity checks of the volume without requiring any additional keys. In this configuration, BitLocker protects the system volume by encrypting it. This configuration does the following:

- Grants access to the volume to users who can log on to the operating system
- Prevents those who have physical access to the server from booting to an alternative operating system to gain access to the data on the volume
- Allows the server to be used with or without a TPM for additional boot security
- Does not require a password or a smart card with a PIN

To use BitLocker without any additional keys, follow these steps:

1. On the Choose How To Unlock Your Drive At Startup page, click Let BitLocker Automatically Unlock My Drive.
2. On the How Do You Want To Back Up Your Recovery Key? page, click Save To A File.
3. In the Save BitLocker Recovery Key As dialog box, choose the location of your USB flash drive or an appropriate network share, and then click Save. Do not use a USB flash drive that is BitLocker-encrypted.
4. You can now optionally save the recovery key to another location, print the recovery key, or both. Click an option, and then follow the wizard steps to set the location for saving or printing the recovery key. When you have finished, click Next.
5. If allowed in Group Policy, you can elect to encrypt used disk space only or the entire drive, and then click Next. Encrypting the used disk space only is faster than encrypting an entire volume. It's also the recommended option for newer servers and drives (except in high-security environments).
6. On the Encrypt The Drive page, click Start Encrypting. How long the encryption process takes depends on the amount of data to encrypt and other factors.

To enhance security, you can require additional authentication at startup. This configuration does the following:

- Grants access to the volume only to users who can provide a valid key
- Prevents those who have physical access to the server from booting to an alternative operating system to gain access to the data on the volume
- Allows the server to be used with or without a TPM for additional boot security
- Requires a password or a smart card with a PIN
- 



Optionally uses network unlock to unlock the volume when the server joined to and connected to the domain

You can enable BitLocker encryption for use with a startup key by following these steps:

1. Insert a USB flash drive in the server (if one is not already there). Do not use a USB flash drive that is BitLocker-encrypted.
2. On the Choose How To Unlock Your Drive At Startup page, click the Insert A USB Flash Drive option.
3. On the Back Up Your Startup Key page, click the USB flash drive, and then click Save.
4. Next, you need to save the recovery key. Because you should not store the recovery key and the startup key on the same medium, remove the USB flash drive and insert a second USB flash drive.

**NOTE** The startup key is different from the recovery key. If you create a startup key, this key is required to start the server. The recovery key is required to unlock the server if BitLocker enters Recovery mode, which might happen if BitLocker suspects the server has been tampered with while the server was offline.

5. On the How Do You Want To Back Up Your Recovery Key? page, click Save To A File.
6. In the Save BitLocker Recovery Key As dialog box, choose the location of your USB flash drive, and then click Save. Do not remove the USB drive with the recovery key.
7. You can now optionally save the recovery key to a network folder, print the recovery key, or both. Click an option, and then follow the wizard's steps to set the location for saving or printing the recovery key. When you have finished, click Next.
8. If allowed in Group Policy, you can elect to encrypt used disk space only or the entire drive, and then click Next. Encrypting the used disk space only is faster than encrypting an entire volume. It's also the recommended option for newer servers and drives (except in high-security environments).
9. On the Encrypt The Volume page, confirm that Run BitLocker System Check is selected, and then click Continue.
10. Confirm that you want to restart the server by clicking Restart Now. The server restarts, and BitLocker ensures that the server is BitLocker-compatible and ready for encryption. If the server is not ready for encryption, you will get an





error and need to resolve the error status before you can complete this procedure. If the server is ready for encryption, the Encryption In Progress status bar is displayed. You can monitor the status of the disk volume encryption by pointing to the BitLocker Drive Encryption icon in the notification area. By double-clicking this icon, you can open the Encrypting dialog box and monitor the encryption process more closely. You also have the option to pause the encryption process. Volume encryption takes approximately 1 minute per gigabyte to complete.

By completing this procedure, you have encrypted the operating system volume and created a recovery key unique to that volume. The next time you turn on your server, either the USB flash drive with the startup key must be plugged into a USB port on the server or the server must be connected to the domain network and using network unlock. If the USB flash drive is required for startup and you do not have the USB flash drive containing your startup key, you will need to use Recovery mode and supply the recovery key to gain access to the data.

You can enable BitLocker encryption for use with a startup PIN by following these steps:

1. On the Choose How To Unlock Your Drive At Startup page, select the Enter A PIN option.
2. On the Enter A PIN page, enter and confirm the PIN. The PIN can be any number you choose and must be 4 to 20 digits in length. The PIN is stored on the server.
3. Insert a USB flash drive on which you want to save the recovery key, and then click Set PIN. Do not use a USB flash drive that is BitLocker-encrypted.
4. Continue with steps 5-9 in the previous procedure.

When the encryption process is complete, you have encrypted the entire volume and created a recovery key unique to this volume. If you created a PIN or a startup key, you are required to use the PIN or startup key to start the server (or the server must be connected to the domain network and using network unlock). Otherwise, you will observe no change to the server unless the TPM changes or cannot be accessed, or if someone tries to modify the disk while the operating system is offline. In this case, the server enters Recovery mode, and you need to enter the recovery key to unlock the server.



## Managing and Troubleshooting BitLocker

You can determine whether a system volume, data volume, or inserted removable drive uses BitLocker by clicking System And Security in Control Panel, and then doubleclicking BitLocker Drive Encryption. You'll find the status of BitLocker on each volume.

The BitLocker Drive Encryption service must be started for BitLocker to work properly. Usually, this service is configured for manual startup and runs under the LocalSystem account.

To use smart cards with BitLocker, the Smart Card service must be started. Typically, this service is configured for manual startup and runs under the LocalService account.

After you create a startup key or PIN and a recovery key for a server, you can create duplicates of the startup key, startup PIN, or recovery key as necessary for backup or replacement purposes by using the options on the BitLocker Drive Encryption page in Control Panel.

**NOTE** With fixed data and operating system drives, another way to access this page is to right-click the volume in File Explorer, and then click Manage BitLocker. If BitLocker is turned off, the Turn On BitLocker option is displayed instead.

The management options provided depend on the type of volume with which you are working and the encryption settings you choose. The available options include the following:

- **Back Up Recovery Key** Enables you to save or print the recovery key. Click this option, and then follow the prompts.
  - **Change Password** Allows you to change the encryption password. Click this option, enter the old password, and then enter and confirm the new password. Click Change Password.
  - **Remove Password** Click this option to remove the encryption password requirement for unlocking the drive. You can do this only if another unlocking method is configured first.
  - **Add Smart Card** Enables you to add a smart card for unlocking the drive. Click this option, and then follow the prompts.
  - **Remove Smart Card** Click this option to remove the smart card requirement for unlocking the drive.
  - **Change Smart Card** Enables you to change the smart card used to unlock the drive. Click this option, and then follow the prompts.
- 

- **Turn On Auto-Unlock** Click this option to turn on automatic unlocking of the drive.
- **Turn Off Auto-Unlock** Click this option to turn off automatic unlocking of the drive.
- **Turn Off BitLocker** Click this option to turn off BitLocker and decrypt the drive.

## Recovering Data Protected by BitLocker Drive Encryption

If you've configured BitLocker Drive Encryption and the server enters Recovery mode, you need to unlock the server. To unlock the server by using a recovery key stored on a USB flash drive, follow these steps:

1. Turn on the server. If the server is locked, the server opens the BitLocker Drive Encryption Recovery console.
2. When you are prompted, insert the USB flash drive that contains the recovery key, and then press Enter.
3. The server will unlock and reboot automatically. You do not need to enter the recovery key manually.

If you saved the recovery key file in a folder on another server or on removable media, you can use another server to open and validate the recovery key file. To locate the correct file, find Password ID on the recovery console displayed on the locked server and write down this number. The file containing the recovery key uses this Password ID as the file name. Open the file and locate the recovery key.

To unlock the server by typing the recovery key, follow these steps:

1. Turn on the server. If the server is locked, the BitLocker Drive Encryption Recovery console opens.
2. Type the recovery key, and then press Enter. The server will unlock and reboot automatically.

A server can become locked if a user tries to enter the recovery key but is repeatedly unsuccessful. In the recovery console, you can press Esc twice to exit the recovery prompt and turn off the server. A server might also become locked if an error related to TPM occurs or boot data is modified. In this case, the server halts very early in the boot process, before the operating system starts. At this point, the locked server might not be able to accept standard keyboard numbers. If that is the case, you must use the



function keys to enter the recovery password. Here, the function keys F1–F9 represent the digits 1 through 9, and the F10 function key represents 0.

## Disabling or Turning Off BitLocker Drive Encryption

When you need to make changes to TPM or make other changes to the system, you might first need to temporarily turn off BitLocker encryption on the system volume. You cannot temporarily turn off BitLocker encryption on data volumes; you can only decrypt data volumes.

To temporarily turn off BitLocker encryption on the system volume, follow these steps:

1. In Control Panel, click System And Security, and then double-click BitLocker Drive Encryption.
2. For the system volume, click Turn Off BitLocker Drive Encryption.
3. When prompted, click Disable BitLocker Drive Encryption.

By completing this procedure, you have temporarily disabled BitLocker on the operating system volume.

To turn off BitLocker Drive Encryption and decrypt a data volume or USB drive, follow these steps:

1. In Control Panel, click System And Security, and then double-click BitLocker Drive Encryption.
2. For the appropriate volume, click Turn Off BitLocker Drive Encryption.
3. When prompted to confirm, click Turn Off BitLocker Drive Encryption.

## Chapter 22. Using Storage Spaces

Storage Spaces provide the primary means for protecting critical data. Storage management and the ways that Windows Server works with disks have changed substantially since the introduction of Storage Spaces. Although traditional storage management techniques relate to physical drives located inside the server, many servers today use attached storage and virtual disks. Generally, when you work with internal fixed drives, you often need to perform advanced disk setup procedures, such as creating a volume set or setting up a redundant array of independent disks (RAID) array. Here, you create volumes or arrays that can span multiple drives and you know the exact physical layout of those drives.





On the other hand, when you work with attached storage, you might not know which actual physical disk or disks the volume you are working with resides on. Instead, you are presented with a virtual disk, also referred to as a logical unit number (LUN), which is a logical reference to a portion of the storage subsystem. Although the virtual disk can reside on one or more physical disks (spindles), the layout of the physical disks is controlled separately from the operating system (by the storage subsystem).

Standards-based, storage management approaches rely on the storage subsystem to manage the physical disk architecture and can be used with a server's internal, external or attached disks. When internal disks are used in this way, the internal disks—such as virtual disks on attached storage—are resources to be allocated by using standardsbased approaches. This means you can create virtual disk volumes on the physical disks, add the physical disks to storage pools, and create Internet SCSI (iSCSI) virtual disks that can be targeted. You can also enable data deduplication on your virtual disks.

Standards-based storage management focuses on the storage volumes themselves rather than the underlying physical layout, relying on hardware to handle the architecture particulars for data redundancy and the portions of disks that are presented as usable disks. This means the layout of the physical disks is controlled by the storage subsystem instead of by the operating system.

## Getting Started with Storage Spaces

With Storage Spaces and other standards-based management approaches, the physical layout of disks (spindles) is abstracted, so a “disk” can be a logical reference to a portion of a storage subsystem (a virtual disk) or an actual physical disk. This means a disk simply becomes a unit of storage and volumes can be created to allocate space within disks for file systems.

Taking this concept a few steps further, you can pool available space on disks so that units of storage (virtual disks) can be allocated from this pool on an as-needed basis. These units of storage, in turn, are apportioned with volumes to allocate space and create usable file systems.

Technically, the pooled storage is referred to as a storage pool and the virtual disks created within the pool are referred to as storages spaces . Given a set of “disks,” you can create a single storage pool by allocating all the disks to the pool or create multiple storage pools by allocating disks separately to each pool.

**REAL WORLD** When you throw storage subsystems into the mix, it's really a threelayered architecture. In Layer 1, the layout of the physical disks is





controlled by the storage subsystem. The storage system likely will use some form of RAID to ensure that data is redundant and recoverable in case of failure. In Layer 2, the virtual disks created by the arrays are made available to servers. The servers simply see the disks as storage that can be allocated. Windows Server can apply software-level RAID or other redundancy approaches to help protect against failure. In Layer 3, the server creates volumes on the virtual disks, and these volumes provide the usable file systems for file and data storage.


## Grouping Disks into Storage Pools

When you are working with File And Storage Services, you can group available physical disks into storage pools so that you can create virtual disks from available capacity. Each virtual disk you create is a storage space. Storage Spaces are made available through the Storage Services role service, which is automatically installed on every server running Windows Server 2016.

To integrate Storage Spaces with standards-based storage management frameworks, you'll want to add the Windows Standards-Based Storage Management feature to your file servers. When a server is configured with the File Services And Storage role, the Windows Standards-Based Storage Management feature adds components and updates Server Manager with the options for working with standards-based volumes. You might also want to do the following:

- [Add the Data Deduplication role service if you want to enable data deduplication.](#)  
Data deduplication works with volumes up to 64 TB and file sizes up to 1 TB. •
- [Add the iSCSI Target Server and iSCSI Target Storage Provider role services if you want the server to host iSCSI virtual disks.](#)

After you configure your servers as appropriate for your environment, you can select the File And Storage Services node in Server Manager to work with your storage volumes, and additional options will be available as well. The Servers subnode lists file servers that have been configured for standards-based management.




The Volumes subnode lists allocated storage on each server according to how volumes are provisioned and how much free space is available. Volumes are listed regardless of whether the underlying disks are physical or virtual. Right-click a volume to display management options, including the following:

- **Configure Data Deduplication** Allows you to enable and configure data deduplication for NTFS volumes. If the feature is enabled, you can then use this option to disable data deduplication as well.
- **Delete Volume** Allows you to delete the volume. The space that was used is then marked as unallocated on the related disk.
- **Extend Volume** Allows you to extend the volume to unallocated space of the related disk.
- **Format** Allows you to create a new file system on the volume that overwrites the existing volume.
- **Manage Drive Letter And Access Paths** Allows you to change the drive letter or access path associated with the volume.
- **New iSCSI Virtual Disk** Allows you to create a new iSCSI virtual disk that is stored on the volume.
- **New Share** Allows you to create new Server Message Block (SMB) or Network File System (NFS) shares on the volume.
- **Properties** Displays information about the volume type, file system, health, capacity, used space, and free space. You can also use this option to set the volume label.




- **Repair File System Errors** Allows you to repair errors detected during an online scan of the file system.
- **Scan File System For Errors** Allows you to perform an online scan of the file system. Although Windows attempts to repair any errors that are found, some errors can be corrected only by using a repair procedure.




The Disks subnode lists the disks available to each server according to total capacity, unallocated space, partition style, subsystem, and bus type. Server Manager attempts to differentiate between physical disks and virtual disks by showing the virtual disk label (if one was provided) and the originating storage subsystem. Right-click a disk to display management options, including the following:

- **Bring Online** Enables you to take an offline disk and make it available for use.
- **Take Offline** Enables you to take a disk offline so that it can no longer be used.
- **Reset Disk** Enables you to completely reset the disk, which deletes all volumes on the disk and removes all available data on the disk.
- **New Volume** Enables you to create a new volume on the disk.





In Server Manager, you can work with storage pools and allocate space by selecting the File And Storage Services node, and then selecting the related Storage Pools subnode. The Storage Pools subnode lists the available storage pools, the virtual disks created within storage pools, and the available physical disks. Keep in mind that what's presented as physical disks might actually be LUNs (virtual disks) from a storage subsystem.




# Storage Pool Options, Versions and Features

Working with storage pools is a multistep process:

1. You create storage pools to pool available space on one or more disks.
2. You allocate space from this pool to create one or more virtual disks.
3. You create one or more volumes on each virtual disk to allocate storage for file systems.

The version of Storage Spaces that ships with Windows Server 2016 is different from the version that ships with earlier releases of Windows Server. If you have any questions about which version of Storage Spaces is being used, complete the following steps to check the version:

1. In Server Manager, select the File And Storage Services node and then select the Storage Pools node.
2. Right-click the storage pool you want to examine and then select Properties.
3. In the Properties dialog box, the storage pool version is listed on the General tab.




You also can use this technique to check capacity, total used space and available space. By default, Storage Space alerts you when storage is approaching capacity and when a storage space reaches 70 percent of the total provisioned size. When you get such an alert, you should consider allocating additional storage.

**NOTE** The Health tab provides details about the operational status of the storage pool. If you are using a storage pool from an earlier version of Windows Server, you'll see a warning status and the storage pool will be listed as having a degraded operational status. This status simply means the storage pool cannot

use any of the new functionality of the current Storage Spaces version. While in a degraded state, the storage pool continues to operate normally and can be used with servers running the specified version of Windows Server as well as Windows Server 2016.

You can upgrade the version of Storage Spaces being used by a storage pool by rightclicking a storage pool, and then selecting Upgrade Storage Pool Version. As the warning prompt states, the version upgrade cannot be reversed once completed.



The Windows Server 2016 version of Storage Spaces does the following:


- Supports using tiered storage when you combine hard disk drives and solid state drives in a storage pool. With storage tiers, the most frequently used files are moved from hard disk storage to solid state storage to improve performance.
- Supports using local storage and removes the need for shared SAS fabric. This allows a server to use the network as the storage fabric and take advantage of the improvements provided with SMB3 and SMB Direct. Microsoft calls this feature Storage Spaces Direct.
- Supports block-level synchronous or asynchronous replication between servers using SMB3 as the transport. Microsoft calls this feature Storage Replica.
- Supports storage spaces that have dual parity, and parity and dual parity spaces on failover clusters. With dual parity, storage spaces are protected against two simultaneous drive failures.
- Supports automatic rebuild of storage spaces from storage pool free space instead of having to use hot spares to recover from drive failures. Here, instead of writing a copy of data that was on a failed disk to a hot spare, the parity or mirrored data is copied to multiple drives in the pool to achieve the previous level of resiliency automatically. As a result, you don't need to specifically allocate hot spares in storage pools, provided that a sufficient number of drives is assigned to the pool to allow for automatic resiliency recovery.



- Supports storage tiers to automatically move frequently used files from slower physical disks to faster Solid State Drive (SSD) storage. This feature is applicable only when a storage space has a combination of SSD storage and hard disk drive (HDD) storage. Additionally, the storage type must be set as fixed, the volumes created on virtual disks must be the same size as the virtual disk, and enough free space must be available to accommodate the preference. For fine-grained management, use the `Set-FileStorageTier` cmdlet to assign files to standard physical drive storage or faster SSD storage.
- Supports write-back caching when a storage pool uses SSD storage. Write-back cache buffers small random writes to SSD storage before later writing the data to HDD storage. Buffering writes in this way helps to protect against data loss in the event of power failures. For write-back cache to work properly, storage spaces with simple volumes must have at least one SSD, storage spaces with two-way mirroring or single-parity must have at least two SSDs, and storage spaces with three-way mirroring or dual parity must have at least three SSDs. When these requirements are met, the volumes automatically will use a 1-GB write-back cache by default. You can designate SSDs that should be used for write-back caching by setting the usage as Journal (the default in this configuration). If enough SSDs are not configured for journaling, the write-back cache size is set to 0 (meaning write-back caching will not be used). The only exception is for parity spaces, which will then have the write-back cache size set to 32 MB.

If you have any question about the size of the write-back cache, complete the following steps to check the cache size:

1. Right-click the virtual disk you want to examine, and then select Properties.
2. In the Properties dialog box, select Details in the left pane, and then choose WriteCacheSize on the Property selection list.



You also can use this technique to check allocated size, status, provisioned size, provision type, redundancy type, and more.





# Creating Storage Pools

Storage pools allow you to pool available space on disks so that units of storage (virtual disks) can be allocated from this pool. To create a storage pool, you must have at least one unused disk and a storage subsystem to manage it. This storage subsystem can be the one included with the Storage Spaces feature or a subsystem associated with attached storage.



# Using Drives and Volumes in a Storage Pool

When a computer has extra drives in addition to the drive on which Windows is installed, you can allocate one or more of the additional drives to a storage pool. However, keep in mind that if you use a formatted drive with a storage pool, Windows permanently deletes all the files on that drive. Additionally, it's important to point out that physical disks with the MBR partition style are converted automatically to the GPT partition style when you add them to a storage pool and create volumes on them.

Each physical disk allocated to the pool can be handled in one of three ways:

- As a data store that is available for use
- As a data store that can be manually allocated for use
- As a hot spare in case a disk in the pool fails or is removed from the subsystem

Types of volumes you can create are as follows:

- **Simple Volumes** Creates a simple volume by writing one copy of your data to one or more drives. With simple volumes, there is no redundancy and no associated overhead. As an example, you can create a single volume that spans two 2-TB drives, making 4 TB of storage available. However, because there is no resiliency, a failure of any drive in a simple volume causes the entire volume to fail.
- **Two-way mirrors** Creates a mirrored set by writing two copies of a computer's data, helping to protect against a single drive failure. Two-way mirrors require at least two drives. With two-way mirrors, there is a 1/2 (50 percent) overhead for redundancy with two drives. As an example, you could allocate two 2-TB drives as a two-way mirror, giving you 2 TB of mirrored storage.
- **Parity volumes** Creates a volume that uses disk striping with parity, helping to provide fault tolerance with less overhead than mirroring. Parity volumes require at least three drives. With parity volumes there is a 1/3 (33.33 percent) overhead for redundancy with three drives. As an example, you could allocate three 2-TB drives as a parity volume, giving you 4 TB of protected storage.
- **Dual parity volumes** Creates a volume that uses disk striping with two sets of parity data, helping to protect against two simultaneous drive failures while requiring less overhead than three-way mirroring. Dual parity volumes require at least seven drives.
- **Three-way mirrors** Creates a mirrored set by writing three copies of a computer's data and by using disk striping with mirroring, helping to protect against two simultaneous drive failures. Although three-way mirrors do not have a penalty for read operations, they do have a performance penalty for write operations because of the overhead associated with having to write data to three separate disks. This





overhead can be mitigated by using multiple drive controllers. Ideally, you'll want to ensure that at least three drive controllers are used. Three-way mirrors require at least five drives.

If you are familiar with RAID 6, you are familiar with dual parity volumes. Although dual parity does not have a penalty for read operations, it does have a performance penalty for write operations because of the overhead associated with calculating and writing dual parity values. With standard dual parity volumes, the usable capacity of dual parity volumes is calculated as the sum of the number of volumes minus two times the size of the smallest volume in the array or  $(N - 2) \times \text{MinimumVolumeSize}$ . For example, with 7 volumes and the smallest volume size of 1TB, the usable capacity typically is 5 TB [calculated as  $(7 - 2) \times 1 \text{ TB} = 5 \text{ TB}$ ].

Although logically it would seem that you need at least six drives to have three mirrored copies of data, mathematically you need only five. Why? If you want three copies of your data, you need at least 15 logical units of storage to create those three copies. Divide 15 by 3 to come up with the number of disks required, and the answer is that you need 5 disks. Thus, Storage Spaces uses 1/3 of each disk to store original data and 2/3 of each disk to store copies of data. Following this, a three-way mirror with five volumes has a 2/3 (66.66 percent) overhead for redundancy. Or put another way, you could allocate five 3-TB drives as a three-way mirror, giving you 5 TB of mirrored storage (and 10 TB of overhead).

With single parity volumes, data is written out horizontally with parity calculated for each row of data. Dual parity differs from single parity in that row data is not only stored horizontally, it is also stored diagonally. If a single disk fails or a read error from a bad bit or block error occurs, the data is re-created by using only the horizontal row parity data (just as in single parity volumes). In the case of a multiple drive issue, the horizontal and diagonal row data are used for recovery.

To understand how dual parity typically works, consider the following simplified example:

Each horizontal row of data has a parity value, the sum of which is stored on the parity disk for that row (and calculated by using an exclusive OR). Each horizontal parity stripe misses one and only one disk. If the parity value is 2, 3, 1, and 4 on disks 0, 1, 2, and 3 respectively, the parity sum stored on the disk 4 (the parity disk for this row) is 10 ( $2 + 3 + 1 + 4 = 10$ ). If disk 0 were to have a problem, the parity value for the row on this disk could be restored by subtracting the remaining horizontal values from the horizontal parity sum ( $10 - 3 - 1 - 4 = 2$ ).





The second set of parity data is written diagonally (meaning in different data rows on different disks). Each diagonal row of data has a diagonal parity value, the sum of which is stored on the diagonal parity disk for that row (and calculated by using an exclusive OR). Each diagonal parity stripe misses two disks: one disk in which the diagonal parity sum is stored and one disk that is omitted from the diagonal parity striping. Additionally, the diagonal parity sum includes a data row from the horizontal row parity as part of its diagonal parity sum.

If the diagonal parity value is 1, 4, 3, and 7 on disks 1, 2, 3, and 4 respectively (with four associated horizontal rows), the diagonal parity sum stored on disk 5 (the diagonal parity disk for this row) is 15 ( $4 + 1 + 3 + 7 = 15$ ) and the omitted disk is disk 0. If disk 2 and disk 4 were to have a problem, the diagonal parity value for the row can be used to restore both of the lost values. The missing diagonal value is restored first by subtracting the remaining diagonal values from the diagonal parity sum. The missing horizontal value is restored next by subtracting the remaining horizontal values for the subject row from the horizontal parity sum for that row.

**NOTE** Keep in mind that dual parity as implemented in Storage Spaces uses 7 disks and the previous example was simplified. Although parity striping with 7 disks works differently than discussed in this example, the basic approach uses horizontal and vertical stripes.






# Creating a Storage Pool


You can create a storage pool by completing the following steps:

1. In Server Manager, select the File And Storage Services node, and then select the related Storage Pools subnode.
2. Select Tasks in the Storage Pools panel, and then select New Storage Pool. This starts the New Storage Pool Wizard. If the wizard displays the Before You Begin page, read the Welcome text, and then select Next.
3. On the Specify A Storage Pool Name And Subsystem page, enter a name and description of the storage pool, and then select the primordial pool with which you want to work. (A primordial pool is simply a group of disks managed by and available to a specific server via a storage subsystem.) Select Next.



**TIP** Select the primordial pool for the server you want to associate the pool with and allocate storage for. For example, if you are configuring storage for CorpServer38, select the primordial pool that is available to CorpServer38.

4. On the Select Physical Disks For The Storage Pool page, select the unused physical disks that should be part of the storage pool, and then specify the type of allocation for each disk. A storage pool must have more than one disk to use the mirroring and parity features available to protect data in case of error or failure. When setting the Allocation value, keep the following in mind:
- Choose Automatic to allocate the disk to the pool and make it available for use as needed.
  - Choose Manual to allocate the disk to the pool but not allow it to be used until it is manually allocated.
  - Choose Hot Spare to allocate the disk to the pool as a spare disk that is made available for use if another disk in the pool fails or is removed from the subsystem.



**TIP** Storage tiers allow you to configure Storage Pools that automatically move frequently used files to the fastest storage available. Generally, this means files are moved from hard disk drives to solid state drives as appropriate. Tiered storage is only available when a storage pool has a mix of hard disk drives and solid state drives. Specifically, the storage pool must have at least one automatically added physical disk of each media type.

- When you are ready to continue, select Next. After you confirm your selections, select Create. The wizard tracks the progress of the pool creation. When the wizard finishes creating the pool, the View Results page will be updated to reflect this. Review the details to ensure that all phases were completed successfully, and then select Close.



If any portion of the configuration failed, note the reason for the failure and take corrective actions as appropriate before repeating this procedure.

- If one of the physical disks is currently formatted with a volume, you'll get the following error:

Could not create storage pool. One of the physical disks specified is not supported by this operation.

This error occurs because physical disks that you want to add to a storage pool cannot contain existing volumes. To resolve the problem, you'll need to repeat the procedure and select a different physical disk, or remove any existing volumes on the physical disk, repeat the procedure, and then select the disk



again. Keep in mind that deleting a volume permanently erases all data it contains.

- If one of the physical disks is unavailable after being selected, you'll see the error:

Could not create storage pool. One or more parameter values passed to the method were invalid.

This error occurs because a physical disk that was available when you started the New Storage Pool Wizard has become unavailable or is offline. To resolve the problem, you'll need to a) repeat the procedure and select a different physical disk, or b) bring the physical disk online or otherwise make it available for use, repeat the procedure, and then select the disk again.

**NOTE** External storage can become unavailable for a variety of reasons. For example, an external connected cable might have been disconnected or a LUN previously allocated to the server might have been reallocated by a storage administrator.



# Creating Virtual Disks and Volumes

After you create a storage pool, you can allocate space from the pool to virtual disks that are available to your servers. You can then create volumes on the disks to enable data storage.



# Creating a Virtual Disk in a Storage Space

Each physical disk allocated to the pool can be handled in one of three ways:


- As a data store that is available for use
- As a data store that can be manually allocated for use
- As a hot spare in case a disk in the pool fails or is removed from the subsystem

When a storage pool has a single disk, your only option for allocating space on that disk is to create virtual disks with a simple layout. A simple layout does not protect against disk failure. If a storage pool has multiple disks, you have these additional layout options:

- **Mirror** With a mirror layout, data is duplicated on disks by using a mirroring technique similar to what I discussed previously in this chapter. However, the mirroring technique is more sophisticated in that data is mirrored onto two or three disks at a time. Like standard mirroring, this approach has its advantages and disadvantages. If a storage space has two or three disks, you are fully protected against a single disk failure, and if a storage space has five or more disks, you are fully protected against two simultaneous disk failures. The disadvantage is that mirroring reduces capacity by up to 50 percent. For example, if you mirror two 1-TB disks, the usable space is 1 TB.
- **Parity** With a parity layout, data and parity information are striped across physical disks by using a striping-with-parity technique similar to what I discussed previously in this chapter. Like standard striping with parity, this approach has its advantages and disadvantages. You need at least three disks to fully protect yourself against a single disk failure. You lose some capacity to the striping, but not as much as with mirroring.

You can create a virtual disk in a storage pool by completing the following steps:

1. In Server Manager, select the File And Storage Services node, and then select the related Storage Pools subnode.
2. Select Tasks in the Virtual Disks panel, and then select New Virtual Disk. This starts the New Virtual Disk Wizard. If the wizard displays the Before You Begin page, read the Welcome text, and then select Next.
3. On the Storage Pool page, select the storage pool in which you want to create the virtual disk, and then select Next. Each available storage pool is listed according to the server it is managed by and available to. Make sure the pool has enough free space to create the virtual disk.



**TIP** Select the storage pool for the server you want to associate the virtual disk with and allocate storage from. For example, if you are configuring storage for CorpServer38, select a storage pool that is available to CorpServer38.


- On the Specify The Virtual Disk Name page, enter a name and description for the virtual disk. If you are using a combination of SSD storage and HDD storage, use the check box provided to specify whether you want to create storage tiers. With storage tiers, the most frequently accessed files are automatically moved from slower HDD to faster SSD storage. This option is not applicable when the server has only HDD or SSD storage. To continue, select Next.



- If you allocated storage from 3 or more different enclosures using the automatic option, you can enable enclosure awareness on the Specify Enclosure Awareness page. With enclosure awareness, copies of your data are stored on separate enclosures to help protect your data if an entire enclosure fails.



6. On the Select The Storage Layout page, select the storage layout as appropriate for your reliability and redundancy requirements. The simple layout is the only option for storage pools that contain a single disk. If the underlying storage pool has multiple disks, you can choose a simple layout, a mirror layout, or a parity layout. Select Next.



**REAL WORLD** If there aren't enough available disks to implement the storage layout, you'll get the error: The storage pool does not contain enough physical disks to support the selected storage layout. Select a different layout or repeat this procedure and select a different storage pool to work with initially.

Keep in mind the storage pool might have one or more disks allocated as hot spares. Hot spares are made available automatically to recover from disk failure when you use mirroring or parity volumes—and cannot otherwise be used. To force Windows to use a hot spare, you can remove the hot spare from the storage pool by right-clicking it and selecting Remove, and then adding the drive back to the storage pool as an automatically allocated disk by right-clicking the storage pool and selecting Add Physical Drive. Unfortunately, doing so might cause a storage pool created with a hot spare to report that it is in an Unhealthy state. If you subsequently try to add the drive again in any capacity, you'll get an error stating “Error adding task: The




storage pool could not complete the operation because its configuration is read-only.” The storage pool is not, in fact, in a read-only state. If the storage pool were in a read-only state you could enter the following command at an administrator PowerShell prompt to clear this state:

```
Get-Storagepool "PoolName" | Set-Storagepool -IsReadonly $false
```


However, entering this command likely will not resolve the problem. To clear this error, I needed to reset Storage Spaces and the related subsystem. You might find it easier to simply restart the server. After you reset or restart the server, the storage pool will transition from an error state (where a red circle with an ‘x’ is showing) to a warning state (where a yellow triangle with an ‘!’ is showing). You can then remove the physical disk from the storage pool by right-clicking it and selecting Remove. Afterward, you will be able to add the physical disk as an automatically-allocated disk by right-clicking the storage pool and selecting Add Physical Drive.

7. On the Specify The Provisioning Type page, select the provisioning type. Storage can be provisioned in a thin disk or a fixed disk. With thin-disk provisioning, the volume uses space from the storage pool as needed, up to the volume size. With fixed provisioning, the volume has a fixed size and uses space from the storage pool equal to the volume size. Select Next.




8. On the Specify The Size Of The Virtual Disk page, use the options provided to set the size of the virtual disk. With fixed provisioning, selecting Maximum Size ensures that the disk is created and sized with the maximum space possible given the available space. For example, if you use a 2-TB disk and a 1.5-TB disk with a mirrored layout, a 1.5-TB fixed disk will be created as this is the maximum mirrored size possible.





9. When you are ready to continue, select Next. After you confirm your selections, select Create. The wizard tracks the progress of the disk creation. When the wizard finishes creating the disk, the View Results page will be updated to reflect this. Review the details to ensure that all phases were completed successfully. If any portion of the configuration failed, note the reason for the failure and take corrective actions as appropriate before repeating this procedure.



10. When you select Close, the New Volume Wizard should start automatically. Use the wizard to create a volume on the disk as discussed in the following section.

## Creating a Standard Volume


Standard volumes can be created on any physical or virtual disk available. You use the same technique regardless of how the disk is presented to the server. This allows you to create standard volumes on a server's internal disks, on virtual disks in a storage subsystem available to a server, and on virtual iSCSI disks available to a server. If you add the data deduplication feature to a server, you can enable data deduplication for standard volumes created for that server.

You can create a standard volume by completing the following steps:

1. Start the New Volume Wizard. If you just created a storage space, the New Volume Wizard might start automatically. If it did not, do one of the following:

- On the Disks subnode, all available disks are listed in the Disks panel. Select the disk with which you want to work, and then under Tasks, select New Volume.
- On the Storage Pools subnode, all available virtual disks are listed in the Virtual Disks panel. Select the disk with which you want to work, and then under Tasks, select New Volume.

2. On the Select The Server And Disk page, select the server for which you are provisioning storage, select the disk where the volume should be created, and then select Next. If you just created a storage space and then New Volume Wizard started automatically, the related server and disk are selected automatically and you simply need to select Next.




3. On the Specify The Size Of The Volume page, use the options provided to set the volume size. By default, the volume size is set to the maximum available on the related disk. Select Next.



4. On the Assign To A Drive Letter Or Folder page, specify whether you want to assign a drive letter or path to the volume, and then select Next. You use these options as follows:

- **Drive Letter** To assign a drive letter, choose this option, and then select an available drive letter in the list provided.
- **The Following Folder** To assign a drive path, choose this option, and then enter the path to an existing folder on an NTFS drive, or select Browse to search for or create a folder.
- **Don't Assign To A Drive Letter Or Drive Path** To create the volume without assigning a drive letter or path, choose this option. You can assign a drive letter or path later if necessary.




5. On the Select File System Settings page, specify how the volume should be formatted by using the following options:


- **File System** Sets the file system type, such as NTFS or ReFS.
- **Allocation Unit Size** Sets the cluster size for the file system. This is the basic unit in which disk space is allocated. The default allocation unit size is based on the volume's size and is set dynamically prior to formatting. To override this feature, you can set the allocation unit size to a specific value.
- **Volume Label** Sets a text label for the partition. This label is the partition's volume name.



- If you elected to create an NTFS volume and added data deduplication to the server, you can enable and configure data deduplication. When you are ready to continue, select Next.



- After you confirm your selections, select Create. The wizard tracks the progress of the volume creation. When the wizard finishes creating the volume, the View Results page will be updated to reflect this. Review the details to ensure that all phases were completed successfully. If any portion of the configuration failed, note the reason for the failure and take corrective actions as appropriate before repeating this procedure.



## 8. Select Close.

**REAL WORLD** In the Registry under `HKLM\SYSTEM\CurrentControlSet\Control\FileSystem`, the `NtfsDisableLastAccessUpdate` and `RefsDisableLastAccessUpdate` values control whether NTFS and ReFS update the last-access timestamp on each directory when it lists directories on a volume. If you notice that a busy server with a large number of directories isn't very responsive when you list directories, this could be because the filesystem log buffer in physical memory is getting filled with timestamp update records. To prevent this, you can set the value to 1. When the value is set to 1, the filesystem does not update the last-access timestamp, and it does not record time stamp updates in the file system log. Otherwise, when the value is set to 0 (the default), the filesystem updates the last-access timestamp on each directory it detects, and it records each time change in the filesystem log.



# Troubleshooting Storage Spaces

Typical problems creating storage spaces and allocating storage were discussed previously. You also might find that a physical disk that should be available for use isn't available. With the Storage Pools node selected in Server Manager, you can add a physical disk that has been detected but isn't listed as available by selecting Tasks on the Physical Disks panel, and then selecting Add Physical Disk. Next, in the Add Physical Disk dialog box, select the physical disk, and then select OK. Alternatively, if the physical disk has not been detected by the storage system, select Tasks on the Storage Pools panel, and then select Rescan Storage.

Other problems you might experience with storage spaces relate to drive failures and a loss of resiliency. When a storage space uses two-way mirroring, three-way mirroring, parity, or dual parity, you can recover resiliency by reconnecting a disconnected drive or replacing a failed drive. When a storage space uses a simple volume and drives were disconnected, you can recover the volume by reconnecting the drives.

Selecting the notification icon for Action Center displays the related notifications. If there is a problem with storage spaces, Action Center updates the related notification panel in the desktop notification area with a message stating "Check Storage Spaces for issues." To open Server Manager, select the notification icon, and then select the link provided. In Server Manager, you'll need to select the File And Storage Services node, and then select Storage Pools to get the relevant error and warning icons.

To view errors and warnings for storage pools, right-click the storage pool with the error or warning icon, and then select Properties. In the Properties dialog box, select Health in the left pane to display the health status and operational status in the main pane. For example, you might find that the health status is listed as Warning and the operation status is listed as Degraded. Degraded is a status you'll get when there is a loss of redundancy.

To view errors and warnings for virtual disks and their associated physical disks, rightclick the virtual disk with the error or warning icon, and then select Properties. In the Properties dialog box, select Health in the left pane to display the health status and operational status in the main pane. Note the storage layout and the physical disks in use as well. If there is a problem with a physical disk, such as a loss of communication, this status will be displayed. You'll get a Loss of Communication status when a physical disk is missing, failed, or disconnected.

When storage spaces use external drives, a missing drive might be a common problem you encounter. In this case, users can continue to work and redundancy will be





restored when you reconnect the drive. On the other hand, if a drive failed, you'll need to complete the following steps to restore redundancy:

1. Physically remove the failed drive. If the drive is connected internally, you'll need to shut down and unplug the computer before you can remove the drive; otherwise, simply disconnect an externally connected drive.
2. Physically add or connect a replacement drive. Next, add the drive to the storage space by doing the following:

On the Storage Spaces panel, right-click the storage space you want to configure, and then select Add Physical Drive.

In the Add Physical Disk dialog box, select the drive that should be allocated to the storage pool.

When you select OK, Windows Server will prepare the drive and allocate it to the storage pool.

3. At this point, the failed drive should be listed as "Retired." Remove the failed drive from the storage space by selecting the related Remove Disk option, and then confirm that you want to remove the drive by selecting Yes when prompted.

Windows Server restores redundancy by copying data as necessary to the new disk. During this process, the status of the storage space normally is listed as "Repairing." You'll also see a value depicting how much of the repair task is completed. When this value reaches 100 percent, the repair is complete.





## Chapter 23. Using RAID

Although Storage Spaces are preferred for protecting critical data volumes, Windows Server 2016 continues to support traditional software-based RAID for system volumes and other volumes. You can use RAID technology to span volumes and add fault tolerance. With RAID, you increase data integrity and availability by creating redundant copies of the data. You can also use RAID to improve your disks' performance.

Different implementations of RAID technology are available, and these implementations are described in terms of levels. Each RAID level offers different features. Windows Server 2016 supports RAID levels 0, 1, and 5:

- **RAID 0 (Disk striping)** Two or more volumes, each on a separate drive, are configured as a striped set. Data is broken into blocks, called stripes, and then written sequentially to all drives in the striped set.
- **RAID 1 (Disk mirroring)** Two volumes on two drives are configured identically. Data is written to both drives. If one drive fails, no data loss occurs because the other drive contains the data. (This level doesn't include disk striping.)
- **RAID 5 (Disk striping with parity)** Uses three or more volumes, each on a separate drive, to create a striped set with parity error checking. In the case of failure, data can be recovered.

You can use RAID-0 to improve the performance of your drives, and you use RAID-1 and RAID-5 to provide fault tolerance for data. In addition to redundancy, RAID 1 offers better write performance than disk striping with parity. RAID 5 offers fault tolerance with less overhead than mirroring and better read performance than disk mirroring.

## Implementing RAID

Windows Server 2016 supports disk mirroring, disk striping, and disk striping with parity. Implementing these RAID techniques is discussed in the sections that follow.

The most common RAID levels in use on servers running Windows Server 2016 are level 1 (disk mirroring), and level 5 (disk striping with parity). With respect to upfront costs, disk mirroring is the least expensive way to increase data protection with redundancy. Here, you use two identically sized volumes on two different drives to create a redundant data set. If one of the drives fails, you can still obtain the data from the other drive.





On the other hand, disk striping with parity requires more disks—a minimum of three—but offers fault tolerance with less overhead than disk mirroring. If any of the drives fail, you can recover the data by combining blocks of data on the remaining disks with a parity record. Parity is a method of error checking that uses an exclusive OR operation to create a checksum for each block of data written to the disk. This checksum is used to recover data in case of failure.

**REAL WORLD** Although it's true that the upfront costs for mirroring should be less than the upfront costs for disk striping with parity, the actual cost per gigabyte might be higher with disk mirroring. With disk mirroring, you have an overhead of 50 percent. For example, if you mirror two 750-gigabyte (GB) drives (a total storage space of 1500 GB), the usable space is only 750 GB. With disk striping with parity, on the other hand, you have an overhead of around 33 percent. For example, if you create a RAID-5 set by using three 500-GB drives (a total storage space of 1500 GB), the usable space (with one-third lost for overhead) is 1000 GB.

## Implementing RAID-0: Disk Striping

RAID level 0 is disk striping. With disk striping, two or more volumes—each on a separate drive—are configured as a striped set. Data written to the striped set is broken into blocks called stripes. These stripes are written sequentially to all drives in the striped set. You can place volumes for a striped set on up to 32 drives, but in most circumstances sets with 2 to 5 volumes offer the best performance improvements. Beyond this, the performance improvement decreases significantly.

The major advantage of disk striping is speed. Data can be accessed on multiple disks by using multiple drive heads, which improves performance considerably. However, this performance boost comes with a price tag. As with volume sets, if any hard disk drive in the striped set fails, the striped set can no longer be used, which essentially means that all data in the striped set is lost. You need to re-create the striped set and restore the data from backups.

**CAUTION** The boot and system volumes shouldn't be part of a striped set. Don't use disk striping with these volumes.

When you create striped sets, you should use volumes that are approximately the same size. Disk Management bases the overall size of the striped set on the smallest volume size. Specifically, the maximum size of the striped set is a multiple of the smallest volume size. For example, if you want to create a three-volume striped set but the



smallest volume is 20 GB, the maximum size for the striped set is 60 GB even if the other two values are 2-TB each.

You can maximize performance of the striped set in a couple of ways:

- Use disks that are on separate disk controllers. This allows the system to simultaneously access the drives.
- Don't use the disks containing the striped set for other purposes. This allows the disk to dedicate its time to the striped set.

You can create a striped set by following these steps:

1. In Disk Management's Graphical View, right-click an area marked Unallocated on a dynamic disk, and then click New Striped Volume. This starts the New Striped Volume Wizard. Read the Welcome page, and then click Next.
2. Select the disks that you want to be part of the striped set, and then size the volume segments on those disks. You need at least two dynamic disks to create a striped volume.
3. Available disks are shown in the Available list. If necessary, select a disk in this list, and then click Add to add the disk to the Selected list. If you make a mistake, you can remove disks from the Selected list by selecting the disk, and then clicking Remove.

**CAUTION** The disk wizards in Windows Server 2016 show both basic and dynamic disks with available disk space. If you add space from a basic disk, the wizard converts the disk to a dynamic disk before creating the volume set. Before clicking Yes to continue, be sure you really want to do this because it can affect how the disk is used by the operating system.

4. Select a disk in the Selected list, and then specify the size of the volume on the disk in the Select The Amount Of Space In MB box. The Maximum Available Space In MB box shows you the largest area of free space available on the disk. The Total Volume Size In Megabytes box shows you the total disk space selected for use with the volume. Click Next.

**TIP** Although you can size a volume set any way you want, consider how you'll use volume sets on the system. Simple and spanned volumes aren't fault tolerant; rather than creating one monstrous volume with all the available free space, you might want to create several smaller volumes to help ensure that losing one volume doesn't mean losing all your data.

5. Specify whether you want to assign a drive letter or path to the volume, and then click Next. You use the available options as follows:

- **Assign The Following Drive Letter** To assign a drive letter, choose this option, and then select an available drive letter in the list provided.
- **Mount In The Following Empty NTFS Folder** To assign a drive path, choose this option, and then type the path to an existing folder on an NTFS drive, or click Browse to search for or create a folder.
- **Do Not Assign A Drive Letter Or Drive Path** To create the volume without assigning a drive letter or path, choose this option. You can assign a drive letter or path later if necessary.

6. Specify whether the volume should be formatted. If you elect to format the volume, set the following formatting options:

- **File System** Specifies the file system type, such as NTFS or ReFS.
- **Allocation Unit Size** Specifies the cluster size for the file system. This is the basic unit in which disk space is allocated. The default allocation unit size is based on the volume's size and is set dynamically prior to formatting. Although you can't change the default size if you select ReFS, you can set the allocation unit size to a specific value with other formats. If you use a lot of small files, you might want to use a smaller cluster size, such as 512 or 1024 bytes. With these settings, small files use less disk space.
- **Volume Label** Specifies a text label for the partition. This label is the partition's volume name.
- **Perform A Quick Format** Tells Windows to format without checking the partition for errors. With large partitions, this option can save you a few minutes. However, it's more prudent to check for errors, which allows Disk Management to mark bad sectors on the disk and lock them out.

7. Click Next, and then click Finish.

After you create a striped volume, you can use the volume as you would any other volume. You can't extend a striped set after it's created; therefore, you should carefully consider the setup before you implement it.

## Implementing RAID-1: Disk Mirroring

RAID level 1 is disk mirroring. With disk mirroring, you use identically sized volumes on two different drives to create a redundant data set. The drives are written with identical sets of information, and if one of the drives fails, you can still obtain the data from the other drive.


Disk mirroring offers about the same fault tolerance as disk striping with parity. Because mirrored disks don't need to write parity information, they can offer better write performance in most circumstances. However, disk striping with parity usually offers better read performance because read operations are spread over multiple drives.

The major drawback to disk mirroring is that it effectively cuts the amount of storage space in half. For example, to mirror a 500-GB drive, you need another 500-GB drive. That means you use 1000 GB of space to store 500 GB of information.

**TIP** If possible, you should mirror boot and system volumes. Mirroring these volumes ensures that you are able to boot the server in case of a single drive failure.

As with disk striping, you'll often want the mirrored disks to be on separate disk controllers to provide increased protection against failure of the disk controller. If one of the disk controllers fails, the disk on the other controller is still available.

Technically, when you use two separate disk controllers to duplicate data, you're using a technique known as disk duplexing. The figure that follows shows difference between the two techniques. Where disk mirroring typically uses a single drive controller, disk duplexing uses two drive controllers; otherwise, the two techniques are essentially the same.





If one of the mirrored drives in a set fails, disk operations can continue. Here, when users read and write data, the data is written to the remaining disk. You need to break the mirror before you can fix it. To learn how, see “Managing RAID and Recovering From Failures” later in this chapter.

## Creating a Mirror Set in Disk Management

You create a mirror set by following these steps:

1. In Disk Management’s Graphical View, right-click an area marked Unallocated on a dynamic disk, and then click New Mirrored Volume. This starts the New Mirrored Volume Wizard. Read the Welcome page, and then click Next.
2. Create the volume as described in “Implementing RAID-0: Disk Striping” earlier in this chapter. The key difference when creating the mirror set is that you must create two identically sized volumes, and these volumes must be on separate dynamic drives. You won’t be able to continue past the Select Disks window until you select the two disks with which you want to work.

Like other RAID techniques, mirroring is transparent to users. Users see the mirrored set as a single drive they can access and use like any other drive.

**NOTE** The status of a normal mirror is Healthy. During the creation of a mirror, you’ll get a status of Resynching, which tells you that Disk Management is creating the mirror.

## Mirroring an Existing Volume

Rather than create a new mirrored volume, you can use an existing volume to create a mirrored set. To do this, the volume you want to mirror must be a simple volume, and you must have an area of unallocated space on a second drive of equal or larger size than the existing volume.

In Disk Management, you mirror an existing volume by following these steps:

1. Right-click the simple volume you want to mirror, and then click Add Mirror. This displays the Add Mirror dialog box.
2. In the Disks list, select a location for the mirror, and then click Add Mirror. Windows Server 2016 begins the mirror creation process. In Disk Management, you’ll get a status of Resynching on both volumes. The disk on which the mirrored volume is being created has a warning icon.





## Implementing RAID-5: Disk Striping with Parity

RAID level 5 is disk striping with parity. With this technique, you need a minimum of three hard disk drives to set up fault tolerance. Disk Management sizes the volumes on these drives identically.

RAID-5 distributes data and parity data sequentially across the disks in the array. Fault tolerance ensures that the failure of a single drive won't bring down the entire drive set. Instead, the set continues to function with disk operations directed at the remaining volumes in the set.

To allow for fault tolerance, RAID-5 writes parity checksums with the blocks of data. If any of the drives in the striped set fails, you can use the parity information to recover the data. (This process, called regenerating the striped set, is covered in "Managing RAID and Recovering From Failures" later in the chapter.) If two disks fail, however, the parity information isn't sufficient to recover the data, and you need to rebuild the striped set from backup.

## Creating a Striped Set with Parity in Disk Management

In Disk Management, you can create a striped set with parity by following these steps:

1. In Disk Management's Graphical View, right-click an area marked Unallocated on a dynamic disk, and then click New RAID-5 Volume. This starts the New RAID-5 Volume Wizard. Read the Welcome page, and then click Next.
2. Create the volume as described previously in "Implementing RAID-0: Disk Striping." The key difference when creating a striped set with parity is that you must select free space on three separate dynamic drives.

After you create a striped set with parity (RAID-5), users can use the set just like they would a normal drive. Keep in mind that you can't extend a striped set with parity after you create it; therefore, you should carefully consider the setup before you implement it.

## Managing RAID and Recovering From Failures

Managing mirrored drives and striped sets is somewhat different from managing other drive volumes, especially when it comes to recovering from failure. The techniques you use to manage RAID arrays and to recover from failure are covered in this section.





# Breaking a Mirrored Set

You might want to break a mirror for two reasons:

- If one of the mirrored drives in a set fails, disk operations can continue. When users read and write data, these operations use the remaining disk. At some point, however, you need to fix the mirror, and to do this you must first break the mirror, replace the failed drive, and then reestablish the mirror.
- If you no longer want to mirror your drives, you might also want to break a mirror. This allows you to use the disk space for other purposes.

**BEST PRACTICES** Although breaking a mirror doesn't delete the data in the set, you should always back up the data before you perform this procedure to ensure that if you have problems, you can recover your data.

In Disk Management, you can break a mirrored set by following these steps:

1. Right-click one of the volumes in the mirrored set, and then click **Break Mirrored Volume**.
2. Confirm that you want to break the mirror by clicking **Yes**. If the volume is in use, you'll get another warning dialog box. Confirm that it's okay to continue by clicking **Yes**.

Windows Server 2016 breaks the mirror, creating two independent volumes.

## Resynchronizing and Repairing a Mirrored Set

Windows Server 2016 automatically synchronizes mirrored volumes on dynamic drives; however, data on mirrored drives can become out of sync. For example, if one of the drives goes offline, data is written only to the drive that's online.

You can resynchronize and repair mirrored sets, but you must rebuild the set by using disks with the same partition style—either master boot record (MBR) or GUID partition table (GPT). You need to get both drives in the mirrored set online. The mirrored set's status should read **Failed Redundancy**. The corrective action you take depends on the failed volume's status:

- If the status is **Missing or Offline**, be sure that the drive has power and is connected properly. Then start Disk Management, right-click the failed volume, and click **Reactivate Volume**. The drive status should change to **Regenerating** and then to **Healthy**. If the volume doesn't return to the **Healthy** status, right-click the volume, and then click **Resynchronize Mirror**.
- If the status is **Online (Errors)**, right-click the failed volume, and then click **Reactivate Volume**. The drive status should change to **Regenerating** and then to

Healthy. If the volume doesn't return to the Healthy status, right-click the volume, and then click Resynchronize Mirror.

- If one of the drives shows a status of Unreadable, you might need to rescan the drives on the system by choosing Rescan Disks from Disk Management's Action menu. If the drive status doesn't change, you might need to reboot the computer.
- If one of the drives still won't come back online, right-click the failed volume, and then click Remove Mirror. Next, right-click the remaining volume in the original mirror, and then click Add Mirror. You now need to mirror the volume on an unallocated area of free space. If you don't have free space, you need to create space by deleting other volumes or replacing the failed drive.

## Repairing a Mirrored System Volume to Enable Boot

The failure of a mirrored drive might prevent your system from starting. Typically, this happens when you're mirroring the system or boot volume, or both, and the primary mirror drive has failed. In previous versions of the Windows operating system, you often had to go through several procedures to get the system back up and running.

With

Windows Server 2016, the failure of a primary mirror is usually much easier to resolve.

When you mirror a system volume, the operating system should add an entry to the system's boot manager that allows you to start to the secondary mirror. Resolving a primary mirror failure is much easier with this entry in the boot manager file than without it because all you need to do is select the entry to start to the secondary mirror. If you mirror the boot volume and a secondary mirror entry is not created for you, you can modify the boot entries in the boot manager to create one by using the BCD Editor (Bcdedit.exe).

If a system fails to start to the primary system volume, restart the system and select the Windows Server 2016—Secondary Plex option for the operating system you want to start. The system should start normally. After you successfully start the system to the secondary drive, you can schedule the maintenance necessary to rebuild the mirror as described in the following steps:

1. Shut down the system, and replace the failed volume or add a hard disk drive. Then restart the system.
2. Break the mirror set, and then re-create the mirror on the drive you replaced, which is usually drive 0. Right-click the remaining volume that was part of the original mirror, and then click Add Mirror. Next, follow the technique in "Mirroring an Existing Volume" earlier in the chapter.





3. If you want the primary mirror to be on the drive you added or replaced, use Disk Management to break the mirror again. Be sure that the primary drive in the original mirror set has the drive letter that was previously assigned to the complete mirror. If it doesn't, assign the appropriate drive letter.
4. Right-click the original system volume, and then click Add Mirror. Now recreate the mirror.
5. Check the boot entries in the boot manager and use the BCD Editor to ensure that the original system volume is used during startup.

## Removing a Mirrored Set

Using Disk Management, you can remove one of the volumes from a mirrored set. When you do this, all data on the removed mirror is deleted and the space it used is marked as Unallocated.

To remove a mirror, follow these steps:

1. In Disk Management, right-click one of the volumes in the mirrored set, and then click Remove Mirror to display the Remove Mirror dialog box.
2. In the Remove Mirror dialog box, select the disk from which to remove the mirror.
3. Confirm the action when prompted. All data on the removed mirror is deleted.





## Repairing a Striped Set Without Parity

A striped set without parity doesn't have fault tolerance. If a drive that's part of a striped set fails, the entire striped set is unusable. Before you try to restore the striped set, you should repair or replace the failed drive. Then you need to re-create the striped set and recover the data contained on the striped set from backup.





## Regenerating a Striped Set with Parity

With RAID-5, you can recover the striped set with parity if a single drive fails. You'll know that a striped set with parity drive has failed because the set's status changes to Failed Redundancy and the individual volume's status changes to Missing, Offline, or Online (Errors).

You can repair RAID-5 disks, but you must rebuild the set by using disks with the same partition style—either MBR or GPT. You need to get all drives in the RAID-5 set online. The set's status should read Failed Redundancy. The corrective action you take depends on the failed volume's status:

- If the status is Missing or Offline, make sure the drive has power and is connected properly. Then start Disk Management, right-click the failed volume, and select Reactivate Volume. The drive's status should change to Regenerating and then to Healthy. If the drive's status doesn't return to Healthy, right-click the volume and select Regenerate Parity.
- If the status is Online (Errors), right-click the failed volume, and select Reactivate Volume. The drive's status should change to Regenerating and then to Healthy. If the drive's status doesn't return to Healthy, right-click the volume and select Regenerate Parity.
- If one of the drives shows as Unreadable, you might need to rescan the drives on the system by choosing Rescan Disks from Disk Management's Action menu. If the drive status doesn't change, you might need to restart the computer.
- If one of the drives still won't come back online, you need to repair the failed region of the RAID-5 set. Right-click the failed volume, and then select Remove Volume. You now need to select an unallocated space on a separate dynamic disk for the RAID-5 set. This space must be at least as large as the region to repair, and it can't be on a drive that the RAID-5 set is already using. If you don't have enough space, the Repair Volume command is unavailable, and you need to free space by deleting other volumes or by replacing the failed drive.

## Chapter 24. Maintaining Partitions and Drives

Disk Management provides many features to manage existing partitions and drives. Use these options to assign drive letters, delete partitions, set the active partition, and more. In addition, Windows Server 2016 provides other utilities to carry out common tasks such as converting a volume to NTFS, checking a drive for errors, and cleaning up unused disk space. These options and tools are discussed in this chapter.






# Drive Maintenance Essentials

As part of routine maintenance, you may need to assign drive letters and paths or perform other tasks, such as changing drive labels. You may also need to resize or delete partitions.



# Assigning Drive Letters and Paths

You can assign drives one drive letter and one or more drive paths, provided that the drive paths are mounted on NTFS drives. Drives don't have to be assigned a drive letter or path. A drive with no designators is considered to be unmounted, and you can mount it by assigning a drive letter or path at a later date. You need to unmount a drive before moving it to another computer.



To manage drive letters and paths, right-click the drive you want to configure in Disk Management, and then click Change Drive Letter And Paths. You can now do the following:

- **Add a drive path** Click Add, select Mount In The Following Empty NTFS Folder, and then type the path to an existing folder, or click Browse to search for or create a folder.
- **Remove a drive path** Select the drive path to remove, click Remove, and then click Yes.
- **Assign a drive letter** Click Add, select Assign The Following Drive Letter, and then choose an available letter to assign to the drive.
- **Change the drive letter** Select the current drive letter, and then click Change. Select Assign The Following Drive Letter, and then choose a different letter to assign to the drive.
- **Remove a drive letter** Select the current drive letter, click Remove, and then click Yes.

**NOTE** If you try to change the letter of a drive that's in use, Windows Server 2016 displays a warning. You need to exit programs that are using the drive and try again, or allow Disk Management to force the change by clicking Yes when prompted.



## Changing or Deleting the Volume Label

The volume label is a text descriptor for a drive. With FAT, the volume label can be up to 11 characters and can include spaces. With NTFS and ReFS, the volume label can be up to 32 characters. Additionally, although FAT doesn't allow you to use some special characters—including \* / \ [ ] : ; | = , . + “ ? < >—NTFS and ReFS do allow you to use these special characters.

Because the volume label is displayed when the drive is accessed in various Windows Server 2016 utilities, including File Explorer, it can provide information about a drive's contents. You can change or delete a volume label by using Disk Management or File Explorer.

Using Disk Management, you can change or delete a label by following these steps:

1. Right-click the partition, and then click Properties.
2. On the General tab of the Properties dialog box, enter a new label for the volume in the Label text box or delete the existing label. Click OK.

Using File Explorer, you can change or delete a label by following these steps:

1. Right-click the drive icon, and then click Properties.
2. On the General tab of the Properties dialog box, enter a new label for the volume in the Label text box or delete the existing label. Click OK.

## Resizing Partitions and Volumes

With Windows Server 2016, you can extend and shrink both basic and dynamic disks using Disk Management, DiskPart, or PowerShell. In extending a volume, you convert areas of unallocated space and add them to the existing volume. For spanned volumes on dynamic disks, the space can come from any available dynamic disk, not only from those on which the volume was originally created. Thus, you can combine areas of free space on multiple dynamic disks and use those areas to increase the size of an existing volume.

**CAUTION** Before you try to extend a volume, be aware of several limitations. You cannot shrink or extend volumes that use traditional software RAID. Although you can extend simple and spanned volumes formatted with NTFS or ReFS, you can't extend striped volumes, volumes that aren't formatted, or volumes that are formatted with FAT. Additionally, you can't extend a system or boot volume, regardless of its configuration.

You can shrink a simple volume or a spanned volume by following these steps:



1. In Disk Management, right-click the volume you want to shrink, and then click Shrink Volume. This option is available only if the volume meets the previously discussed criteria.
2. In the box provided in the Shrink dialog box, enter the amount of space to shrink.

The screenshot shows the 'Shrink M:' dialog box with the following fields and values:

| Field                                      | Value   |
|--------------------------------------------|---------|
| Total size before shrink in MB:            | 2097152 |
| Size of available shrink space in MB:      | 2094030 |
| Enter the amount of space to shrink in MB: | 2000000 |
| Total size after shrink in MB:             | 97152   |

Below the fields, there is a warning icon and text: "You cannot shrink a volume beyond the point where any unmovable files are located. See the 'defrag' event in the Application log for detailed information about the operation when it has completed."

At the bottom, there is a link: "See 'Shrink a basic volume' in Disk Management help for more information"

At the bottom right, there are two buttons: "Shrink" and "Cancel".

The Shrink dialog box provides the following information:

- **Total Size Before Shrink In MB** Lists the total capacity of the volume in megabytes. This is the formatted size of the volume.
  - **Size Of Available Shrink Space In MB** Lists the maximum amount by which the volume can be shrunk. This doesn't represent the total amount of free space on the volume; rather, it represents the amount of space that can be removed, not including any data reserved for the master file table, volume snapshots, page files, and temporary files.
  - **Enter The Amount Of Space To Shrink In MB** Lists the total amount of space that will be removed from the volume. The initial value defaults to the maximum amount of space that can be removed from the volume. For optimal drive performance, you'll want to ensure that the drive has at least 10 percent of free space after the shrink operation.
  - **Total Size After Shrink In MB** Lists what the total capacity of the volume will be (in megabytes) after the shrink. This is the new formatted size of the volume.
3. Click Shrink to shrink the volume.

You can extend a simple volume or a spanned volume by following these steps:

1. In Disk Management, right-click the volume you want to extend, and then click Extend Volume. This option is available only if the volume meets the previously discussed criteria and free space is available on one or more of the system's dynamic disks.
2. In the Extend Volume Wizard, read the introductory message, and then click Next.
3. On the Select Disks page, select the disk or disks from which you want to allocate free space. Any disks currently being used by the volume are automatically selected. By default, all remaining free space on those disks is selected for use.

The screenshot shows the 'Extend Volume Wizard' dialog box, specifically the 'Select Disks' step. The title bar reads 'Extend Volume Wizard' with a close button. Below the title, the text 'Select Disks' is followed by the instruction: 'You can use space on one or more disks to extend the volume.' The main area is divided into two sections: 'Available:' on the left and 'Selected:' on the right. The 'Available:' section contains an empty list box. The 'Selected:' section contains a list box with one entry: 'Disk 3 6167325 MB'. Between these two sections are three buttons: 'Add >', '< Remove', and '< Remove All'. Below the list boxes, there are three rows of text with corresponding input fields: 'Total volume size in megabytes (MB):' with the value '6291325', 'Maximum available space in MB:' with the value '6167325', and 'Select the amount of space in MB:' with the value '6167325' and a spin button. At the bottom of the dialog are three buttons: '< Back', 'Next >' (which is highlighted with a blue border), and 'Cancel'.

4. With dynamic disks, you can specify the additional space you want to use on other disks by performing the following tasks:
  - Click the disk, and then click Add to add the disk to the Selected list.
  - Select each disk in the Selected list, and then, in the Select The Amount Of Space In MB list, specify the amount of unallocated space to use on the selected disk.
5. Click Next, confirm your options, and then click Finish.

## Deleting Partitions and Drives

To change the configuration of a drive that's fully allocated, you might need to delete existing partitions and logical drives. Deleting a partition or a drive removes the




associated file system, and all data in the file system is lost. Before you delete a partition or a drive, you should back up any files and directories that the partition or drive contains.

**NOTE** To protect the integrity of the system, you can't delete the system or boot partition. However, Windows Server 2016 does let you delete the active partition or volume if it is not designated as boot or system. Always check to be sure that the partition or volume you are deleting doesn't contain important data or files.

You can delete a primary partition, volume, or logical drive by following these steps:

1. In Disk Management, right-click the partition, volume, or drive you want to delete, and then click Explore. Using File Explorer, move all the data to another volume or verify an existing backup to ensure that the data was properly saved.
2. In Disk Management, right-click the partition, volume, or drive again, and then click the related Delete option.
3. Confirm that you want to delete the selected item by clicking Yes.



The steps for deleting an extended partition differ slightly from those for deleting a primary partition or a logical drive. To delete an extended partition, follow these steps:

4. Delete all the logical drives on the partition following the steps listed in the previous procedure.
5. Select the extended partition area itself and delete it.

## Converting a Volume to NTFS

Windows Server 2016 provides a utility for converting FAT volumes to NTFS. This utility, Convert (Convert.exe), is located in the %SystemRoot% folder. When you convert a volume by using this tool, the file and directory structure is preserved and no data is lost. Keep in mind, however, that Windows Server 2016 doesn't provide a utility for converting NTFS to FAT. The only way to go from NTFS to FAT is to delete the partition by following the steps listed in the previous section, and then re-create the partition as a FAT volume.



# Understanding the Convert Utility Syntax

Convert is run at the command prompt. If you want to convert a drive, use the following syntax: `convert volume /FS:NTFS`

Here volume is the drive letter followed by a colon, drive path, or volume name. For example, if you want to convert the D drive to NTFS, use the following command:

```
convert D: /FS:NTFS
```

If the volume has a label, you are prompted to enter the volume label for the drive, but you are not prompted if the disk doesn't have a label.

The complete syntax for Convert is shown here:

```
convert volume /FS:NTFS [/V] [/X] [/CvtArea: filename ] [/NoSecurity]
```

The options and switches for Convert are used as follows:

- **volume** Sets the volume with which to work, such as G:.
- **/FS:NTFS** Converts to NTFS
- **/V** Sets verbose mode
- **/X** Forces the volume to dismount before the conversion (if necessary)
- **/CvtArea: filename** Sets the name of a contiguous file in the root directory to be a placeholder for NTFS system files
- **/NoSecurity** Removes all security attributes, and makes all files and directories accessible to the group Everyone

The following sample statement uses Convert:

```
convert C: /FS:NTFS /V
```

## Using the Convert Utility

Before you use the Convert utility, determine whether the partition is being used as the active boot partition or a system partition containing the operating system. You can convert the active boot partition to NTFS. Doing so requires that the system gain exclusive access to this partition, which can be obtained only during startup. Thus, if you try to convert the active boot partition to NTFS, Windows Server 2016 displays a prompt asking if you want to schedule the drive to be converted the next time the system starts. If you click Yes, you can restart the system to begin the conversion process.

**TIP** Often, you will need to restart a system several times to completely convert the active boot partition. Don't panic. Let the system proceed with the conversion.





Before the Convert utility actually converts a drive to NTFS, the utility checks whether the drive has enough free space to perform the conversion. Generally, Convert needs a block of free space that's roughly equal to 25 percent of the total space used on the drive. For example, if the drive stores 200 GB of data, Convert needs about 50 GB of free space. If the drive doesn't have enough free space, Convert aborts and tells you that you need to free up some space. On the other hand, if the drive has enough free space, Convert initiates the conversion. Be patient. The conversion process takes several minutes (longer for large drives). Don't access files or applications on the drive while the conversion is in progress.

You can use the /CvtArea option to improve performance on the volume so that space for the master file table (MFT) is reserved. This option helps to prevent fragmentation of the MFT. How? Over time, the MFT might grow larger than the space allocated to it. The operating system must then expand the MFT into other areas of the disk. Although the Optimize Drives utility can defragment the MFT, it cannot move the first section of the MFT, and it is very unlikely that there will be space after the MFT because this will be filled by file data.

To help prevent fragmentation in some cases, you might want to reserve more space than the default (12.5 percent of the partition or volume size). For example, you might want to increase the MFT size if the volume will have many small or average-size files rather than a few large files. To specify the amount of space to reserve, you can use FSUtil to create a placeholder file equal in size to that of the MFT you want to create. You can then convert the volume to NTFS and specify the name of the placeholder file to use with the /CvtArea option.

In the following example, you use FSUtil to create a 1.5-GB (1,500,000,000 bytes) placeholder file named Temp.txt:

```
fsutil file createnew c:\temp.txt 1500000000
```

To use this placeholder file for the MFT when converting drive C to NTFS, you would then type the following command:

```
convert c: /fs:ntfs /cvtarea:temp.txt
```

Notice that the placeholder file is created on the partition or volume that is being converted. During the conversion process, the file is overwritten with NTFS metadata and any unused space in the file is reserved for future use by the MFT.



# Repairing Disks

Windows Server 2016 includes feature enhancements that reduce the amount of manual maintenance you must perform on disk drives. The following enhancements have the most impact on the way you work with disks:

- [Transactional NTFS Self-healing NTFS](#)

Transactional NTFS allows file operations on an NTFS volume to be performed transactionally, which means programs can use a transaction to group sets of file and registry operations so that all of them succeed or none of them succeed. While a transaction is active, changes are not visible outside the transaction. Changes are committed and written fully to disk only when a transaction is completed successfully. If a transaction fails or is incomplete, the program rolls back the transactional work to restore the file system to the state it was in prior to the transaction.

**REAL WORLD** Resilient File System (ReFS) takes the transactional and selfhealing features of NTFS a few steps further. With ReFS, several background processes are used to maintain disk integrity automatically. The scrubber process checks the disk for inconsistencies and errors. If any are found, a repair process localizes the problems and performs automatic online correction. In the rare case that a physical drive has bad sectors that are causing the problem, ReFS uses a salvage process to mark the bad sectors and remove them from the file system—and all while the volume is online.

Transactions that span multiple volumes are coordinated by the Kernel Transaction Manager (KTM). The KTM supports independent recovery of volumes if a transaction fails. The local resource manager for a volume maintains a separate transaction log and is responsible for maintaining threads for transactions separate from threads that perform the file work.

Traditionally, you had to use the Check Disk tool to fix errors and inconsistencies in NTFS volumes on a disk. Because this process can disrupt the availability of Windows systems, Windows Server 2016 uses self-healing NTFS to protect file systems without requiring you to use separate maintenance tools to fix problems. Because much of the self-healing process is enabled and performed automatically, you might need to perform volume maintenance manually only when you are notified by the operating system that a problem cannot be corrected automatically. If such an error occurs, Windows Server 2016 notifies you about the problem and provides possible solutions.

Self-healing NTFS has many advantages over Check Disk, including the following:



- Check Disk must have exclusive access to volumes, which means system and boot volumes can be checked only when the operating system starts up. On the other hand, with self-healing NTFS, the file system is always available and does not need to be corrected offline (in most cases).
- Self-healing NTFS attempts to preserve as much data as possible if corruption occurs and reduces failed file system mounting that previously could occur if a volume was known to have errors or inconsistencies. During restart, self-healing NTFS repairs the volume immediately so that it can be mounted.
- Self-healing NTFS reports changes made to the volume during repair through existing Chkdsk.exe mechanisms, directory notifications, and update sequence number (USN) journal entries. This feature also allows authorized users and administrators to monitor repair operations through Verification, Waiting For Repair Completion, and Progress Status messages.
- Self-healing NTFS can recover a volume if the boot sector is readable but does not identify an NTFS volume. In this case, you must run an offline tool that repairs the boot sector, and then allow self-healing NTFS to initiate recovery.


Although self-healing NTFS is a terrific enhancement, at times you might want to (or might have to) manually check the integrity of a disk. In these cases, you can use Check Disk (Chkdsk.exe) to check for and (optionally) repair problems found on FAT, FAT32, exFAT, and NTFS volumes.

**IMPORTANT** Because ReFS is self-correcting, you do not need to use Check Disk to check ReFS volumes for errors. However, it's important to point out that ReFS as originally released did not efficiently correct corruption on parity spaces. With Windows Server 2016, this deficiency has been corrected. ReFS automatically corrects corruption on parity spaces when integrity streams are enabled to detect corrupt data. When corruption is detected, ReFS examines the data copies that the parity spaces contain and then uses the correct version of the data to correct the problem. As ReFS now supports concurrent I/O requests to the same file, the performance of integrity streams also has been improved.

Although Check Disk can check for and correct many types of errors, the utility primarily looks for inconsistencies in the file system and its related metadata. One of the ways Check Disk locates errors is by comparing the volume bitmap to the disk sectors assigned to files in the file system. Beyond this, the usefulness of Check Disk is rather limited. For example, Check Disk can't repair corrupted data within files that appear to be structurally intact.

As part of automated maintenance, Windows Server 2016 performs a proactive scan of

NTFS volumes. As with other automated maintenance, Windows scans disks using Check Disk at 2:00 A.M. if the computer is running on AC power and the operating system is idle. Otherwise, Windows scans disks the next time the computer is running on AC power and the operating system is idle. Although automated maintenance triggers the disk scan, the process of calling and managing Check Disk is handled by a separate task. In Task Scheduler, you'll find the ProactiveScan task in the scheduler library under Microsoft\Windows\Chkdsk, and you can get detailed run details by reviewing the information provided on the task's History tab.



**REAL WORLD** Automatic Maintenance is built on the Windows Diagnostics framework. By default, Windows periodically performs routine maintenance at 2:00 A.M. if the computer is running on AC power and the operating system is idle. Otherwise, maintenance will start the next time the computer is running on AC power and the operating system is idle. Because maintenance runs only when the operating system is idle, maintenance is allowed to run in the background for up to three days. This allows Windows to complete complex maintenance tasks automatically. Maintenance tasks include software updates, security scanning, system diagnostics, checking disks, and disk optimization. You can change the run time for automated maintenance by opening Action Center, expanding the Maintenance panel, selecting Change Maintenance Settings, and then selecting a new run schedule.



# Checking Disks Manually

With Windows Server 2016, Check Disk performs enhanced scan and repair automatically, instead of using the legacy scan and repair available with earlier releases of Windows. Here, when you use Check Disk with NTFS volumes, Check Disk performs an online scan and analysis of the disk for errors. Check Disk writes information about any detected corruptions in the \$corrupt system file. If the volume is in use, detected corruptions can be repaired by taking the volume offline temporarily; however, unmounting the volume for the repair invalidates all open file handles. With the boot/system volume, the repairs are performed the next time you start the computer.

Storing the corruption information and then repairing the volume while it is dismounted enables Windows to rapidly repair volumes. You can also keep using the disk while a scan is being performed. Typically, offline repair takes only a few seconds, compared to what otherwise would have been hours for very large volumes using the legacy scan and repair technique.

**NOTE** FAT, FAT32, and exFAT do not support the enhanced features. When you use Check Disk with FAT, FAT32, or exFAT, Windows Server 2016 uses the legacy scan and repair process. This means the scan and repair process typically requires taking the volume offline and preventing it from being used. You can't use Check Disk with ReFS.

You can run Check Disk from the command prompt or within other utilities. At a command prompt, you can test the integrity of the E drive by entering the following command:

```
chkdsk /scan E:
```

Check Disk then performs an analysis of the disk and returns a status message regarding any problems it encounters. Check Disk won't repair problems, however, unless you specify further options. To repair errors on drive E, use this command:

```
chkdsk /spotfix E:
```

Fixing the volume requires exclusive access to the volume. The way this works depends on the type of volume:

- For nonsystem volumes, you'll get a prompt asking whether you would like to force a dismount of the volume for the repair. In this case, you can enter **Y** to proceed or **N** to cancel the dismount. If you cancel the dismount, you'll get the prompt asking whether you would like to schedule the volume for the repair the next time the computer is started. In this case, you can enter **Y** to schedule the repair or **N** to cancel



the repair. • For system volumes, you'll get a prompt asking whether you would like to schedule the volume for the repair the next time the computer is started. In this case, you can enter **Y** to schedule the repair or **N** to cancel the repair.

You can't run Check Disk with both the `/scan` and `/spotfix` options because the scan and repair tasks are now independent of each other.

The complete syntax for Check Disk is as follows:

```
CHKDSK [ volume [[ path ] filename ]] [/F] [/V] [/R] [/X] [/I] [/C] [/B]
  [/L: size ] [/scan] [/forceofflinefix] [/perf] [/spotfix]
[/sdcleanup] [/offlinescanandfix]
```

The options and switches for Check Disk are used as follows:

- **volume** Sets the volume with which to work.
- **[path]filename** FAT only. It specifies files to check for fragmentation.
- **/B** Reevaluates bad clusters on the volume (NTFS only; implies `/R` ).
- **/C** NTFS only. It skips the checking of cycles within the folder structure.
- **/F** Fixes errors on the disk by using the offline (legacy) scan and fix behavior.
- **/I** NTFS only. It performs a minimum check of index entries.
- **/L: size** NTFS only. It changes the log file size.
- **/R** Locates bad sectors, and recovers readable information (implies `/F` ).
- **/V** On FAT, it displays the full path and name of every file on the disk. On NTFS, it displays cleanup messages, if there are any.
- **/X** Forces the volume to dismount first if necessary (implies `/F` ).

For NTFS volumes, Check Disk supports these enhanced options:

- **/forceofflinefix** Must be used with `/scan` . It bypasses all online repair and queues errors for offline repair.
- **/offlinescanandfix** Performs an offline scan and fix of the volume.
- **/perf** Performs the scan as fast as possible by using more system resources.
- **/scan** Performs an online scan of the volume (the default). Errors detected during the scan are added to the \$corrupt system file.
- **/sdcleanup** Cleans up unneeded security descriptor data. It implies `/F` (with legacy scan and repair).
- **/spotfix** Allows certain types of errors to be repaired online (the default).


## Running Check Disk Interactively

You can run Check Disk interactively by following these steps:






1. In File Explorer or Disk Management, right-click the drive, and then click Properties.



2. On the Tools tab of the Properties dialog box, click Check. This displays the Error Checking dialog box.



3. Click Scan Drive to start the scan. If no errors are found, Windows confirms this; otherwise,. if errors are found, you'll be prompted with additional options. As with checking disks at a prompt, the way this works depends on whether you are working with a system or nonsystem volume.

**NOTE** For FAT, FAT32, and exFAT volumes, Windows uses the legacy Check Disk. You click Scan And Repair Drive to start the scan. If the scan finds errors, you might need to restart the computer to repair them.

## Analyzing and Optimizing Disks

Any time you add files to or remove files from a drive, the data on the drive can become fragmented. When a drive is fragmented, large files can't be written to a




single continuous area on the disk. As a result, the operating system must write the file to several smaller areas on the disk, which means more time is spent reading the file from the disk. To reduce fragmentation, Windows Server 2016 can manually or automatically analyze and optimize disks by using the Optimize Drives utility.

With manual optimization, Optimize Drives performs an online analysis of volumes, and then reports the percentage of fragmentation. If defragmentation is needed, you can then elect to perform online defragmentation. System and boot volumes can be defragmented online as well, and Optimize Drives can be used with FAT, FAT32, exFAT, NTFS, and ReFS volumes.

You can manually analyze and optimize a disk by following these steps:

1. In File Explorer or Disk Management, right-click the drive, and then click Properties.
2. On the Tools tab, click Optimize. In the Optimize Drives dialog box, select a disk, and then click Analyze. Optimize Drives then analyzes the disk to determine whether it needs to be defragmented. If so, it recommends that you defragment at this point.
3. If a disk needs to be defragmented, select the disk, and then click Optimize.



Automatic analysis and optimization of disks can occur while the disks are online, so long as the computer is on AC power and the operating system is running but otherwise idle. By default, disk optimization is a weekly task rather than a daily task—and there’s a good reason for this. Normally, you need only to periodically optimize a




server's disks, and optimization once a week is sufficient in most cases. Note, however, that although nonsystem disks can be rapidly analyzed and optimized, it can take significantly longer to optimize system disks online.

You can control the approximate start time for the analysis and optimization of disks by changing the automated maintenance start time. Windows Server also notifies you if three consecutive runs are missed. All internal drives and certain external drives are optimized automatically as part of the regular schedule, as are new drives you connect to the server.

**NOTE** Windows Server 2016 automatically performs cyclic pickup defragmentation. With this feature, when a scheduled defragmentation pass is stopped and rerun, the computer automatically picks up where it left off or starts with the next unfinished volume in line to be defragmented.

You can configure and manage automated defragmentation by following these steps:


1. In Computer Management, select the Storage node and then the Disk Management node. Right-click a drive, and then click Properties.
2. On the Tools tab, click Optimize. This displays the Optimize Drives dialog box.
3. If you want to change how optimization works, click Change Settings. To cancel automated defragmentation, clear the Run On A Schedule check box. To enable automated defragmentation, select Run On A Schedule.



4. The default run frequency is set as shown. In the Frequency list, you can choose Daily, Weekly, or Monthly as the run schedule. If you don't want to be notified

about missed runs, clear the Notify Me check box.

5. If you want to manage which disks are defragmented, click Choose, and then select the volumes to defragment. By default, all disks installed within or connected to the computer are defragmented, and any new disks are defragmented automatically as well. Select the check boxes for disks that should be defragmented automatically, and clear the check boxes for disks that should not be defragmented automatically. Click OK to save your settings.



6. Click OK, and then click Close.



## Chapter 25. Implementing File Sharing

The Server Message Block (SMB) protocol is the primary file sharing protocol used by computers running Microsoft Windows. When folders are shared over a network, an SMB client reads and writes to files and requests services from computers hosting SMB-shared folders. With SMB, Windows Server 2016 supports standard file sharing and public folder sharing. Standard file sharing enables remote users to access network resources such as files, folders, and drives. When you share a folder or a drive, you make all its files and subfolders available to a specified set of users. Because you don't need to move files from their current location, standard file sharing is also referred to as in-place file sharing .

You can enable standard file sharing on disks formatted with FAT, FAT32, exFAT, NTFS, or Resilient File System (ReFS). One set of permissions apply to disks formatted with exFAT, FAT, or FAT32. These permissions are called share permissions . Two sets of permissions apply to disks formatted with NTFS or ReFS: NTFS permissions (also referred to as access permissions ) and share permissions . Having two sets of permissions enables you to determine precisely who has access to shared files and the level of access assigned. With either NTFS permissions or share permissions, you do not need to move the files you are sharing.

### Working with SMB

SMB was designed to enable clients and servers to negotiate and then use the highest version supported by both the client attempting to connect an SMB share and the server hosting the share. At the time of this writing, the current version of SMB is version 3.11, which is supported by Windows 10 and Windows Server 2016. Thus, when a Windows 10 computer connects to an SMB share hosted on a server running Windows Server 2016, SMB 3.11 is the version used for the SMB session.





# SMB Essentials

Standard (in-place) file sharing allows remote users to access files, folders, and drives over the network. When you share a folder or a drive, you make all its files and subfolders available to a specified set of users. Share permissions and access permissions together enable you to control who has access to shared files and the level of access assigned. You do not need to move the files you are sharing.

With standard file sharing, local users don't have automatic access to any data stored on a computer. You control local access to files and folders by using the security settings on the local disk. Windows Server 2016 adds layers of security through compound identities, claims-based access controls, and central access policies. With both Windows 10 and Windows Server 2016, you can assign claims-based access controls to file and folder resources on NTFS and ReFS volumes. Users are granted access to file and folder resources, either directly with access permissions and share permissions or indirectly with claims-based access controls and central access policies.

The earliest implementation of SMB was called CIFS, which was introduced with Windows NT 4.0, followed by SMB 1.0, which was used by all versions of Windows from Windows 2000 to Windows Server 2003 R2. In Windows Server 2016, support for CIFS and SMB 1.0 is an optional feature that must be enabled. Because CIFS and SMB 1.0 are outdated, perform poorly, and are less secure than their predecessors, SMB 1.0/CIFS File Sharing Support should not be enabled unless required.

You can enter **Get-SmbConnection** at an administrator PowerShell prompt to determine the version of SMB a client has negotiated with a file server. In the command output, the version is listed in the Dialect column, as shown in the following sample output:

| ServerName | ShareName   | UserName         | Credential       | Dialect | NumOpens |
|------------|-------------|------------------|------------------|---------|----------|
| -----      | -----       | -----            | -----            | -----   | -----    |
| Server36   | IPC\$       | TVPRESS\williams | TVPRESS\williams | 3.11    | 0        |
| Server36   | PrimaryData | TVPRESS\williams | TVPRESS\williams | 3.11    | 14       |

# SMB Encryption and Bandwidth Limits

Both Windows 10 and Windows Server 2016 support SMB 3.0. (They have an SMB 3.0 client.) SMB 3.0 makes it possible to encrypt data being transferred over the network. You can enable SMB encryption for shares configured on NTFS and ReFS volumes. SMB encryption works only when the computer requesting data from an SMB-based share (either a standard file share or a DFS share) and the server supplying the data support SMB 3.0.





If you add the SMB Bandwidth Limit feature to a server, you can use Windows PowerShell to define bandwidth limits for three categories of file transfers:

- **Default** Standard file transfers
- **VirtualMachine** File transfers related to virtual machines
- **LiveMigration** File transfers related to live migration

The cmdlet you use to set limits is `Set-SMBBandwidthLimit`. The basic syntax is:

```
Set-SMBBandwidthLimit -Category Category -BytesPerSecond Limit
```

Where `Category` sets the type of transfers to limit and `Limit` sets the maximum number of bytes per second that the server can send for the specified traffic category. Add KB, MB or GB to define a limit in kilobytes, megabytes or gigabytes as shown in these examples:

```
Set-SMBBandwidthLimit -Category Default -BytesPerSecond 1 GB
```

```
Set-SMBBandwidthLimit -Category VirtualMachine -BytesPerSecond 100 MB
```

```
Set-SMBBandwidthLimit -Category LiveMigration -BytesPerSecond 900 KB
```

## SMB Replication

Windows Server 2016 also supports server to server storage replication over SMB. Here, domain-joined servers, each with separate storage, synchronously or asynchronously replicate data to ensure critical data is always protected. The storage can be local SCSI/SATA, Just a Bunch of Disks (JBOD) SAS, iSCSI target or fibre channel SAN but must be available only to each of the servers with no shared access. Although both hard disk drives and solid state drives can be used for data drives, SSDs are recommended for log drives.

You enable replication by installing the Storage Replica feature, configuring storage and then establishing a replication partnership using Windows PowerShell. Storage on each server must be configured with two virtual disks: one for data and one for logs. The data disks and log disks must be identically sized and the underlying physical disks must have the same sector size.

The virtual disks can be created using Storage Spaces, either in a mirrored or parity configuration. They also can be created as volumes on disks protected with software or hardware RAID.

**NOTE** As neither disk should contain system or boot files, you can't replicate the operating system volume or volumes containing page files or dump files. Before setting up a replication partnership, you can use `Test-SRTopology` to ensure your servers and environment meet all of the requirements.







The cmdlet you use to establish replication partnerships is New-SRPartnership. The basic syntax is:

```
New-SRPartnership -SourceComputerName ComputerName -SourceRGName  
RGName  
-SourceVolumeName Drive -SourceLogVolumeName Drive -  
DestinationComputerName  
ComputerName -DestinationRGName RGName -DestinationVolumeName Drive -  
DestinationLogVolumeName Drive
```

Here, you specify the source and destination servers and assign each a replication group name while also specifying the source data volume, the source log volume, the destination data volume and the destination log volume, such as:

```
New-SRPartnership -SourceComputerName FileServer23 -SourceRGName SR01 -  
SourceVolumeName g: -SourceLogVolumeName h: -DestinationComputerName  
FileServer87 -DestinationRGName SR02 -DestinationVolumeName s:  
-DestinationLogVolumeName t:
```

The replication group name is used to get information about and work with the storage replica. When replication is being configured, the destination volumes are and their drive letters or mount points are dismounted. This ensures the destination volumes are protected and otherwise inaccessible while replication is active.

You can get information about replication using Get-SRGroup, Get-SRPartnership and Get-SRGroup replicas. These commands also help you track the status of replication. As related events also are recorded in the event logs, you can use the following command to monitor related events as well:

```
Get-WinEvent -ProviderName Microsoft-Windows-StorageReplica
```

You can use Set-SRPartnership to modify the replication settings. To stop replication, you must remove the replication partnership and related groups as shown in this example:

```
Get-SRPartnership | Remove-SRPartnership Get-  
SRGroup | Remove-SRGroup
```

Storage Replica is designed to replace DFS replication, as it is much simpler to set up and manage. However, storage replication does have its limitations:

- Storage Replica only allows one-to-one replication between designated volumes. If you want to replicate multiple volumes between servers, you must configure multiple replication partnerships.
- Storage Replica doesn't allow users to access the replicated data on the destination server while replication is active. You must stop replication if you want users to be able to access the replicated data on the destination server.






# Public Folder Sharing

Windows Server 2016 supports public folder sharing in addition to in-place file sharing. Public folder sharing is designed to enable users to share files and folders from a single location. With public folder sharing, you copy or move files you want to share to a computer's %SystemDrive%\Users\Public folder. You can access public folders in File Explorer by double-clicking the system drive, and then accessing the Users\Public folder.

The Public folder has several subfolders you can use to help organize public files:

- **Public Desktop** Used for shared desktop items. Any files and program shortcuts placed in the Public Desktop folder appear on the desktop of all users who log on to the computer (and to all network users if network access has been granted to the Public folder).
- **Public Documents, Public Music, Public Pictures, Public Videos** Used for shared document and media files. All files placed in one of these subfolders are available to all users who log on to the computer (and to all network users if network access has been granted to the Public folder).
- **Public Downloads** Used for shared downloads. Any downloads placed in the Public Downloads subfolder are available to all users who log on to the computer (and to all network users if network access has been granted to the Public folder).

**NOTE** By default, the Public Desktop folder is hidden from view. If hidden items aren't being displayed in File Explorer, click View, and then select Hidden Items.





By default, anyone with a user account and password on a computer can access that computer's Public folder. When you copy or move files to the Public folder, access permissions are changed to match that of the Public folder, and some additional permissions are added as well.


You can change the default Public folder sharing configuration in two key ways:

- Allow users logged on to the computer to view and manage public files but restrict network users from accessing public files. When you configure this option, the implicit groups Interactive, Batch, and Service are granted special permissions on public files and public folders.
- Allow users with network access to view and manage public files. This allows network users to open, change, create, and delete public files. When you configure this option, the implicit group Everyone is granted Full Control permission to public files and public folders.

Windows Server 2016 can use either or both sharing models at any time. However, standard file sharing offers more security and better protection than public folder sharing, and increasing security is essential to protecting your organization's data. With standard file sharing, share permissions are used only when a user attempts to access a file or folder from a different computer on the network. Access permissions are always used, whether the user is logged on to the console or is using a remote system to access a file or folder over the network. When data is accessed remotely, first the share permissions are applied, and then the access permissions are applied.

You can configure the basic file sharing settings for a server by using Advanced Sharing Settings in Network And Sharing Center. Separate options are provided for network discovery, file and printer sharing, and public folder sharing.





You can manage a computer's sharing configuration by following these steps:

1. In Control Panel, click View Network Status And Tasks under the Network And Internet heading to open Network And Sharing Center.
2. In Network And Sharing Center, click Change Advanced Sharing Settings in the left pane. Select the network profile for the network on which you want to enable file and printer sharing. Typically, this will be the Domain profile.
3. Standard file and printer sharing controls network access to shared resources. To configure standard file sharing, do one of the following:
  - Select Turn On File And Printer Sharing to enable file sharing.
  - Select Turn Off File And Printer Sharing to disable file sharing.
4. Public folder sharing controls access to a computer's Public folder. To configure public folder sharing, expand the All Networks panel by clicking the related expand button. On the Public Folder Sharing panel, choose one of the following options:
  - **Turn On Sharing So Anyone With Network Access Can Read And Write Files In The Public Folders** Enables public folder sharing by granting access to the Public folder and all public data to anyone who can access the computer over the network. Windows Firewall settings might prevent external access.
  - **Turn Off Public Folder Sharing** Disables public folder sharing, preventing local network access to the Public folder. Anyone who logs on locally to your computer can still access the Public folder and its files.

- 
5. Click Save Changes.

## Configuring Standard File Sharing

You use shares to control access for remote users. Permissions on shared folders have no effect on users who log on locally to a server or to a workstation that has shared folders.

# Viewing Existing Shares

You can use both Computer Management and Server Manager to work with shares. You also can view current shares on a computer by entering **net share** at a command prompt or by entering **get-smbshare** at a PowerShell prompt.


**TIP** The get-smbshare cmdlet is only one of many cmdlets associated with the smbshare module. To get a list of other cmdlets available for working with SMB shares, enter **get-command -module smbshare** at a PowerShell prompt.

**NOTE** Computer Management, net share, and get-smbshare display information about

SMB-based shares, including standard SMB folder shares, hidden SMB folder shares (those ending with the \$ suffix), and SMB folders shared by using Distributed File System (DFS). Server Manager displays information about standard SMB folder shares, SMB folders shared by using DFS, and folders shared by using Network File System (NFS). Server Manager does not display information about hidden SMB folder shares.

In Computer Management, you can view the shared folders on a local or remote computer by following these steps:

1. To connect to a remote computer (you're connected to the local computer by default), right-click the Computer Management node, and then click **Connect To Another Computer**. Choose Another Computer, enter the name or IP address of the computer to which you want to connect, and then click OK.
2. In the console tree, expand System Tools, expand Shared Folders, and then select Shares. The current shares on the system are displayed.




3. The columns for the Shares node provide the following information:

- **Share Name** Name of the shared folder.
- **Folder Path** Complete path to the folder on the local system.
- **Type** What kind of computers can use the share. This normally shows as Windows because SMB shares are for Windows-based computers.
- **# Client Connections** Number of clients currently accessing the share.
- **Description** Description of the share.

In Server Manager, you can view the shared folders on a local or remote computer by following these steps:

4. Select the File And Storage Services node, and then select the related Shares subnode.
5. The Shares subnode provides information about shares on each file server that has been added for management. The columns for the Shares subnode provide the following information:
  - **Share** Name of the shared folder.
  - **Local Path** Complete path to the folder on the local system.
  - **Protocol** What protocol the share uses, either SMB or NFS.
  - **Cluster Role** If the server sharing the folder is part of a cluster, the cluster role is shown here. Otherwise, the cluster role is listed as None.



6. When you click a share in the Shares pane, information about the related volume is displayed in the Volume pane.

**REAL WORLD** NFS is the file sharing protocol used by UNIX-based systems, which includes computers running Apple OS X. As discussed in “Configuring NFS sharing” later in this chapter, you can enable support for NFS by installing the Server For NFS role service as part of the file server configuration.

# Creating Shares in Computer Management

Windows Server 2016 provides several ways to share folders. You can share local folders by using File Explorer, and you can share local and remote folders by using Computer Management or Server Manager.

When you create a share with Computer Management, you can configure its share permissions and offline settings. When you create a share with Server Manager, you can provision all aspects of sharing, including NTFS permissions, encrypted data access, offline settings for caching, and share permissions. Normally, you create shares on NTFS volumes because NTFS offers the most robust solution.

In Computer Management, you share a folder by following these steps:

1. Right-click Shares, and then click New Share. This starts the Create A Shared Folder Wizard. Click Next.
2. In the Folder Path text box, enter the local file path to the folder you want to share, such as **E:\Reports** . If you don't know the full path, click Browse, use the Browse For Folder dialog box to find the folder you want to share, and then click OK. Click Next.



**TIP** If the file path you specified doesn't exist, the wizard can create it for you. Click Yes when prompted to create the necessary folder or folders.

3. In the Share Name text box, enter a name for the share. This is the name of the folder to which users will connect. Share names must be unique for each system.




**TIP** If you want to hide a share from users (which means that they won't be able to view the shared resource when they try to browse to it in File Explorer or at the command line), enter a dollar sign ( \$ ) as the last character of the shared resource name. For example, you could create a share called PrivEngData\$, which would be hidden from File Explorer, NET VIEW, and other similar utilities. Users can still connect to the share and access its data if they've been granted access permission and they know the share's name. Note that the \$ must be typed as part of the share name when mapping to the shared resource.

4. If you want to, enter a description of the share in the Description text box. When you view shares on a particular computer, the description is displayed in Computer Management.
  5. By default, the share is configured so that only files and programs that users specify are available for offline use. Normally, this is the option you want to use because this option also enables users to take advantage of the new Always Offline feature. If you want to use different offline file settings, click Change, select the appropriate options in the Offline Settings dialog box, and then click OK. The offline availability settings available include the following:
- **Only The Files And Programs That Users Specify Are Available Offline**  
Select this option if you want client computers to cache only the files and programs that users specify for offline use. Optionally, if the BranchCache For Network Files role service is installed on the file server, select Enable BranchCache to enable computers in a branch office to cache files that are downloaded from the shared folder, and then securely share the files to other computers in the branch office.



- **No Files Or Programs From The Shared Folder Are Available Offline** Select this option if you don't want cached copies of the files and programs in the share to be available on client computers for offline use.
- **All Files And Programs That Users Open From The Shared Folder Are Automatically Available Offline** Select this option if you want client computers to automatically cache all files and programs that users open from the share. Optionally, select Optimize For Performance to run cached program files from the local cache instead of the shared folder on the server.



6. Click Next, and then set basic permissions for the share. You'll find helpful pointers in "Managing Share Permissions" later in the chapter. The available options are as follows:

- **All Users Have Read-Only Access** Gives users access to view files and read data. They can't create, modify, or delete files and folders.
- **Administrators Have Full Access; Other Users Have Read-Only Access** Gives administrators complete control over the share. Full access allows administrators to create, modify, and delete files and folders. On an NTFS volume or partition, it also gives administrators the right to change permissions and to take ownership of files and folders. Other users can view files and read data; however, they can't create, modify, or delete files and folders.
- **Administrators Have Full Access; Other Users Have No Access** Gives administrators complete control over the share, but prevents other users from accessing the share.
- **Customize Permissions** Allows you to configure access for specific users and groups, which is usually the best technique to use. Setting share permissions is discussed fully in "Managing Share Permissions."



7. When you click Finish, the wizard creates the share and displays a status report, which should state “Sharing Was Successful.” If an error is displayed instead, note the error and take corrective action as appropriate before repeating this procedure to create the share. Click Finish.

Individual folders can have multiple shares. Each share can have a different name and a different set of access permissions. To create additional shares on an existing share, simply follow the preceding steps for creating a share with these changes:


- In step 3, when you name the share, make sure that you use a different name. In
- step 4, when you add a description for the share, use a description that explains what the share is used for and how it’s different from the other shares for the same folder.

## Creating Shares in Server Manager

In Server Manager, you share a folder by following these steps:

1. The Shares subnode of the File And Storage Services node shows existing shares for file servers that have been added for management. In the Shares pane, click Tasks, and then click New Share to start the New Share Wizard.
  2. Choose one of the available file share profiles, and then click Next. The New Share Wizard has the following file share profiles:
- **SMB Share—Quick** A basic profile for creating SMB file shares that allows you to configure the settings and permissions of the shares.
  - **SMB Share—Advanced** An advanced profile for creating SMB file shares that allows you to configure the settings, permissions, management properties, and NTFS quota profile (if applicable) of the shares.
  - **SMB Share—Applications** A custom profile for creating SMB file shares with settings appropriate for Hyper-V, certain databases, and other server applications.


It's essentially the same as the quick profile, but it doesn't allow you to enable access-based enumeration or offline caching.



**NOTE** If you are using the Server For NFS role service, options are available for creating NFS shares as well.

**REAL WORLD** SMB 3 includes enhancements for server-based applications. These enhancements improve performance for small random reads and writes, which are common with server-based applications, such as Microsoft SQL Server OLTP. With SMB 3, packets use large Maximum Transmission Units (MTUs) as well, which enhance performance for large, sequential data transfers, such as those used for deploying and copying virtual hard disks (VHDs) over the network, database backup and restore over the network, and SQL Server data warehouse transactions over the network.

3. On the Select The Server And Path For This Share page, select the server on which you want the share to be created. Only file servers you've added for management are available.



4. Select the volume on which you want the share to be created. By default, Server Manager creates the file share as a new folder in the \Shares directory on the selected volume. To override this, choose the Type A Custom Path option, and then either enter the share path, such as C:\CorpData, or click Browse to use the Select Folder dialog box to select the share path. When you are ready to continue, click Next.

Share location:

☐ Select by volume:

| Volume | Free Space | Capacity | File System |
|--------|------------|----------|-------------|
| C:     | 135 GB     | 160 GB   | NTFS        |
| D:     | 71.8 GB    | 72.2 GB  | ReFS        |

The location of the file share will be a new folder in the \Shares directory on the selected volume.

☒ Type a custom path:

C:\CorpData

NOTE If you use a custom path, the folder path on the designated drive doesn't have to already be created. The folder path will be created as part of the set up process.

5. On the Specify Share Name page, enter a name for the share. This is the name of the folder to which users will connect. Share names must be unique for each system.

Specify share name

Select Profile  
Share Location  
**Share Name**  
Other Settings  
Permissions  
Confirmation  
Results

Share name: CorpData

Share description:

Local path to share: C:\CorpData


Remote path to share: \\CorpNet28\CorpData

6. If you want to, enter a description of the share in the Description text box. When you view shares on a particular computer, the description is displayed in Computer Management.
7. Note the local and remote paths to the share. These paths are set based on the share location and share name you specified. When you are ready to continue, click Next.
8. If you used a custom path and the folder hasn't been created yet, you are prompted to create path. Click OK.



9. On the Configure Share Settings page, use the following options to configure the way the share is used:
  - **Enable Access-Based Enumeration** Configures permissions so that when users browse the folder, only files and folders a user has been granted at least Read access to are displayed. If a user doesn't have at least Read (or equivalent) permission for a file or folder within the shared folder, that file or folder is hidden from view. (This option is dimmed if you are creating an SMB share optimized for applications.)
  - **Allow Caching Of Share** Configures the share to cache only the files and programs that users specify for offline use. Although you can later edit the share properties and change the offline files' availability settings, you normally want to select this option because it allows users to take advantage of the new Always Offline feature. Optionally, if the BranchCache For Network Files role service is installed on the file server, select Enable BranchCache to enable computers in a branch office to cache files that are downloaded from the shared folder and then securely share the files to other computers in the branch office. (This option is dimmed if you are creating an SMB share optimized for applications.)
  - **Encrypt Data Access** Configures the share to use SMB encryption, which protects file data from eavesdropping while being transferred over the network. This option is useful on untrusted networks.
10. On the Specify Permissions To Control Access page, the default permissions assigned to the share are listed. By default, the special group Everyone is granted the Full Control share permission and the underlying folder permissions are as listed. To change share, folder, or both permissions, click Customize Permissions, and then use the Advanced Security Settings dialog box to configure the required permissions. Setting share permissions is discussed fully in "Managing Share Permissions" later in this chapter. Setting folder permissions is discussed fully in "File and Folder Permissions" in Chapter 27 "Managing Permissions and Auditing."





**NOTE** If the share will be used for Hyper-V, you might need to enable constrained delegation for remote management of the Hyper-V host.


11. If you are using the advanced profile, optionally set the folder management properties, and then click Next. These properties specify the purpose of the folder and the type of data stored in it so that data management policies, such as classification rules, can then use these properties.
12. If you are using the advanced profile, optionally apply a quota based on a template to the folder, and then click Next. You can select only quota templates that have already been created. For more information, see “Managing Disk Quota Templates” in Chapter 28.
13. On the Confirm Selections page, review your selections. When you click Create, the wizard creates the share, configures it, and sets permissions. The status should state, “The share was successfully created.” If an error is displayed instead, note the error and take corrective action as appropriate before repeating this procedure to create the share. Click Close.

## Changing Share Settings

When you create a share, you can configure many basic and advanced settings, including those for access-based enumeration, encrypted data access, offline settings for caching, and management properties. In Server Manager, you can modify these settings by following these steps:

1. The Shares subnode of the File And Storage Services node shows existing shares for file servers that have been added for management. Right-click the share with which you want to work, and then click Properties.

2. In the Properties dialog box, you have several options panels that can be accessed by using controls in the left pane. You can expand the panels one by one or click Show All to expand all the panels at the same time.
3. Use the options provided to modify the settings as necessary, and then click OK. The options available are the same whether you use the basic, advanced, or applications profile to create the shared folder.



**TIP** If you're creating a share for general use and general access, you can publish the shared resource in Active Directory. Publishing the resource in Active Directory makes finding the share easier for users; however, this option is not available in Server Manager. To publish a share in Active Directory, right-click the share in Computer Management, and then click Properties. On the Publish tab, select the Publish This Share In Active Directory check box, add an optional description and owner information, and then click OK.



# Managing Share Permissions

Share permissions set the maximum allowable actions available within a shared folder. By default, when you create a share, everyone with access to the network has Read access to the share's contents. This is an important security change—in previous editions of Windows Server, the default permission was Full Control.

With NTFS and ReFS volumes, you can use file and folder permissions and ownership, in addition to share permissions, to further constrain actions within the share. With FAT volumes, share permissions control only access.







# Understanding the Various Share Permissions

From the most restrictive to the least restrictive, the share permissions available are as follows:

- **No Access** No permissions are granted for the share.
- **Read** Users can do the following:
  - View file and subfolder names
  - Access the subfolders in the share
  - Read file data and attributes
  - Run program files
- **Change** Users have Read permission and the ability to do the following:
  - Create files and subfolders
  - Modify files
  - Change attributes on files and subfolders
  - Delete files and subfolders
- **Full Control** Users have Read and Change permissions, in addition to the following capabilities on NTFS volumes:
  - Change file and folder permissions
  - Take ownership of files and folders


You can assign share permissions to users and groups. You can even assign permissions to implicit groups. For details on implicit groups, see Chapter 7, “Accounts: The Essentials.”

## Viewing and Configuring Share Permissions


You can view and configure share permissions in Computer Management or Server Manager. To view and configure share permissions in Computer Management, follow these steps:

1. In Computer Management, connect to the computer on which the share is created. In the console tree, expand System Tools, expand Shared Folders, and then select Shares.
2. Right-click the share with which you want to work, and then click Properties.
3. In the Properties dialog box, click the Share Permissions tab. You can now view the users and groups that have access to the share and the type of access they have.





4. Users or groups that already have access to the share are listed in the Group Or User Names list. You can remove permissions for these users and groups by selecting the user or group you want to remove, and then clicking Remove. You can change permissions for these users and groups by selecting the user or group you want to change and then allowing or denying access permissions in the Permissions list.
5. To add permissions for another user or group, click Add. This opens the Select Users, Computers, Service Accounts, Or Groups dialog box.





6. Enter the name of a user, computer, or group in the current domain, and then click Check Names. This produces one of the following results:
  - If a single match is found, the dialog box is automatically updated and the entry is underlined.
  - If no matches are found, you either entered an incorrect name part or you're working with an incorrect location. Modify the name and try again, or click Locations to select a new location.
  - If multiple matches are found, select the name or names you want to use, and then click OK. To assign permissions to other users, computers, or groups, enter a semicolon ( ; ) and then repeat this step.

**NOTE** The Locations button enables you to access account names in other domains. Click Locations to find a list of the current domains, trusted domains, and other resources you can access. Because of the transitive trusts in Windows Server, you can usually access all the domains in the domain tree or forest.

7. Click OK. The users and groups are added to the Group Or User Names list for the share.
8. Configure access permissions for each user, computer, and group by selecting an account name and then allowing or denying access permissions. Keep in mind that you're setting the maximum allowable permissions for a particular account.
9. Click OK. To assign additional security permissions for NTFS, see "File and Folder Permissions" in Chapter 27.


**IMPORTANT** Keep in mind that you can select the opposite permission to override an inherited permission. Note also that Deny normally overrides Allow, so if you explicitly deny permission to a user or group for a child folder or file, this permission should be denied to that user or group of users.

To view and configure share permissions in Server Manager, follow these steps:


1. The Shares subnode of the File And Storage Services node shows existing shares for file servers that have been added for management.
2. Right-click the share with which you want to work, and then click Properties.
3. In the Properties dialog box, click the Permissions in the left pane. You can now view the users and groups that have access to the share and the type of access they have.



4. To change share, folder, or both permissions, click Customize Permissions. Next, select the Share tab in the Advanced Security Settings dialog box.



5. Users or groups that already have access to the share are listed in the Permission Entries list. Although you can't modify inherited permissions, you can block inheritance by clicking Disable Inheritance and then either converting the inherited permissions to explicit permissions that can be managed or removing all inherited permissions.



While you are working with the Advanced Security Settings For... dialog box you can manage permissions in a variety of ways. You can change, remove or add non-inherited permissions. To change permissions for users and groups, double-click the user or group you want to work with, and then use the Permissions For... dialog box to manage

the permissions. To remove permissions for users or groups, select the user or group by clicking it and then click Remove.



To add permissions for another user or group, click Add. This opens the Permission Entry dialog box. Next, click Select A Principal to display the Select User, Computer, Service Account Or Group dialog box. Enter the name of a user or a group account. Be sure to reference the user account name rather than the user's full name. Only one name can be entered at a time.

Click Check Names. If a single match is found for each entry, the dialog box is automatically updated, and the entry is underlined. Otherwise, you'll get an additional dialog box. If no matches are found, you either entered the name incorrectly or you're working with an incorrect location. Modify the name in the Name Not Found dialog box and try again, or click Locations to select a new location. When multiple matches are found, in the Multiple Names Found dialog box, select the name you want to use, and then click OK twice.

The user and group is added as the Principal, and the Permission Entry dialog box is updated to show this. Use the Type list to specify whether you are configuring allowed or denied permissions, and then select the permissions you want to allow or deny. Click OK to return to the Advanced Security Settings dialog box. You can assign additional security permissions for NTFS as discussed in "File and Folder Permissions" in Chapter 27.



# Connecting to Network Drives

Users can connect to a network drive and to shared resources available on the network. This connection is shown as a network drive that users can access like any other drive on their systems.

**NOTE** When users connect to network drives, they're subject not only to the permissions set for the shared resources, but also to Windows Server 2016 file and folder permissions. Differences in these permission sets are usually the reason users might not be able to access a particular file or subfolder within the network drive.



# Mapping a Network Drive

In Windows Server 2016, you connect to a network drive by mapping to it using NET USE and New-PsDrive. The syntax for NET USE is the following syntax: `net use DeviceName \\ ComputerName \ ShareName`

DeviceName specifies the drive letter or an asterisk (\*) to use the next available drive letter, and \\ ComputerName \ ShareName is the UNC path to the share, such as either of the following: `net use g: \\ROMEO\DOCS` OR `net use * \\ROMEO\DOCS`

**NOTE** To ensure that the mapped drive is available each time the user logs on, make the mapping persistent by adding the /Persistent:Yes option.

The syntax for New-PsDrive is:

```
New-PsDrive -Name DriveLetter -Root \\ ServerName \ ShareName -  
PsProvider FileSystem
```

Where DriveLetter is the driver letter to use and ServerName is the DNS name or IP address of the server hosting the share and ShareName is the name of the share, such as:

```
New-PsDrive -Name g: -Root \\CorpServer21\CorpData  
-PsProvider FileSystem
```

**NOTE** To ensure that the mapped drive is available each time the user logs on, add the -Persist parameter.

If the client computer is running Windows 10, you can map network drives by completing the following steps:

1. When you open File Explorer, the This PC node should be opened by default. If you have an open Explorer window and This PC is not the selected node, select the leftmost option button in the address list, and then select This PC.
2. Next, click the Map Network Drive button in the Computer panel, and then click Map Network Drive.
3. Use the Drive list to select a free drive letter to use, and then click the Browse button to the right of the Folder list. In the Browse For Folder dialog box, expand the network folders until you can select the name of the workgroup or the domain with which you want to work.
4. When you expand the name of a computer in a workgroup or a domain, you'll get



a list of shared folders. Select the shared folder with which you want to work, and then click OK.

5. Select Reconnect At Logon if you want Windows to connect to the shared folder automatically at the start of each session.
6. Click Finish. If the currently logged-on user doesn't have appropriate access permissions for the share, select Connect Using Different Credentials, and then click Finish. After you click Finish, you can enter the user name and password of the account with which you want to connect to the shared folder. Enter the user name in Domain\Username format, such as **Tvpress\Williams** . Before clicking OK, select Remember My Credentials if you want the credentials to be saved.

Otherwise, you'll need to provide credentials in the future.







## Disconnecting a Network Drive

In Windows Server 2016, you disconnect to a network drive using NET USE and Remove-PsDrive. The syntax for NET USE is:

```
net use DeviceName /delete
```

DeviceName specifies the network drive to remove, such as:

```
net use g: /delete
```

The syntax for Remove-PsDrive is:

```
Remove-PsDrive -Name DriveLetter
```

Where DriveLetter is the network drive to remove, such as:

```
Remove-PsDrive -Name g:
```

**NOTE** If the network drive has open connections, you can force remove the network drive using `-Force` parameter.

In File Explorer, you can disconnect a network drive by following these steps:

1. When you open File Explorer, the This PC node should be opened by default. If you have an open Explorer window and This PC is not the selected node, select the leftmost option button in the address list, and then select This PC.
2. Under Network Location, right-click the network drive icon, and then click Disconnect.

## Managing Existing Shares

As an administrator, you often have to manage shared folders. This section covers the common administrative tasks of managing shares.



# Understanding Special Shares

When you install Windows Server, the operating system creates special shares automatically. These shares are known as administrative shares and hidden shares, and they are designed to help make system administration easier. You can't set access permissions on automatically created special shares; Windows Server assigns access permissions. You can create your own hidden shares by adding the \$ symbol as the last character of the share name.

You can delete special shares temporarily if you're certain the shares aren't needed; however, the shares are re-created automatically the next time the operating system starts. To permanently disable the administrative shares, change the following registry values to 0 (zero):

- 
- 
- `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters`
- `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters`

Which special shares are available depends on your system configuration. A list of special shares you might find and how they're used follows:

- **ADMIN\$** A share used during remote administration of a system. It provides access to the operating system %SystemRoot%. On workstations and servers, administrators and backup operators can access these shares. On domain controllers, server operators also have access.
- **FAX\$** Supports network faxes. Used by fax clients when sending faxes.
- **IPC\$** Supports named pipes during remote interprocess communications (IPC) access. Used by programs when performing remote administration and when viewing shared resources.
- **NETLOGON** Supports the Net Logon service. Used by the Net Logon service when processing domain logon requests. Everyone has Read access.
- **PRINT\$** Supports shared printer resources by providing access to printer drivers. Used by shared printers. Everyone has Read access. Administrators, server operators, and printer operators have Full Control.
- **SYSVOL** Supports Active Directory. Used to store data and objects for Active Directory.
- **Driveletter \$** A share that allows administrators to connect to a drive's root folder. These shares are shown as C\$, D\$, E\$, and so on. On workstations and servers,

administrators and backup operators can access these shares. On domain controllers, server operators also have access.

## Connecting to Special Shares

Most special shares end with the \$ symbol. Although these shares aren't displayed in File Explorer, administrators and certain operators can connect to them (except for NETLOGON and SYSVOL). If your current logon account has appropriate permissions, you can connect directly to a special share or any standard share simply by typing the UNC path for the share in File Explorer's address box. The basic syntax is:


`\\ServerName\ShareName`

Where ServerName is the DNS name or IP address of the server and ShareName is the name of the share. In the following example, you connect to the D\$ share on CorpServer34:

`\\CorpServer34\D$`

If you always want the drive to be listed as a network location in This PC or need to specify credentials, you can connect to a special share by following these steps:

1. When you open File Explorer, the This PC node should be opened by default. If you have an open Explorer window and This PC is not the selected node, select the leftmost option button in the address list, and then select This PC.
2. Next, click the Map Network Drive button on the Computer panel, and then click Map Network Drive. This displays the Map Network Drive dialog box.



3. In the Drive list, select a free drive letter. This drive letter is used to access the special share.

- 
4. In the Folder text box, enter the Universal Naming Convention (UNC) path to the share. For example, to access the C\$ share on a server called Twiddle, you would use the path \\TWIDDLE\C\$.
  5. The Reconnect At Sign-In check box is selected automatically to ensure that the network drive is connected each time you log on. If you need to access the share only during the current logon session, clear this check box.
  6. If you need to connect to the share using different user credentials, select the Connect Using Different Credentials check box.
  7. Click Finish. If you are connecting using different credentials, enter the user name and password when prompted. Enter the user name in Domain\Username format, such as **Imaginedlands\Williams** . Before clicking OK, select Remember My Credentials if you want the credentials to be saved. Otherwise, you'll need to provide credentials in the future.

After you connect to a special share, you can access it as you would any other drive. Because special shares are protected, you don't have to worry about ordinary users accessing these shares. The first time you connect to the share, you might be prompted for a user name and password. If you are prompted, provide that information.





## Viewing User and Computer Sessions

You can use Computer Management to track all connections to shared resources on a Windows Server 2016 system. Whenever a user or computer connects to a shared resource, Windows Server 2016 lists a connection in the Sessions node.

To view connections to shared resources, enter **net session** at an administrator command prompt or Get-SMBSession at an administrator PowerShell prompt. You also can follow these steps:

1. In Computer Management, connect to the computer on which you created the shared resource.
2. In the console tree, expand System Tools, expand Shared Folders, and then select Sessions. You can now view connections to shares for users and computers.

The columns for the Sessions node provide the following important information about user and computer connections:

- **User** The names of users or computers connected to shared resources. Computer names are shown with a \$ suffix to differentiate them from users.
- **Computer** The name of the computer being used.
- **Type** The type of network connection being used.
- **# Open Files** The number of files with which the user is actively working. For more detailed information, access the Open Files node.
- **Connected Time** The time that has elapsed since the connection was established.
- **Idle Time** The time that has elapsed since the connection was last used.
- **Guest** Whether the user is logged on as a guest.

As shown in the following example, the output of Get-SMBSession provides the session ID, client computer name, client user name and the number of open files for each session:

| SessionId    | ClientComputerName | ClientUserName   | NumOpens |
|--------------|--------------------|------------------|----------|
| 834295421326 | 10.0.0.60          | TVPRESS\williams | 21       |

## Managing Sessions and Shares

Managing sessions and shares is a common administrative task. Before you shut down a server or an application running on a server, you might want to disconnect users from shared resources. You might also need to disconnect users when you plan to change





access permissions or delete a share entirely. Another reason to disconnect users is to break locks on files. You disconnect users from shared resources by ending the related user sessions.

## Ending Individual Sessions

To disconnect individual users from shared resources, enter **net session** \computername/**delete** at an administrator command prompt or Close-SMBSession at – ComputerName computername an elevated PowerShell prompt. In both instances, computername is the DNS name or IP address of computer from which the session originates. You also can disconnect users by following these steps:

1. In Computer Management, connect to the computer on which you created the share.
2. In the console tree, expand System Tools, expand Shared Folders, and then select Sessions.
3. Right-click the user sessions you want to end, and then click Close Session.
4. Click Yes to confirm the action.

## Ending All Sessions

To disconnect all users from shared resources, follow these steps:

1. In Computer Management, connect to the computer on which you created the share.
2. In the console tree, expand System Tools, expand Shared Folders, and then rightclick Sessions.
3. Click Disconnect All Sessions, and then click Yes to confirm the action.

**NOTE** Keep in mind that you're disconnecting users from shared resources, not from the domain. You can use only logon hours and Group Policy to force users to log off after they've logged on to the domain. Thus, disconnecting users doesn't log them off the network. It simply disconnects them from the shared resource.

To disconnect individual users from shared resources, enter **net session** \computername/**delete** at an administrator command prompt or Close-SMBSession at – ComputerName computername an elevated PowerShell prompt. In both instances, computername is the DNS name or IP address of computer from which the session originates. You also can disconnect users by following these steps:



You also can use Windows PowerShell to disconnect multiple users from a shared resource. The key here is to ensure you only close the sessions you want to close. Consider the following example:

```
ForEach-Object ($Session in (Get-SMBSession)) { Close-SMBSession -force}
```

This example uses a ForEach loop to get all active SMB sessions and then close each SMB session in turn. Thus, if you enter this example at an administrator PowerShell prompt, you will disconnect all users from all shared resources.

To close all connections only for a specific share, you must create a ForEach loop that only examines the connections for that share, such as:

```
ForEach-Object ($Session in (Get-SMBShare CorpData | Get-SMBSession)) {Close-SMBSession -force}
```

This example uses a ForEach loop to get all active SMB sessions for the CorpData share and then close each of those sessions in turn. Thus, if you enter this example at an administrator PowerShell prompt, you only disconnect users from the CorpData share.

## Managing Open Resources

Any time users connect to shares, the individual file and object resources they are working with are displayed in the Open Files node. The Open Files node might show the files the user has open but isn't currently editing.

You can access the Open Files node by following these steps:

1. In Computer Management, connect to the computer on which you created the share.
  2. In the console tree, expand System Tools, expand Shared Folders, and then select Open Files. This displays the Open Files node, which provides the following information about resource usage:
- **Open File** The file or folder path to the open file on the local system. The path might also be a named pipe, such as \PIPE\spools, which is used for printer spooling.
  - **Accessed By** The name of the user accessing the file.
  - **Type** The type of network connection being used.
  - **# Locks** The number of locks on the resource.
  -



**Open Mode** The access mode used when the resource was opened, such as read, write, or write+read.

You also can use Get-SMBOpenFile to list open files. As shown in the following example, Get-SMBOpenFile provides the file ID, session ID, path, share relative path, client computer name, and client user name for each open file:

| FileId       | SessionId    | Path             | ShareRelativePath | ClientComputerName | ClientUserName     |
|--------------|--------------|------------------|-------------------|--------------------|--------------------|
| 601295424973 | 601295421497 | C:\PrimaryData\  |                   | 10.0.0.60          | TVPRESS\williams   |
| 601295425045 | 601295421577 | C:\Windows\SYVOL | tv...             | 10.0.0.60          | TVPRESS\CORPPC29\$ |

### Closing an Open File

To close an open file on a computer’s shares, follow these steps:

1. In Computer Management, connect to the computer with which you want to work.
2. In the console tree, expand System Tools, expand Shared Folders, and then select Open Files.
3. Right-click the open file you want to close, and then click Close Open File.
4. Click Yes to confirm the action.

You also can use Close-SMBOpenFile to close open files. When you close a file, you use the –FileID parameter to specify the identifier for the file to close, such as:

```
Close-SMBOpenFile –FileID 601295424973
```

Add the –Force parameter to force close the file if needed. However, if the file has been modified by a user, any changes to the file could be lost.

### Closing All Open Files

To close all open files on a computer’s shares, follow these steps:

1. In Computer Management, connect to the computer on which the share is created.
2. In the console tree, expand System Tools, expand Shared Folders, and then rightclick Open Files.
3. Click Disconnect All Open Files, and then click Yes to confirm the action.

You also can use Windows PowerShell to close all open files on a computer’s share. The key here is to ensure you only close the files you want to close. Consider the following example:

```
ForEach-Object ($Session in (Get-SMBOpenFile)) { Close-SMBOpenFile –force}
```







This example uses a `ForEach` loop to get all open SMB files and then close each SMB file in turn. Thus, if you enter this example at an administrator PowerShell prompt, you will close all open files for all shared resources.

To close open files on a specific share, you must create a `ForEach` loop that only examines the open files for that share, such as:

```
ForEach-Object ($Session in (Get-SMBShare CorpData | Get-SMBOpenFile)) {Close-SMBOpenFile -force}
```

This example uses a `ForEach` loop to get all open SMB files for the `CorpData` share and then close each of those files in turn. Thus, if you enter this example at an administrator PowerShell prompt, you only close open files for the `CorpData` share.

## Stopping File and Folder Sharing

To stop sharing a folder, follow these steps:

1. Do one of the following:
  - In Server Manager, select the share you want to manage on the Shares subnode of the File And Storage Services node.
  - In Computer Management, connect to the computer on which you created the share, and then access the Shares node.
2. Right-click the share you want to remove, click Stop Sharing, and then click Yes to confirm the action.

**CAUTION** You should never delete a folder containing shares without first stopping the shares. If you fail to stop the shares, Windows Server 2016 attempts to reestablish the shares the next time the computer is started, and the resulting error is logged in the system event log.






# Configuring NFS Sharing

To enable NFS sharing, you can install Server For NFS as a role service on a file server. Server For NFS provides a file sharing solution for enterprises with mixed Windows, OS X, and UNIX environments, allowing users to transfer files between Windows Server 2016, OS X, and UNIX operating systems by using the NFS protocol.

You can configure NFS sharing using either File Explorer or Server Manager. In File Explorer, follow these steps to enable and configure NFS sharing:

1. Right-click the share you want to manage, and then click Properties to display a Properties dialog box for the share.
2. On the NFS Sharing tab, click Manage NFS Sharing.
3. In the NFS Advanced Sharing dialog box, select the Share This Folder check box.
4. In the Share Name text box, enter a name for the share. This is the name of the folder to which UNIX users will connect. NFS share names must be unique for each system and can be the same as those used for standard file sharing.
5. ANSI is the default encoding for text associated with directory listings and file names. If your UNIX computers use a different default encoding, you can choose that encoding in the Encoding list.





6. UNIX computers use Kerberos v5 authentication by default. Typically, you want to allow Kerberos integrity and authentication in addition to standard Kerberos authentication. Select the check boxes for the authentication mechanisms you want to use. Clear the check boxes for those you don't want to use.
7. The share can be configured so that no server authentication is required. If you want to require server authentication, select the No Server Authentication check box, and then choose additional options as appropriate. Unmapped user access can be allowed and enabled. If you want to allow anonymous access to the NFS share, select the Allow Anonymous Access option, and then enter the anonymous user UID and anonymous group GID.
8. For UNIX computers, you configure access primarily based on the computer names (also referred to as host names ). By default, no UNIX computers have access to the NFS share. If you want to grant read-only or read/write permissions, click Permissions, set the permissions you want to use in the NFS Share Permissions dialog box, and then click OK. You can configure no access, read-only access, or read/write access by client computer name and client computer groups.

9. Click OK twice to close the open dialog boxes and save your settings.

In File Explorer, you can disable NFS sharing by following these steps:

1. Right-click the share you want to manage, and then click Properties. This displays a Properties dialog box for the share.
2. On the NFS Sharing tab, click Manage NFS Sharing.
3. In the NFS Advanced Sharing dialog box, clear the Share This Folder check box, and then click OK twice.

With Server Manager, you can configure NFS permissions as part of the initial share configuration when you are provisioning a share. On the Shares subnode of the File And Storage Services node, you can create an NFS share by following these steps:

1. In the Shares pane, click Tasks, and then click New Share to start the New Share Wizard. Choose NFS Share—Quick or NFS Share—Advanced as the share profile, and then click Next.
2. Specify the share name and location as you would for an SMB share.
3. On the Specify Authentication Methods page, configure Kerberos v5 Authentication and No Server Authentication. The options provided are similar to those discussed previously in this section.
4. On the Specify Share Permissions page, configure access for UNIX hosts. Hosts can be set for no access, read-only access, or read/write access to the share.
5. On the Specify Permissions To Control Access, optionally set NTFS permissions for the share.
6. On the Confirm Selections page, review your selections. When you click Create, the wizard creates the share, configures it, and sets permissions. The status should state, “The share was successfully created.” If an error is displayed instead, note the error and take corrective action. However, because typical errors relate to configuring host access, you probably won’t need to repeat this procedure to create the share. Instead, you might need to modify only the share permissions. Click Close.



# Chapter 26. Using Shadow Copies and Work Folders

Shadow Copies and Work Folders are extensions to file sharing. Any time your organization uses shared folders, you should consider creating shadow copies of these shared folders as well. Shadow copies are point-in-time backups of data files that users can access directly in shared folders.

Point-in-time backups can save you and the other administrators in your organization a lot of work, especially if you routinely have to retrieve lost, overwritten, or corrupted data files from backups. The normal procedure for retrieving shadow copies is to use the Previous Versions or Shadow Copy client. Windows Server 2016 includes a feature enhancement that enables you to revert an entire (nonsystem) volume to a previous shadow copy state.

Although the standard approach to sharing files requires a computer that is joined and connected to a domain, synced sharing does not. With sync shares, users can use an Internet or corporate network connection to sync data to their devices from folders located on enterprise servers. You implement synced sharing by using Work Folders.

Work Folders is a feature that you can add to servers running Windows Server 2016 or later. Work Folders use a client-server architecture. A Work Folders client is natively integrated into Windows 10 and is also available for iPad and other devices.

## Using Shadow Copies

You can create shadow copies only on NTFS volumes. You use the Shadow Copy feature to create automatic backups of the files in shared folders on a per-volume basis. For example, on a file server that has three NTFS volumes, each containing shared folders, you need to configure this feature for each volume separately.


If you enable this feature in its default configuration, shadow copies are created twice each weekday (Monday–Friday) at 7:00 A.M. and 12:00 P.M. You need at least 300 megabytes (MB) of free space to create the first shadow copy on a volume. The total disk space used beyond this depends on the amount of data in the volume's shared folders. You can restrict the total amount of disk space used by Shadow Copy by setting the allowable maximum size of the point-in-time backups.



# Creating Shadow Copies

When shadow copies are enabled on a disk, you configure and view current Shadow Copy settings via the disk's Properties dialog box. In File Explorer or Computer Management, right-click the icon for the disk with which you want to work, click Properties, and then click the Shadow Copies tab. The Select A Volume panel shows the following:

- **Volume** The volume label of NTFS volumes on the selected disk drive
- **Next Run Time** The status of Shadow Copy as Disabled, or the next time a shadow copy of the volume will be created
- **Shares** The number of shared folders on the volume
- **Used** The amount of disk space used by Shadow Copy



Individual shadow copies of the currently selected volume are listed in the Shadow Copies Of Selected Volume panel by date and time.

Disks have the Shadow Copies tab only when shadow copies are enabled. To enable shadow copies on a disk and create the first shadow copy, follow these steps:

1. In Computer Management, right-click the Disk Management node, point to All Tasks, and then click Configure Shadow Copies.
2. In the Shadow Copies dialog box, volumes listed with a Disabled status in the





Next Run Time column haven't been configured to use shadow copies. Select a disabled volume by clicking it.

3. Click Settings. In the Settings dialog box, click Details to display information about the total and free disk space on the selected volume. Then close this dialog box by clicking OK. Next, configure the maximum size of all shadow copies for this volume. For maximum size, you can set an exact limit or specify that there is no size limit. Typically, the default maximum size is approximately 10% of the total disk space available.
4. In the Settings dialog box, click Schedule. As discussed previously, two default schedules are defined. If you want to change the default schedule, select a schedule in the Schedule list and then change its configuration using the options provided. Click OK to close the Schedule dialog box, then click OK to close the Settings dialog box.







5. After you configure the volume for shadow copying, click Enable if necessary. When prompted to confirm this action, click Yes. Enabling shadow copying creates the first shadow copy and sets the schedule for later shadow copies. Click OK to close the Shadow Copies dialog box.

**NOTE** If you create a run schedule when configuring the Shadow Copy settings, shadow copying is enabled automatically for the volume when you click OK to close the Settings dialog box. However, the first shadow copy won't be created until the next scheduled run time. If you want to create a shadow copy of the volume now, select the volume and then click Create Now.

Once you've enabled shadows copies on a disk, the copies are created automatically according to the schedule defined. You can create a shadow copy at any time by following these steps:





1. In Computer Management, right-click the Disk Management node, point to All Tasks, and then click Configure Shadow Copies.
2. In the Shadow Copies dialog box, select the volume with which you want to work by clicking it.
3. Click Create Now.





## Restoring a Shadow Copy

Users working on client computers access shadow copies of individual shared folders by using the Previous Versions or Shadow Copy client. The best way to access shadow copies on a client computer is to follow these steps:

1. In File Explorer, right-click the share for which you want to access previous file versions, click Properties, and then click the Previous Versions tab.
  2. On the Previous Versions tab, select the folder version with which you want to work. Each folder has a date and time stamp. Click the button corresponding to the action you want to perform:
- Click Open to open the shadow copy in File Explorer.
  - Click Copy to display the Copy Items dialog box, which lets you copy the snapshot image of the folder to the location you specify.
  - Click Restore to roll back the shared folder to its state at the time of the snapshot image you selected.





## Reverting an Entire Volume

Windows Server 2016 features a shadow copy enhancement that enables you to revert an entire volume to the state it was in when a particular shadow copy was created. Because volumes containing operating system files can't be reverted, the volume you want to revert must not be a system volume. The same goes for volumes on a cluster shared disk.

To revert an entire volume to a previous state, follow these steps:

1. In Computer Management, right-click Disk Management, point to All Tasks, and then click Configure Shadow Copies.
2. On the Shadow Copies tab, select the volume with which you want to work in the Select A Volume list.
3. Individual shadow copies of the currently selected volume are listed by date and time in the Shadow Copies Of Selected Volume panel. Select the shadow copy with the date and time stamp to which you want to revert, and then click Revert.
4. To confirm this action, select the Check Here If You Want To Revert This Volume check box, and then click Revert Now. Click OK to close the Shadow Copies dialog box.





## Deleting Shadow Copies

Each point-in-time backup is maintained separately. You can delete individual shadow copies of a volume as necessary, and this recovers the disk space used by the shadow copies.

To delete a shadow copy, follow these steps:

1. In Computer Management, right-click Disk Management, point to All Tasks, and then click Configure Shadow Copies.
2. On the Shadow Copies tab, select the volume with which you want to work in the Select A Volume list.
3. Individual shadow copies of the currently selected volume are listed by date and time in the Shadow Copies Of Selected Volume panel. Select the shadow copy you want to delete, and then click Delete Now. Click Yes to confirm the action.





## Disabling Shadow Copies

If you no longer want to maintain shadow copies of a volume, you can disable the Shadow Copy feature. Disabling this feature turns off the scheduling of automated point-in-time backups and removes any existing shadow copies.

To disable shadow copies of a volume, follow these steps:

1. In Computer Management, right-click Disk Management, point to All Tasks, and then click Configure Shadow Copies.
2. On the Shadow Copies tab, select the volume with which you want to work in the Select A Volume list, and then click Disable.
3. When prompted, confirm the action by clicking Yes. Click OK to close the Shadow Copies dialog box.





# Configuring Synced Sharing

Although the standard approach to sharing files requires a computer that is joined and connected to a domain, synced sharing does not. With sync shares, users can use an Internet or corporate network connection to sync data to their devices from folders located on enterprise servers. You implement synced sharing by using Work Folders.

Work Folders is a feature that you can add to servers running Windows Server 2016. Work Folders use a client-server architecture. A Work Folders client is natively integrated into Windows 10 and is also available for iPad and other devices.





# Getting Started with Work Folders

You deploy Work Folders in the enterprise by performing these procedures:

1. Add the Work Folders role to servers that you want to host sync shares.
2. Use Group Policy to enable discovery of Work Folders.
3. Create sync shares on your sync servers and optionally, enable SMB access to sync shares.
4. Configure clients to access Work Folders.

**NOTE** Group Policy is discussed in detail in Chapter 29. For detailed information about configuring Group Policy to enable discovery of Work Folders, see “Automatically Configuring Work Folders,” in that chapter.

Work Folders use a remote web gateway configured as part of the IIS hostable web core. When users access a sync share via a URL provided by an administrator and configured in Group Policy, a user folder is created as a subfolder of the sync share and this subfolder is where the user’s data is stored. The folder naming format for the userspecific folder is set when you create a sync share. The folder can be named by using only the user alias portion of the user’s logon name or the full logon name in alias@domain format. The format you choose primarily depends on the level of compatibility required. Using the full logon name eliminates potential conflicts when users from different domains have identical user aliases, but this format is not compatible with redirected folders.

To maintain compatibility with redirected folders, you should configure sync folders to use aliases. However, in enterprises with multiple domains, the drawback to this approach is that there could be conflicts between identical user aliases in different domains. Although the automatically configured permissions for a user folder would prevent amyh from the imaginedlands.com domain from accessing a user folder created for amyh from the pocket-consultant.com domain, the conflict would cause problems. If there was an existing folder for amyh from the imaginedlands.com domain, the server would not be able to create a user folder for amyh from the pocket-consultant.com.

With Work Folders, you have several important options during initial setup. You can encrypt files in Work Folders on client devices and ensure that the screens on client devices lock automatically and require an access password. Encryption is implemented by using the Encrypting File System (EFS). EFS encrypts files with an enterprise encryption key rather than an encryption key generated by the client device. The enterprise encryption key is specific to the enterprise ID of the user (which by default





is the primary SMTP address of the user). Having an enterprise encryption key that is separate from a client's standard encryption key is important to ensure that encrypted personal files and encrypted work files are managed separately.

When files are encrypted, administrators can use a selective wipe to remove enterprise files from a client device. The selective wipe simply removes the enterprise encryption key and thus renders the work files unreadable. Selective wipe does not affect any encrypted personal files. As the work files remain encrypted, there's no need to actually delete the work files from the client device. That said, you could run Disk Optimizer on the drive where the work files were stored. During optimization, Disk Optimizer should then overwrite the sectors where the work files were stored. Selective wipe only works when you've enabled the encryption option on Work Folders.

Although encryption is one way to protect enterprise data, another way is to configure client devices to lock screens and require a password for access. The exact policy enforced requires:

- [A minimum password length of 6 characters](#)
- [A maximum password retry of 10](#)
- [A screen that automatically locks in 15 minutes or less](#)

If you enforce the use of automatic lock screens and passwords, any device that doesn't support these requirements is prevented from connecting to the Work Folder.

By default, sync shares are not available in the same way as standard file shares. Because of this, users can only access sync shares by using the Work Folders client. If you want to make sync shares available to users as standard file shares, you must enable SMB access. After you enable SMB access, users can access files stored in Work Folders by using syncing and by mapping network drives.

When a user makes changes to files in Work Folders, the changes might not be immediately apparent to others using the same Work Folders. For example, if a user deletes a file from a Work Folder by using SMB, other users accessing the Work Folder might still see the file as available. This inconsistency can occur because by default clients only poll the sync server every 10 minutes for SMB changes.

A sync server also uses a Work Folders client to check periodically for changes users have made using SMB; the default polling interval is 5 minutes. When the server identifies changes, the server relays the changes the next time a client syncs. Following this, you can determine that it could take up to 15 minutes for a change made using SMB to fully propagate.







**REAL WORLD** To minimize support issues related to Work Folders, you'll want to let readers know how the technology works. Specifically, you'll want to let readers know changes might not be immediately apparent and they'll need to be patient when waiting for changes to propagate.

You can specify how frequently the server checks for changes made locally on the server or through SMB by using the `-MinimumChangeDetectionMins` parameter of the `Set-SyncServerSetting` cmdlet. However, as the server must check the change information for each file stored in the sync share, you need to be careful that you don't configure a server to try to detect changes too frequently. A server that checks for changes too frequently can become overloaded. Remember, change detection uses more resources as the number of files stored in the sync share increases.

If you deploy roles and features that require a full version of the Web (IIS) role, you might find that these roles and features or the Work Folders feature itself don't work together. A conflict can occur because the full version of the Web (IIS) role has a Default Web Site that uses port 80 for HTTP communications and port 443 for secure HTTP communications. For example, running Windows Essentials Experience and Work Folders together on the same server requires a special configuration. Typically, you need to change the ports used by Windows Essentials Experience so that they don't conflict with the ports used by Work Folders.


To enable detailed logging of Work Folders, you can enable and configure the Audit Object Access policy setting for a Group Policy Object (GPO) processed by the server. You'll find this setting in the Administrative Templates for Computer Configuration under Windows Settings\Security Settings\Local Policies Audit Policies. After you enable Audit Object Access, add an audit entry for the specific folders you want to audit. In File Explorer, right-click a folder you want to audit, and then select Properties. In the Properties dialog box, on the Security tab, select Advanced. In the Advanced Security Settings dialog box, use the options on the Auditing tab to configure auditing.

## **Creating Sync Shares and Enabling SMB Access**

You create a sync share to identify a local folder on a sync server that will be synchronized and accessible to domain users via the Work Folders client. As sync shares are mapped to local paths on sync servers, I recommend that you create any folders that you want to use before creating sync shares. This will make it easier to select the exact folders with which you want to work. For details on adding the Work Folders role and configure Work Folders in Group Policy, see "Automatically

Configuring Work Folders” in Chapter 29 “Using Group Policy for Administration.” To create a sync share, complete the following steps:

1. In Server Manager, select File And Storage Services, and then select Work Folders. On the Work Folders panel, select Tasks, and then select New Sync Share to open the New Sync Share Wizard. If the Before You Begin page is displayed, click Next.
2. On the Select The Server And Path page, select the server with which you want to work. Keep in mind that only servers that have the Work Folders role installed are available for selection.




3. When configuring sync shares, you have several options. You can:
  - Add syncing to an existing file share by choosing the Select By File Share option, and then selecting the file share that should also be synced.
  - Add syncing to an existing local folder by choosing Enter A Local Path, selecting Browse, and then using the Select Folder dialog box to locate and chose the folder to sync.
  - Add syncing to a new local folder by choosing Enter A Local Path, and then entering the path to use.
4. When you are ready to continue, click Next. If you specified a new folder location, you are prompted to confirm whether you want to create this folder. Select OK to create the folder and continue.



5. On the Specify The Structure For User Folders page, choose a folder naming format for the subfolders where user data is stored. To use only the user alias portion of the user's logon name for naming user folders, choose User Alias. To use the full logon name for naming user folders, choose User alias@domain.




6. By default, all folders and files stored under the user folder are synced automatically. If you'd prefer that only a specific folder is synced, select the Sync Only The Following Folder check box, and then enter the name of the folder, such as Documents. Click Next to continue
7. On the Enter The Sync Share Name page, enter a share name and description before clicking Next to continue.




8. On the Grant Sync Access To Groups page, use the options provided to specify the users and groups that should be able to access the sync share. To add a user

or group, click Add, and then use the Select User Or Group dialog box to specify the user or group that should have access to the sync share.

**SECURITY ALERT** Any users and groups you specify will be granted permissions on the base folder that allows the users and groups to create folders and access files in their folders. Specifically, Creator/Owner is granted Full Control on subfolders and files only. The users and groups are granted List Folder/Read Data, Create Folders/Append Data, Traverse Folder/Execute File, Read/Write attributes on the base folder. Local System is granted Full Control of the base folder, subfolders, and files. Administrator is granted Read permission on the base folder.



9. By default, inherited permissions are disabled and users have exclusive access to their user folders. Because of this, only the user who stores a file has access to this file on the share. If the base folder for the share has permissions that you want to be applied to user folders, such as those that would grant administrators access to user folders, clear the Disable Inherited Permissions check box. When you are ready to continue, click Next.
10. On the Specify Security Policies page, you have two options. You can select Encrypt Work Folders to encrypt files in Work Folders on client devices. You can select Automatically Lock Screen And Require A Password to ensure that the screens on client devices lock automatically and require a password for access.



11. Click Next to continue, and then confirm your selections. Select Create to create the sync share. If the wizard is unable to create the sync share, you'll get an error and will need to note the error and take appropriate corrective action. A common error you might get occurs when the server hosts both Work Folders (which use the hostable web core) and the full Web (IIS) role. Before you can create sync shares, you'll need to modify the ports used so they do not conflict or install Work Folders on a server that doesn't have the full Web (IIS) role.
12. If you did not select an existing file share during set up and want to enable the sync share for SMB access, open File Explorer. In File Explorer, right-click the folder, select Share With, and then select Specific People. Finally, configure file sharing as discussed earlier in this chapter.

## Accessing Work Folders on Clients

Users with a domain user account can access Work Folders from a client device over the Internet or over the corporate network. You can configure Work Folder Access for a user by completing the following steps:

1. In Control Panel, click System And Security, and then select Work Folders. On the Manage Work Folders page, click Set Up Work Folders.
2. On the Enter Your Work Email Address page, enter the user email address, such as **amyh@imaginedlands.com**, and then click Next. If the client device is joined to the domain, you will not be prompted for the user's credentials. Otherwise, you are prompted for the user's credentials. After the user enters her credentials, you can select Remember My Credentials to store the user's credentials for future use, and then click OK to continue.
3. On the Introducing Work Folders page, note where the work files for the user will be stored. By default, work files are stored in a user profile subfolder called Work Folders. For example, the work files for Amyh would be stored under %SystemDrive%\Users\Amyh\WorkFolders. To store work files in another



location, click Change and then use the options provided to specify a new save location for work files. When you are ready to continue, click Next.

4. On the Security Policies page, review the security policies that will be applied, and then have the user select the I Accept These Policies On My PC check box. You will not be able to continue if you do not select this check box.
5. Select Set Up Work Folders to create Work Folders on the client device.

After you configure Work Folders for initial use on a client device, the user can access Work Folders in File Explorer. When a user opens File Explorer, the This PC node should be opened by default. If so, the user simply needs to double-click Work Folders to view work files. If a user has an open Explorer window and This PC is not the selected node, she simply needs to click the leftmost option button in the address list, and then click This PC.

As the user works with files, the changes the user makes trigger sync actions with the server. If the user doesn't change any files locally for an extended period of time, the client connects to the server every 10 minutes to determine whether there are changes to sync.

## **Chapter 27. Managing Permissions and Auditing**

Data is the heart of any enterprise and few aspects of administration are more important than ensuring that data is protected. Although file and folder permissions protects important resources by restricting access, protecting enterprise data isn't just about file and folder permissions. To secure enterprise data appropriately, you need a firm understanding of object management, ownership, inheritance, and auditing. To help ensure that enterprise data is manageable, you also need to know how to implement quotas that restrict the amount of data that can be stored on servers.



# Object Management, Ownership, and Inheritance

Windows Server 2016 takes an object-based approach to describing resources and managing permissions. Objects that describe resources are defined on NTFS volumes and in Active Directory Domain Services (AD DS). With NTFS volumes, you can set permissions for files and folders. With Active Directory, you can set permissions for other types of objects, such as users, computers, and groups. You can use these permissions to control access with precision.







# Objects and Object Managers

Whether defined on an NTFS volume or in Active Directory, each type of object has an object manager and primary management tools. The object manager controls object settings and permissions. The primary management tools are the tools of choice for working with the object. Objects, their managers, and management tools are summarized in Table 27-1.

**TABLE 27-1** Windows Server 2016 objects

| OBJECT TYPE       | OBJECT MANAGER      | MANAGEMENT TOOL                                                  |
|-------------------|---------------------|------------------------------------------------------------------|
| Files and folders | NTFS                | File Explorer                                                    |
| Printers          | Print spooler       | Printers in Control Panel                                        |
| Registry keys     | Windows registry    | Registry Editor                                                  |
| Services          | Service controllers | Security Configuration Tool Set                                  |
| Shares            | Server service      | File Explorer, Computer Management, Share And Storage Management |

## Object Ownership and Transfer

It’s important to understand the concept of object ownership. In Windows Server 2016, the object owner isn’t necessarily the object’s creator; instead, the object owner is the person who has direct control over the object. Object owners can grant access permissions and give other users permission to take ownership of the object.

As an administrator, you can take ownership of objects on the network to ensure that you can’t be locked out of files, folders, printers, and other resources. After you take ownership of files, however, you can’t return ownership to the original owner (in most





cases). This prevents administrators from accessing files and then trying to hide the fact.

The way ownership is assigned initially depends on the location of the resource being created. In most cases, the Administrators group is listed as the current owner, and the object's actual creator is listed as a person who can take ownership.

Ownership can be transferred in several ways:

- If the Administrators group is initially assigned as the owner, the creator of the object can take ownership, if she does this before someone else takes ownership.
- The current owner can grant the Take Ownership permission to other users, allowing those users to take ownership of the object.
- An administrator can take ownership of an object, if the object is under his administrative control.

To take ownership of an object, follow these steps:

1. Open the management tool for the object. For example, if you want to work with files and folders, start File Explorer.
2. Right-click the object you want to take ownership of, and then click Properties. In the Properties dialog box, click the Security tab.
3. On the Security tab, click Advanced to display the Advanced Security Settings dialog box where the current owner is listed under the file or folder name.
4. Click Change. Use the options in the Select User, Computer, Service Account, Or Group dialog box to select the new owner.
5. Click OK twice when you have finished.

**TIP** If you're taking ownership of a folder, you can take ownership of all subfolders and files within the folder by selecting the Replace Owner On Subcontainers And Objects check box. This option also works with objects that contain other objects, in which case you would take ownership of all child objects.

## Object Inheritance

Objects are defined by using a parent-child structure. A parent object is a top-level object, and a child object is an object defined below a parent object in the hierarchy. For example, the folder C:\ is the parent of the folders C:\Data and C:\Backups. Any subfolders created in C:\Data or C:\Backups are children of these folders and grandchildren of C:\.



Child objects can inherit permissions from parent objects; in fact, all Windows Server 2016 objects are created with inheritance enabled by default. This means that child objects automatically inherit the permissions of the parent; therefore, the parent object permissions control access to the child object. If you want to change permissions on a child object, you must do one of the following:

- Edit the permissions of the parent object.
- Stop inheriting permissions from the parent object, and then assign permissions to the child object.
- Select the opposite permission to override the inherited permission. For example, if the parent allows the permission, you would deny it on the child object.

To stop inheriting permissions from a parent object, follow these steps:

1. Open the management tool for the object. For example, if you want to work with files and folders, start File Explorer.
2. Right-click the object with which you want to work, and then click Properties. In the Properties dialog box, click the Security tab.
3. Click Advanced to display the Advanced Security Settings dialog box.
4. On the Permissions tab, click Change Permissions to display an editable version of the Permissions tab.
5. On the Permissions tab, you'll see a Disable Inheritance button if inheritance currently is enabled. Click Disable Inheritance.
6. You can now either convert the inherited permissions to explicit permissions or remove all inherited permissions and apply only the permissions that you explicitly set on the folder or file.

Keep in mind that if you remove the inherited permissions and no other permissions are assigned, everyone but the owner of the resource is denied access. This effectively locks out everyone except the owner of a folder or file; however, administrators still have the right to take ownership of the resource regardless of the permissions. Thus, if an administrator is locked out of a file or a folder and truly needs access, she can take ownership and then have unrestricted access.

To start inheriting permissions from a parent object, follow these steps:

1. Open the management tool for the object. For example, if you want to work with files and folders, start File Explorer.
  2. Right-click the object with which you want to work, and then click Properties. In the Properties dialog box, click the Security tab.
  3. Click Advanced to display the Advanced Security Settings dialog box.
- 



4. On the Permissions tab, click Enable Inheritance, and then click OK. Note that the Enable Inheritance button is available only if permission inheritance currently is disabled.





# File and Folder Permissions

NTFS permissions are always evaluated when a file is accessed. On NTFS and ReFS volumes, you can set security permissions on files and folders to grant or deny access to the files and folders. Because Windows Server 2016 adds new layers of security, NTFS permissions now encompass the following:

- [Basic permissions](#)
- [Claims-based permissions](#)
- [Special permissions](#)

You can view NTFS permissions for files and folders by following these steps:

1. In File Explorer, right-click the file or folder with which you want to work, and then click Properties. In the Properties dialog box, click the Security tab.
2. In the Group Or User Names list, select the user, computer, or group whose permissions you want to view. If the permissions are not available (dimmed), the permissions are inherited from a parent object.

Shared folders have both share permissions and NTFS permissions. You can view the underlying NTFS permissions for shared folders by following these steps:

1. In Server Manager, the Shares subnode of the File And Storage Services node shows existing shares for file servers that have been added for management.
2. Right-click the folder with which you want to work, and then click Properties to display a Properties dialog box.
3. When you click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
4. To get more information, click Customize Permissions to open the Advanced Security Settings dialog box.

On file servers running Windows Server 2016, you can also use central access policies to precisely define the specific attributes that users and devices must have to access resources.

## Understanding File and Folder Permissions

The basic permissions you can assign to files and folders are summarized in the list that follows. File permissions include Full Control, Modify, Read & Execute, Read, and Write. Folder permissions include Full Control, Modify, Read & Execute, List Folder Contents, Read, and Write.



- **Read** For folders, permits viewing and listing files and subfolders. For files, permits viewing or accessing a file's contents.
- **Write** For folders, permits adding files and subfolders. For files, permits writing to a file.
- **Read & Execute** For folders, does not permit viewing the contents of files. You can list file and folder names, but you can't open files to read, nor can you execute files if that execute requires opening the file (as in a batch or PS1 file). Inherited by files and folders. For files, permits viewing and accessing a file's contents in addition to executing a file.
- **List Folder Contents** For folders, permits viewing and listing file names and subfolder names in addition to executing files; inherited by folders only. Not applicable to files.
- **Modify** For folders, permits reading and writing of files and subfolders; allows deletion of the folder. For files, permits reading and writing of a file; allows deletion of a file.
- **Full Control** For folders, permits reading, writing, changing, and deleting files and subfolders. For files, permits reading, writing, changing, and deleting a file.

Any time you work with file and folder permissions, you should keep the following in mind:

- Read is the only permission needed to run scripts. Execute permission doesn't matter.
- Read access is required to access a shortcut and its target.

Giving a user permission to write to a file but not to delete it doesn't prevent the user from deleting the file's contents.

- If a user has full control over a folder, the user can delete files in the folder regardless of the permission on the files.

The basic permissions are created by combining special permissions in logical groups. By using advanced permission settings, you can assign these special permissions individually, if necessary. As you study the special permissions, keep the following in mind:

- By default, if no access is specifically granted or denied, the user is denied access. Further, if a permission has been explicitly denied, the deny will override any permission grant.
- Actions that users can perform are based on the sum of all the permissions assigned to the user and to all the groups of which the user is a member. For example, if the



user GeorgeJ has Read access and is a member of the group Techies, which has Change access, GeorgeJ will have Change access. If Techies is a member of Administrators, which has Full Control, GeorgeJ will have complete control over the file. However, if GeorgeJ has been explicitly denied a permission, the deny will override any grant.


As you work with special permissions, keep in mind that when you create files and folders, these files and folders inherit certain permission settings from parent objects. These permission settings are shown as the default permissions.



# Setting Basic File and Folder Permissions

To set basic NTFS permissions for files and folders, follow these steps:

1. In File Explorer, right-click the file or folder with which you want to work, and then click Properties. In the Properties dialog box, click the Security tab.
2. Click Edit to display an editable version of the Security tab.



3. Users or groups that already have access to the file or folder are listed in the Group Or User Names list. You can change permissions for these users and groups by selecting the user or group you want to change and then granting or denying access permissions in the Permissions list box.

**TIP** Inherited permissions are shaded (dimmed). If you want to override an inherited permission, select the opposite permission.

4. To set access permissions for additional users, computers, or groups, click Add to display the Select Users, Computers, Service Accounts, Or Groups dialog box.
5. Enter the name of a user, computer, or group in the current domain, and then click Check Names. One of the following actions occurs:

- If a single match is found, the dialog box is updated and the entry is underlined.
- If no matches are found, you entered an incorrect name part or are working with an incorrect location. Modify the name and try again, or click Locations to select a new location.
- If multiple matches are found, select the name or names you want to use, and then click OK. To add more users, computers, or groups, enter a semicolon ( ; ), and then repeat this step.

**NOTE** The Locations button allows you to access account names in other domains. Click Locations to view a list of the current domain, trusted domains, and other resources you can access. Because of the transitive trusts in Windows Server 2016, you can usually access all the domains in the domain tree or forest.

6. In the Group Or User Names list, select the user, computer, or group you want to configure, and in the check boxes in the Permissions list, allow or deny permissions. Repeat for other users, computers, or groups.
7. Click OK.

Because shared folders also have NTFS permissions, you might want to set basic NTFS permissions by using Server Manager. To do this, follow these steps:

1. In Server Manager, select File and Storage Services, select the server you want to work with, and then Select Shares.
2. Right-click the folder with which you want to work, and then click Properties to display a Properties dialog box.
3. When you click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
4. Click Customize Permissions to open the Advanced Security Settings dialog box with the Permissions tab selected.
5. Users or groups that already have access to the file or folder are listed under Permission Entries. Use the options provided to view, edit, add, or remove permissions for users and groups.


## Setting Special Permissions on Files and Folders

To set special NTFS permissions for files and folders, follow these steps:

1. In File Explorer, right-click the file or folder with which you want to work, and then click Properties.
2. In the Properties dialog box, select the Security tab, and then click Advanced to display the Advanced Security Settings dialog box. Before you can modify




permissions, you must click Change Permissions. The permissions are presented much as they are on the Security tab. The key differences are that you see individual allow or deny permission sets, whether permissions are inherited and where they are from, and the resources to which the permissions apply.



3. If a user or group already has permissions set for the file or folder (and those permissions are not being inherited), you can modify the special permissions by selecting the user or group, and then clicking Edit. Afterward, skip steps 4–7, and then follow the rest of the steps in this procedure.
4. To add special permissions for a user or group, click Add to display the Permission Entry dialog box. Click Select A Principal to display the Select User, Computer, Service Account, Or Group dialog box.
5. Enter the name of a user or a group account. Be sure to reference the user account name rather than the user's full name. Only one name can be entered at a time.
6. Click Check Names. If a single match is found for each entry, the dialog box is automatically updated and the entry is underlined; otherwise, you'll get an additional dialog box. If no matches are found, you either entered the name incorrectly, or you're working with an incorrect location. Modify the name in the Name Not Found dialog box and try again, or click Locations to select a new location. When multiple matches are found, in the Multiple Names Found dialog box, select the name you want to use, and then click OK.

7. Click OK. The user and group are added as the Principal, and the Permission Entry dialog box is updated to show this.
8. When you are editing permissions, only basic permissions are listed by default. Click Show Advanced Permissions to display the special permissions.



9. Use the Type list to specify whether you are configuring allowed or denied special permissions, and then select the special permissions you want to allow or deny. If any permissions are dimmed (unavailable), they are inherited from a parent folder.

**NOTE** You allow and deny special permissions separately. Therefore, if you want to both allow and deny special permissions, you need to configure the allowed permissions, and then repeat this procedure starting with step 1 to configure the denied permissions.

10. If the options in the Applies To list are available, choose the appropriate option to ensure that the permissions are properly inherited. The options include the following:
  - **This Folder Only** The permissions apply only to the currently selected folder.
  - **This Folder, Subfolders And Files** The permissions apply to this folder, any subfolders of this folder, and any files in any of these folders.
  - **This Folder And Subfolders** The permissions apply to this folder and any subfolders of this folder. They do not apply to any files in any of these folders.
  - **This Folder And Files** The permissions apply to this folder and any files in this folder. They do not apply to any subfolders of this folder.

- **Subfolders And Files Only** The permissions apply to any subfolders of this folder and any files in any of these folders. They do not apply to this folder itself.
  - **Subfolders Only** The permissions apply to any subfolders of this folder but not to the folder itself or any files in any of these folders.
  - **Files Only** The permissions apply to any files in this folder and any files in subfolders of this folder. They do not apply to this folder itself or to subfolders.
11. When you have finished configuring permissions, click OK.

Because shared folders also have NTFS permissions, you might want to set special NTFS permissions by using Server Manager. To do this, follow these steps:

1. In Server Manager, select File And Storage Services, select the server you want to work with, and then select Shares. Next, right-click the folder with which you want to work, and then click Properties to display a Properties dialog box.
2. When you click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
3. Click Customize Permissions to open the Advanced Security Settings dialog box with the Permissions tab selected.
4. Users or groups that already have access to the file or folder are listed under Permission Entries. Use the options provided to view, edit, add, or remove permissions for users and groups. When you are editing or adding permissions in the Permission Entry dialog box, follow steps 8-11 of the previous procedure to display and work with special permissions.

## Setting Claims-Based Permissions

Claims-based access controls use compound identities that incorporate not only the groups of which a user and the user's computer is a member, but also claim types, which are assertions about objects based on Active Directory attributes, and resource properties, which classify objects and describe their attributes. When resources are remotely accessed, claims-based access controls and central access policies rely on Kerberos with Armoring for authentication of computer device claims. Kerberos with Armoring improves domain security by allowing domain-joined clients and domain controllers to communicate over secure, encrypted channels.

You use claims-based permissions to fine-tune access by defining conditions that limit access as part of a resource's advanced security permissions. Typically, these conditions add device claims or user claims to the access controls. User claims identify users; device claims identify devices. For example, you could define claim types based on business category and country code. The Active Directory attributes

are businessCategory and countryCode, respectively. By using these claim types, you could then fine-tune access to ensure that only users, devices, or both that belong to specific business categories and have certain country codes are granted access to a resource. You could also define a resource property called Project to help fine-tune access even more.

**MORE INFO** With central access policies, you define central access rules in Active Directory and those rules are applied dynamically throughout the enterprise. Central access rules use conditional expressions that require you to determine the resource properties, claim types, and/or security groups required for the policy, in addition to the servers to which the policy should be applied.

Before you can define and apply claim conditions to a computer's files and folders, a claims-based policy must be enabled. For computers that are not joined to the domain, you can do this by enabling and configuring the KDC Support For Claims, Compound Authentication And Kerberos Armoring policy in the Administrative Templates policies for Computer Configuration under System\KDC. The policy must be configured to use one of the following modes:

- **Supported** Domain controllers support claims, compound identities, and Kerberos armoring. Client computers that don't support Kerberos with Armoring can be authenticated.
- **Always Provide Claims** This mode is the same as the Supported mode, but domain controllers always return claims for accounts.
- **Fail Unarmored Authentication Requests** Kerberos with Armoring is mandatory.  
Client computers that don't support Kerberos with Armoring cannot be authenticated.

The Kerberos Client Support For Claims, Compound Authentication And Kerberos Armoring policy controls whether the Kerberos client running on Windows 10 and Windows Server 2016 requests claims and compound authentication. The policy must be enabled for compatible Kerberos clients to request claims and compound authentication for Dynamic Access Control and Kerberos armoring. You'll find this policy in the Administrative Templates policies for Computer Configuration under System\Kerberos.

For application throughout a domain, a claims-based policy should be enabled for all domain controllers in a domain to ensure consistent application. Because of this, you typically enable and configure this policy through the Default Domain Controllers

Group Policy Object (GPO), or the highest GPO linked to the domain controller's organizational unit (OU).

After you've enabled and configured the claims-based policy, you can define claim conditions by completing these steps:

1. In File Explorer, right-click the file or folder with which you want to work, and then click Properties. In the Properties dialog box, select the Security tab, and then click Advanced to display the Advanced Security Settings dialog box.
2. If the user or group already has permissions set for the file or folder, you can edit their existing permissions. Here, click the user with which you want to work, click Edit, and then skip steps 3–6.
3. Click Add to display the Permission Entry dialog box. Click Select A Principal to display the Select User, Computer, Service Account, Or Group dialog box.
4. Enter the name of a user or a group account. Be sure to reference the user account name rather than the user's full name. Only one name can be entered at a time.
5. Click Check Names. If a single match is found for each entry, the dialog box is automatically updated and the entry is underlined. Otherwise, you'll get an additional dialog box. If no matches are found, you either entered the name incorrectly or you're working with an incorrect location. Modify the name in the Name Not Found dialog box and try again, or click Locations to select a new location. When multiple matches are found, in the Multiple Names Found dialog box, select the name you want to use and then click OK.
6. Click OK. The user and group are added as the Principal. Click Add A Condition.
7. Use the options provided to define the condition or conditions that must be met to grant access. With users and groups, set basic claims based on group membership, previously defined claim types, or both. With resource properties, define conditions for property values.
8. When you have finished configuring conditions, click OK.

Because shared folders also have NTFS permissions, you might want to set claimsbased permissions by using Server Manager. To do this, follow these steps:

9. In Server Manager, select File and Storage Services, select the server you want to work with, and then Select Shares.
10. Right-click the folder with which you want to work, and then click Properties to display a Properties dialog box.



11. When you click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
12. Click Customize Permissions to open the Advanced Security Settings dialog box with the Permissions tab selected.
13. Users or groups that already have access to the file or folder are listed under Permission Entries. Use the options provided to view, edit, add, or remove permissions for users and groups. When you are editing or adding permissions in the Permission Entry dialog box, you can add conditions just as I discussed in steps 6 to 8 of the previous procedure.





# Auditing System Resources

Auditing is the best way to track what's happening on your Windows Server 2016 systems. You can use auditing to collect information related to resource usage such as file access, system logons, and system configuration changes. Any time an action occurs that you've configured for auditing, the action is written to the system's security log, where it's stored for your review. The security log is accessible from Event Viewer.

**NOTE** For most auditing changes, you need to be logged on using an account that's a member of the Administrators group or you need to be granted the Manage Auditing And Security Log right in Group Policy.






# Setting Auditing Policies

Auditing policies are essential to help ensure the security and integrity of your systems. Just about every computer system on the network should be configured with some type of security logging. You configure auditing policies for individual computers with local Group Policy and for all computers in domains with Active Directory–based Group Policy. Through Group Policy, you can set auditing policies for an entire site, a domain, or an organizational unit. You can also set policies for an individual workstation or server.

After you access the GPO with which you want to work, you can set auditing policies by following these steps:

1. In the Group Policy Management Editor, access the Audit Policy node by working your way down the console tree. Expand Computer Configuration, Policies, Windows Settings, Security Settings, and Local Policies, and then select Audit Policy.



2. The auditing options are as follows:

- **Audit Account Logon Events** Tracks events related to user logon and logoff.
- **Audit Account Management** Tracks account management by means of Active Directory Users And Computers. Events are generated any time user, computer, or group accounts are created, modified, or deleted.
- **Audit Directory Service Access** Tracks access to Active Directory. Events are generated any time users or computers access the directory.
- **Audit Logon Events** Tracks events related to user logon, logoff, and remote connections to network systems.
- **Audit Object Access** Tracks system resource usage for files, directories, shares, printers, and Active Directory objects.





- **Audit Policy Change** Tracks changes to user rights, auditing, and trust relationships.
- **Audit Privilege Use** Tracks the use of user rights and privileges, such as the right to back up files and directories.

NOTE The Audit Privilege Use policy doesn't track system access-related events, such as the use of the right to log on interactively or the right to access the computer from the network. You track these events with logon and logoff auditing.

- **Audit Process Tracking** Tracks system processes and the resources they use.
- **Audit System Events** Tracks system startup, shutdown, and restart, in addition to actions that affect system security or the security log.

3. To configure an auditing policy, double-click its entry, or right-click the entry, and then click Properties.
4. In the dialog box that is displayed, select the Define These Policy Settings check box, and then select either the Success check box, the Failure check box, or both. Success logs successful events, such as successful logon attempts. Failure logs failed events, such as failed logon attempts.
5. Click OK.

When auditing is enabled, the security event log will reflect the following:

- Event IDs of 560 and 562 detailing user audits
- Event IDs of 592 and 593 detailing process audits


## Auditing Files and Folders

If you configure a GPO to enable the Audit Object Access option, you can set the level of auditing for individual folders and files. This enables you to control precisely how folder and file usage is tracked. Auditing of this type is available only on NTFS volumes.

You can configure file and folder auditing by following these steps:

1. In File Explorer, right-click the file or folder to be audited, and then click Properties.
2. Click the Security tab, and then click Advanced to display the Advanced Security Settings dialog box.
3. On the Auditing tab, click Continue. You can now view and manage auditing settings.





4. The Auditing Entries list shows the users, groups, or computers whose actions you want to audit. To remove an account, select the account in the Auditing Entries list, and then click Remove.
5. To configure auditing for additional users, computers, or groups, click Add. This displays the Select Users, Computers, Service Accounts, Or Groups dialog box.
6. Enter the name of a user, computer, or group in the current domain, and then click Check Names. If a single match is found, the dialog box is automatically updated and the entry is underlined; otherwise, you'll get an additional dialog box. If no matches are found, you either entered the name incorrectly or you're working with an incorrect location. Modify the name in the Name Not Found dialog box and try again, or click Locations to select a new location. When multiple matches are found, in the Multiple Names Found dialog box, select the name you want to use, and then click OK.
7. Click OK. The user and group are added, and the Principal and the Auditing Entry dialog box are updated to show this. Only basic permissions are listed by default. If you want to work with advanced permissions, click Show Advanced Permissions to display the special permissions.
8. As necessary, use the Applies To list to specify where objects are audited. If you are working with a folder and want to replace the auditing entries on all child objects of this folder (and not on the folder itself), select Only Apply These Settings To Objects And/Or Containers Within This Container.



Keep in mind that the Applies To list lets you specify the locations where you want the auditing settings to apply. The Only Apply These Settings To Objects And/Or Containers Within This Container check box controls how auditing settings are applied. When this check box is selected, auditing settings on the parent object replace settings on child objects. When this check box is cleared, auditing settings on the parent are merged with existing settings on child objects.

9. Use the Type list to specify whether you are configuring auditing for success, failure, or both, and then specify which actions should be audited. Success logs successful events, such as successful file reads. Failure logs failed events, such as failed file deletions. The events you can audit are the same as the special permissions, except that you can't audit the synchronizing of offline files and folders. For essential files and folders, you'll typically want to track the following:

- Write Attributes—Successful
- Write Extended Attributes—Successful
- Delete Subfolders And Files—Successful
- Delete—Successful
- Change Permissions—Successful

**TIP** If you want to audit actions for all users, use the special group Everyone; otherwise, select the specific user groups, users, or both that you want to audit.

10. If you're using claims-based policies and want to limit the scope of the auditing entry, you can add claims-based conditions to the auditing entry. For example, if all corporate computers are members of the Domain Computers group, you might want to closely audit access by devices that aren't members of this group.
11. When you have finished configuring auditing, click OK. Repeat this process to audit other users, groups, or computers.

## Auditing the Registry

If you configure a GPO to enable the Audit Object Access option, you can set the level of auditing for keys within the registry. This enables you to track when key values are set, when subkeys are created, and when keys are deleted.

You can configure registry auditing by following these steps:

1. Open the Registry Editor. At a command prompt, enter **regedit** ; or enter **regedit.exe** in the Search box, and then press Enter.
2. Browse to a key you want to audit. On the Edit menu, select Permissions.

- 
3. In the Permissions dialog box, click Advanced. In the Advanced Security Settings dialog box, click the Auditing tab.
  4. Click Add to display the Auditing Entry dialog box. Click Select A Principal to display the Select User, Computer, Service Account, Or Group dialog box.
  5. In the Select User, Computer, Service Account, Or Group dialog box, enter **Everyone**, click Check Names, and then click OK.
  6. In the Auditing Entry dialog box, only basic permissions are listed by default. Click Show Advanced Permissions to display the special permissions.
  7. Use the Applies To list to specify how the auditing entry is to be applied.
  8. Use the Type list to specify whether you are configuring auditing for success, failure, or both, and then specify which actions should be audited. Typically, you'll want to track the following advanced permissions:
    - Set Value—Successful and Failed
    - Create Subkey—Successful and Failed
    - Delete—Successful and Failed
  9. Click OK three times to close all open dialog boxes and apply the auditing settings.

## Auditing Active Directory Objects

If you configure a GPO to enable the Audit Directory Service Access option, you can set the level of auditing for Active Directory objects so that you can control precisely how object usage is tracked.

To configure object auditing, follow these steps:

1. In Active Directory Users And Computers, ensure that Advanced Features is selected on the View menu, and then access the container for the object.
  2. Double-click the object to be audited to open the related Properties dialog box.
  3. Click the Security tab, and then click Advanced.
  4. In the Advanced Settings dialog box, click the Auditing tab. The Auditing Entries list shows the users, groups, or computers whose actions you are auditing currently (if any). To remove an account, select the account in the Auditing Entries list, and then click Remove.
  5. To add specific accounts, click Add to display the Auditing Entry dialog box. Click Select A Principal to display the Select User, Computer, Service Account, Or Group dialog box.
- 



6. Enter the name of a user, computer, or group in the current domain, and then click Check Names. If a single match is found, the dialog box is automatically updated and the entry is underlined; otherwise, you'll get an additional dialog box. If no matches are found, you either entered the name incorrectly or you're working with an incorrect location. Modify the name in the Name Not Found dialog box and try again, or click Locations to select a new location. When multiple matches are found, in the Multiple Names Found dialog box, select the name you want to use, and then click OK.
7. Click OK to return to the Auditing Entry dialog box. Use the Applies To list to specify how the auditing entry is to be applied.
8. Use the Type list to specify whether you are configuring auditing for success, failure, or both, and then specify which actions should be audited. Success logs successful events, such as a successful attempt to modify an object's permissions. Failed logs failed events, such as a failed attempt to modify an object's owner.
9. Click OK. Repeat this process to audit other users, groups, or computers.





## Chapter 28. Configuring Disk Quotas

Windows Server 2016 supports two mutually exclusive types of disk quotas:

- **NTFS disk quotas** NTFS disk quotas are supported on all editions of Windows Server 2016 and enable you to manage disk space usage by users. You configure quotas on a per-volume basis. Although users who exceed limits get warnings, administrators are notified primarily through the event logs.
- **Resource Manager disk quotas** Resource Manager disk quotas are supported on all editions of Windows Server 2016, allowing you to manage disk space usage by folder, by file type, and by volume. Users who are approaching or have exceeded a limit can be automatically notified by email. The notification system also allows for notifying administrators by email, triggering incident reporting, running commands, and logging related events.

Regardless of the quota system being used, you can configure quotas only for NTFS volumes. You can't create quotas for FAT, FAT32, or ReFS volumes.





# Using NTFS Disk Quotas

Administrators use NTFS disk quotas to manage disk space usage for critical volumes, such as those that provide corporate data shares or user data shares. When you enable NTFS disk quotas, you can configure two values:

- **Disk quota limit** Sets the upper boundary for space usage, which you can use to prevent users from writing additional information to a volume, to log events regarding the user exceeding the limit, or both.
- **Disk quota warning** Warns users and logs warning events when users are getting close to their disk quota limit.

**TIP** You can set disk quotas but not enforce them, and you might be wondering why you'd want to do this. Sometimes you want to track disk space usage on a peruser basis and know when users have exceeded some predefined limit, but instead of denying them additional disk space, you log an event in the application log to track the overage. You can then send out warning messages or figure out other ways to reduce the space usage.

NTFS disk quotas apply only to end users and not to administrators. Administrators can't be denied disk space even if they exceed enforced disk quota limits.



# Understanding NTFS Disk Quotas

In a typical environment, you restrict disk space usage in megabytes (MB) or gigabytes (GB). For example, on a corporate data share used by multiple users in a department, you might want to limit disk space usage from 20 to 100 GB. For a user data share, you might want to set the level much lower, such as from 5 to 20 GB, which restricts the user from creating large amounts of personal data. Often you'll set the disk quota warning as a percentage of the disk quota limit. For example, you might set the warning from 90 to 95 percent of the disk quota limit.

Because NTFS disk quotas are tracked on a per-volume, per-user basis, disk space used by one user doesn't affect the disk quotas for other users. Thus, if one user exceeds his limit, any restrictions applied to this user don't apply to other users. For example, if a user exceeds a 5-GB disk quota limit and the volume is configured to prevent writing over the limit, the user can no longer write data to the volume. Users can, however, remove files and folders from the volume to free up disk space. They can also move files and folders to a compressed area on the volume, which might free up space, or they can elect to compress the files themselves. Moving files to a different location on the volume doesn't affect the quota restriction. The amount of file space is the same unless the user moves uncompressed files and folders to a folder with compression. In any case, the restriction on a single user doesn't affect other users' ability to write to the volume (as long as there's free space on the volume).

You can enable NTFS disk quotas on the following:

- **Local volumes** To manage disk quotas on local volumes, you work with the local disk itself. When you enable disk quotas on a local volume, the Windows system files are included in the volume usage for the user who installed those files. Sometimes this might cause the user to go over the disk quota limit so to prevent this, you might want to set a higher limit on a local workstation volume.
- **Remote volumes** To manage disk quotas on remote volumes, you must share the root directory for the volume, and then set the disk quota on the volume. Remember, you set quotas on a per-volume basis, so if a remote file server has separate volumes for different types of data—that is, a corporate data volume and a user data volume—these volumes have different quotas.

Only members of the Domain Admins group or the local system Administrators group can configure disk quotas. The first step in using quotas is to enable quotas in Group Policy, which you can do at two levels:



- **Local** Through local Group Policy, you can enable disk quotas for an individual computer.
- **Enterprise** Through Group Policy that applies to a site, a domain, or an organizational unit, you can enable disk quotas for groups of users and computers.

Having to keep track of disk quotas does cause some overhead on computers. This overhead is a function of the number of disk quotas being enforced, the total size of the volumes and their data, and the number of users to which the disk quotas apply.

Although on the surface disk quotas are tracked per user, behind the scenes Windows Server 2016 manages disk quotas according to security identifiers (SIDs). Because SIDs are tracked by disk quotas, you can safely modify user names without affecting the disk quota configuration. Tracking by SIDs does cause some additional overhead when viewing disk quota statistics for users because Windows Server 2016 must correlate

SIDs to user account names so that the account names can be displayed in dialog boxes. This means contacting the local user manager and the Active Directory domain controller as necessary.

After Windows Server 2016 looks up names, it caches them to a local file so that they can be available immediately the next time they're needed. The query cache is infrequently updated—if you notice a discrepancy between what's displayed and what's configured, you need to refresh the information. Usually, this means choosing Refresh from the View menu or pressing F5 in the current window.

When you apply disk quotas, you need to be particularly careful in the way you enforce quotas, especially with respect to system accounts, service accounts, or other special purpose accounts. Improper application of disk quotas to these types of accounts can cause serious problems that are difficult to diagnose and resolve. Enforcing quotas on the System, NetworkService, and LocalService accounts could prevent the computer from completing important operating system tasks. As an example, if these accounts reach their enforced quota limit, you would not be able to apply changes to Group Policy because the Group Policy client runs within a LocalSystem context by default and would not be able to write to the system disk. If the service can't write to the system disk, Group Policy changes cannot be made, and being unable to change Group Policy could have all sorts of unexpected consequences because you would be stuck with the previously configured settings. For example, you would be unable to disable or modify the quota settings through Group Policy.

In this scenario, where service contexts have reached an enforced quota limit, any other configuration settings that use these service contexts and require making



changes to files on disk would likely also fail. For example, you would be unable to complete the installation or removal of roles, role services, and features. This would leave the server in a state in which Server Manager always includes a warning that you need to restart the computer to complete configuration tasks, but restarting the computer would not resolve these issues.

To address this problem, you need to edit the disk quota entries for the system disk, raise the enforced limits on the service accounts, and then restart the computer. Restarting the computer triggers the finalization tasks and enables the computer to complete any configuration tasks stuck in a pending status. Because the Group Policy client service could process changes and write them to the system disk, changes to Group Policy would then be applied as well.




# Enabling Quotas Through Group Policy

The best way to configure NTFS disk quotas is through Group Policy. When you configure disk quotas through local policy or through unit, domain, and site policy, you define general policies that are set automatically when you enable quota management on individual volumes. Thus, rather than having to configure each volume separately, you can use the same set of rules and apply them in turn to each volume you want to manage.

Policies that control NTFS disk quotas are applied at the system level. You access these policies through the Administrative Templates for Computer Configuration under System\Disk Quotas. The available policies are:

- **Apply Policy To Removable Media** Determines whether quota policies apply to NTFS volumes on removable media. If you don't enable this policy, quota limits apply only to fixed media drives.
- **Enable Disk Quotas** Turns disk quotas on or off for all NTFS volumes of the computer, and prevents users from changing the setting.
- **Enforce Disk Quota Limit** Specifies whether quota limits are enforced. If quotas are enforced, users will be denied disk space if they exceed the quota. This overrides settings on the Quota tab on the NTFS volume.
- **Log Event When Quota Limit Exceeded** Determines whether an event is logged when users reach their limit, and prevents users from changing their logging options.
- **Log Event When Quota Warning Level Exceeded** Determines whether an event is logged when users reach the warning level.
- **Specify Default Quota Limit And Warning Level** Sets a default quota limit and warning level for all users. This setting overrides other settings and affects only new users.





Whenever you work with quota limits, you should use a standard set of policies on all systems; however, you typically, won't want to enable all the policies. Instead, you'll selectively enable policies and then use the standard NTFS features to control quotas on various volumes. If you want to enable quota limits, follow these steps:

1. In the Group Policy Management Console, open the appropriate Group Policy Object for editing or access local policy on the server with which you want to work.
2. In the policy editor, access the Disk Quotas node by expanding System, and then selecting Disk Quotas in the Administrative Templates for Computer Configuration.
3. Double-click Enable Disk Quotas. Select Enabled, and then click OK.
4. Double-click Enforce Disk Quota Limit. If you want to enforce disk quotas on all NTFS volumes residing on this computer, click Enabled. Otherwise, click Disabled, and then set specific limits on a per-volume basis. Click OK.
5. Double-click Specify Default Quota Limit And Warning Level. In the dialog box shown, select Enabled.
6. Under Default Quota Limit, set a default limit that's applied to users when they first write to the quota-enabled volume. The limit doesn't apply to current users or affect current limits in place. On a corporate share, such as a share used by members of a project team, a good limit is between 5 and 10 GB. Of course, this depends on the size of the data files that the users routinely work with, the number of users, and the size of the disk volume. Graphic designers and data engineers might need much more disk space.
7. To set a warning limit, scroll down in the Options window. A good warning limit is about 90 percent of the default quota limit, which means that if you set the default quota limit to 10 GB, you should set the warning limit to 9 GB. Click OK.

Specify default quota limit and warning level

Previous Setting Next Setting

☐ Not Configured
 ☒ Enabled
 ☐ Disabled

Comment:

Supported on: At least Windows 2000

Options:

Help:

Default quota limit:

Value: 1

Units: GB

Default warning level:

Value: 900

Units: MB

This policy setting specifies the default disk quota limit and warning level for new users of the volume.

This policy setting determines how much disk space can be used by each user on each of the NTFS file system volumes on a computer. It also specifies the warning level, the point at which the user's status in the Quota Entries window changes to indicate that the user is approaching the disk quota limit.

This setting overrides new users' settings for the disk quota limit and warning level on their volumes, and it disables the corresponding options in the "Select the default quota limit for new users of this volume" section on the Quota tab.

This policy setting applies to all new users as soon as they write to the volume. It does not affect disk quota limits for current users, or affect customized limits and warning levels set for particular users (on the Quota tab in Volume Properties).

If you disable or do not configure this policy setting, the disk space available to users is not limited. The disk quota

8. Double-click Log Event When Quota Limit Exceeded. Select Enabled so that limit events are recorded in the application log, and then click OK.
9. Double-click Log Event When Quota Warning Level Exceeded. Select Enabled so that warning events are recorded in the application log, and then click OK.
10. Double-click Apply Policy To Removable Media. Select Disabled so that the quota limits apply only to fixed media volumes on the computer, and then click OK.

**TIP** To ensure that the policies are enforced immediately, access the System\Group Policy node in the Administrative Templates for Computer Configuration, and then double-click Configure Disk Quota Policy Processing. Select Enabled, and then select the Process Even If The Group Policy Objects Have Not Changed check box. Click OK.



## Enabling Quotas Through Disk Properties

You can set NTFS disk quotas on a per-volume basis. Only NTFS volumes can have disk quotas. After you configure the appropriate group policies, you can use Computer Management to set disk quotas for local and remote volumes.

**NOTE** If you use the Enforce Disk Quota Limit policy setting to enforce quotas, users are denied disk space if they exceed the quota. This overrides settings on the Quota tab on the NTFS volume.

If you aren't using quotas enforced by policy and want to enable NTFS disk quotas on an NTFS volume, follow these steps:

1. Using Volume List view or Graphical View in Disk Management, right-click the volume with which you want to work, and then click Properties.
2. On the Quota tab, select Show Quota Settings.





3. Select the Enable Quota Management check box. If you already set quota management values through Group Policy, the options are unavailable and you can't change them. You must modify options through Group Policy instead.

**BEST PRACTICES** Whenever you work with the Quota tab, pay particular attention to the Status text and the associated traffic light icon. Both change based on the state of quota management. If quotas aren't configured, the traffic light icon shows a red light and the status is inactive or not configured. If the operating system is working or updating the quotas, the traffic light icon shows a yellow light and the status shows the activity being performed. If quotas are configured, the traffic light icon shows a green light and the status text states that the quota system is active.

4. To set a default disk quota limit for all users, select Limit Disk Space To. In the text boxes provided, set a limit in kilobytes, megabytes, gigabytes, terabytes, petabytes, or exabytes. Then set the default warning limit in the Set Warning Level To text boxes. Again, you'll usually want the disk quota warning limit to be 90–95 percent of the disk quota limit.

**TIP** Although the default quota limit and warning apply to all users, you can configure different levels for individual users. You do this in the Quota Entries dialog box. If you create many unique quota entries and don't want to re-create them on a volume with similar characteristics and usage, you can export the quota entries and import them into a different volume.

5. To enforce the disk quota limit and prevent users from going over the limit, select the Deny Disk Space To Users Exceeding Quota Limit check box. Keep in mind that this creates an actual physical limitation for users (but not for administrators).
6. To configure logging when users exceed a warning limit or the quota limit, select the Log Event check boxes. Click OK to save your changes.
7. If the quota system isn't currently enabled, you'll get a prompt asking you to enable the quota system. Click OK so that Windows Server 2016 can rescan the volume and update disk usage statistics. Actions might be taken against users who exceed the current limit or warning levels. These actions can include preventing additional writing to the volume, notifying them the next time they access the volume, and logging applicable events to the application log.



# Viewing Disk Quota Entries

Disk space usage is tracked on a per-user basis. When disk quotas are enabled, each user storing data on a volume has an entry in the disk quota file. This entry is updated periodically to show the current disk space used, the applicable quota limit, the applicable warning level, and the percentage of allowable space being used. As an administrator, you can modify disk quota entries to set different limits and warning levels for particular users. You can also create disk quota entries for users who haven't yet saved data on a volume. The key reason for creating entries is to ensure that when a user does make use of a volume, the user has an appropriate limit and warning level.

To view the current disk quota entries for a volume, follow these steps:

- 1. Using Volume List view or Graphical View in Disk Management, right-click the volume with which you want to work, and then click Properties.
- 2. On the Quota tab, select Show Quota Settings. Next, click Quota Entries in the Quota Settings dialog box. This displays the Quota Entries dialog box.

| Status      | Name                              | Logon Name                       | Amount Used | Quota Limit | Warning Level | Percent Used |
|-------------|-----------------------------------|----------------------------------|-------------|-------------|---------------|--------------|
| Above Limit |                                   | NT SERVICE\TrustedInstaller      | 6.84 GB     | 1 GB        | 900 MB        | 684          |
| Above Limit |                                   | NT AUTHORITY\SYSTEM              | 2.18 GB     | 1 GB        | 900 MB        | 218          |
| OK          |                                   | BUILTIN\Administrators           | 6.73 GB     | No Limit    | No Limit      | N/A          |
| OK          | williams                          | IMAGINEDLANDS\williams           | 106.77 MB   | No Limit    | No Limit      | N/A          |
| OK          |                                   | NT AUTHORITY\LOCAL SERVICE       | 84.39 MB    | 1 GB        | 900 MB        | 8            |
| OK          |                                   | NT AUTHORITY\NETWORK SER...      | 42.61 MB    | 1 GB        | 900 MB        | 4            |
| OK          |                                   | NT SERVICE\MapsBroker            | 1 KB        | 1 GB        | 900 MB        | 0            |
| OK          | [Account Information Unavailable] | S-1-5-21-397955417-626881126-... | 1 MB        | 1 GB        | 900 MB        | 0            |

- 3. Each quota entry is listed according to a status, which is meant to quickly depict whether a user has gone over a limit. A status of OK means the user is working within the quota boundaries. Any other status usually means the user has reached the warning level or the quota limit.

# Creating Disk Quota Entries

You can create disk quota entries for users who haven't yet saved data on a volume. This enables you to set custom limits and warning levels for a particular user. You usually use this feature when a user frequently stores more information than other users and you want to allow the user to go over the normal limit or when you want to set a specific limit for administrators. As you might recall, administrators aren't subject to disk quota limits, so if you want to enforce limits for individual administrators, you must create disk quota entries for each administrator you want to limit.

REAL WORLD You shouldn't create individual disk quota entries haphazardly. You need to track individual entries carefully. Ideally, you should keep a log that






details any individual entries so that other administrators understand the policies in place and how those policies are applied. When you modify the base rules for quotas on a volume, you should reexamine individual entries to see whether they're still applicable or need to be updated as well. I've found that certain types of users are exceptions more often than not, and that it's sometimes better to put different classes of users on different volumes and then apply disk quotas to each volume. In this way, each class or category of user has a quota limit that's appropriate for its members' typical usage, and you have fewer (perhaps no) exceptions. For example, you might use separate volumes for executives, managers, and standard users, or you might have separate volumes for management, graphic designers, engineers, and all other users.

To create a quota entry on a volume, follow these steps:


1. Using Volume List view or Graphical View in Disk Management, right-click the volume with which you want to work, and then click Properties.
2. On the Quota tab, select Show Quota Settings. Next, click Quota Entries in the Quota Settings dialog box. This displays the Quota Entries dialog box. Current quota entries for all users are listed. To refresh the listing, press F5 or choose Refresh from the View menu.
3. If the user doesn't have an existing entry on the volume, you can create it by choosing New Quota Entry from the Quota menu. This opens the Select Users dialog box.



4. In the Select Users dialog box, enter the name of a user you want to use in the Enter The Object Names To Select text box, and then click Check Names. If a match is found, select the account you want to use, and then click OK. If no matches are found, update the name you entered and try searching again. Repeat

this step as necessary, and then click OK. You also can enter user names separated by semicolons to apply the same quota to multiple users at the same time.

5. After you select a user, the Add New Quota Entry dialog box is displayed. You have two options. You can remove all quota restrictions for this user by selecting Do Not Limit Disk Usage, or you can set a specific limit and warning level by selecting Limit Disk Space To and then entering the appropriate values. Click OK.



## Deleting Disk Quota Entries

When you've created disk quota entries on a volume and a user no longer needs to use the volume, you can delete the associated disk quota entry. When you delete a disk quota entry, all files owned by the user are collected and displayed in a dialog box so that you can permanently delete the files, take ownership of the files, or move the files to a folder on a different volume.

To delete a disk quota entry for a user and manage the user's remaining files on the volume, follow these steps:

1. Using Volume List view or Graphical View in Disk Management, right-click the volume with which you want to work, and then click Properties.
2. On the Quota tab, select Show Quota Settings. Next, click Quota Entries in the Quota Settings dialog box. This displays the Quota Entries dialog box. Current quota entries for all users are listed. To refresh the listing, press F5 or choose Refresh from the View menu.
3. Select the disk quota entry you want to delete, and then press Delete, or choose Delete Quota Entry from the Quota menu. You can select multiple entries by using the Shift and Ctrl keys.



4. When prompted to confirm the action, click Yes to display the Disk Quota dialog box with a list of current files owned by the selected user or users.
5. In the List Files Owned By list, display files for a user whose quota entry you're deleting. You must now specify how the files for the user are to be handled. You can handle each file separately by selecting individual files and then choosing an appropriate option, or you can select multiple files by using the Shift and Ctrl keys. The following options are available:
  - **Permanently Delete Files** Select the files to delete, and then press Delete. When prompted to confirm the action, click Yes.
  - **Take Ownership Of Files** Select the files you want to take ownership of, and then click Take Ownership Of Files.
  - **Move Files To** Select the files you want to move, and then enter the path to a folder on a different volume. If you don't know the path you want to use, click Browse to display the Browse For Folder dialog box. When you find the folder, click Move.
6. Click Close when you have finished managing the files. If you've appropriately handled all user files, the disk quota entries will be deleted.

## Exporting and Importing NTFS Disk Quota Settings

Rather than re-creating custom disk quota entries on individual volumes, you can export the settings from a source volume, and then import the settings to another volume. You must format both volumes by using NTFS. To export and then import disk quota entries, follow these steps:

1. Using Volume List view or Graphical View in Disk Management, right-click the volume with which you want to work, and then click Properties.
2. On the Quota tab, select Show Quota Settings. Next, click Quota Entries in the Quota Settings dialog box. This displays the Quota Entries dialog box.
3. Current quota entries for all users are listed. To refresh the listing, press F5 or choose Refresh from the View menu.
4. Select Export from the Quota menu to display the Export Quota Settings dialog box. Choose the save location for the file containing the quota settings, and then enter a name for the file in the File Name text box. Click Save.

**NOTE** If you save the settings file to a mapped drive on the target volume, you'll have an easier time importing the settings. Quota files are usually small, so you don't need to worry about disk space usage.

5. On the Quota menu, click Close to exit the Quota Entries dialog box.



- 
6. Right-click Computer Management in the console tree, and then click Connect To Another Computer. In the Select Computer dialog box, choose the computer containing the target volume. The target volume is the one on which you want to use the exported settings.
  7. As explained previously, open the Properties dialog box for the target volume. Then click Quota Entries on the Quota tab to display the Quota Entries dialog box for the target volume.
  8. Click Import on the Quota menu. In the Import Quota Settings dialog box, select the quota settings file you saved previously. Click Open.
  9. If the volume had previous quota entries, you are given the choice to replace existing entries or keep existing entries. When prompted about a conflict, click Yes to replace an existing entry, or click No to keep the existing entry. To apply the option to replace or keep existing entries to all entries on the volume, select the Do This For All Quota Entries check box prior to clicking Yes or No.

## Disabling NTFS Disk Quotas

When disk quotas are managed in Group Policy, you'll need to edit the appropriate policy object or objects to disable disk quotas:

1. In the Group Policy Management Console, open the appropriate Group Policy Object for editing or access local policy on the server with which you want to work.
2. In the policy editor, access the Disk Quotas node by expanding System, and then selecting Disk Quotas in the Administrative Templates for Computer Configuration.
3. Double-click Enable Disk Quotas. Select Disabled or Not Configured as appropriate, and then click OK.
4. Double-click Enforce Disk Quota Limit. Select Disabled or Not Configured as appropriate, and then click OK.
5. Double-click Specify Default Quota Limit And Warning Level. Select Disabled or Not Configured as appropriate, and then click OK.
6. Double-click Log Event When Quota Limit Exceeded. Select Disabled or Not Configured as appropriate, and then click OK.
7. Double-click Log Event When Quota Warning Level Exceeded. Select Disabled or Not Configured as appropriate, and then click OK.

Otherwise, you can disable quotas for individual users or all users on per-volume basis. When you disable quotas for a particular user, the user is no longer subject to





the quota restrictions but disk quotas are still tracked for other users. When you disable quotas on a volume, quota tracking and management are completely removed. To disable quotas for a particular user, follow the technique outlined earlier in the chapter in “Viewing Disk Quota Entries.” To disable quota tracking and management on a particular volume, follow these steps:

8. Using Volume List view or Graphical View in Disk Management, right-click the volume with which you want to work, and then click Properties.
9. On the Quota tab, select Show Quota Settings. This displays the Quota Settings dialog box.
10. Clear the Enable Quota Management check box. Click OK. When prompted to confirm, click OK.

**IMPORTANT** Disabling quota tracking doesn't delete existing quota entries on a volume. If you later enable quota management, the previously created quota entries will be available and will be enforced.

## Resource Manager Disk Quotas


Windows Server 2016 supports an enhanced quota management system called Resource Manager disk quotas . By using Resource Manager disk quotas, you can manage disk space usage by folder and by volume.

**TIP** Because you manage Resource Manager disk quotas separately from NTFS disk quotas, you can configure a single volume to use both quota systems; however, it's recommended that you use one quota system or the other. Alternatively, if you've already configured NTFS disk quotas, you might want to continue by using NTFS disk quotas on a per-volume basis and supplement this quota management with Resource Manager disk quotas for important folders.



# Understanding Resource Manager Disk Quotas

If you want to use Resource Manager disk quotas on a server, you must add the File Server Resource Manager role service. This role service is part of the File And Storage Services role.




When you're working with Windows Server 2016, Resource Manager disk quotas are another tool you can use to manage disk usage. You can configure Resource Manager disk quotas on a per-volume or per-folder basis. You can set disk quotas with a specific hard limit—meaning a limit can't be exceeded—or a soft limit, meaning a limit can be exceeded.

Generally, you should use hard limits when you want to prevent users from exceeding a specific disk-usage limitation. Use soft limits when you want to monitor usage and simply warn users who exceed or are about to exceed usage guidelines. All quotas have a quota path, which designates the base file path on the volume or folder to which the quota is applied. The quota applies to the designated volume or folder and all subfolders of the designated volume or folder. The particulars of how quotas work and how users are limited or warned are derived from a source template that defines the quota properties.

Windows Server 2016 includes several default templates including:

- **Templates with 100 MB, 2 GB, 5GB and 10 GB limits** Each of which has a designated hard limit and sends warnings to users as the limit is approached and exceeded.
- **Templates that monitor volume usage** Each of which has a soft limit, monitors volume usage, and warns when the limit is approached and exceeded.
- **Templates that monitor share usage** Each of which has a soft limit, monitors share usage, and warns when the limit is approached and exceeded.





By using the File Server Resource Manager, you can easily define additional templates that would then be available whenever you define quotas, or you can set single-use custom quota properties when defining a quota.

Quota templates or custom properties define the following:

- **Limit** The disk space usage limit
- **Quota type** Hard or soft
- **Notification thresholds** The types of notification that occur when usage reaches a specific percentage of the limit

Although each quota has a specific limit and type, you can define multiple notification thresholds as either a warning threshold or a limit threshold. Warning thresholds are considered to be any percentage of the limit that is less than 100 percent. Limit thresholds occur when the limit reached is 100 percent. For example, you could define warning thresholds that are triggered at 85 percent and 95 percent of the limit and a limit threshold that is triggered when 100 percent of the limit is reached.

Users who are approaching or have exceeded a limit can be automatically notified by email. The notification system also allows for notifying administrators by email, triggering incident reporting, running commands, and logging related events.

## Managing Disk Quota Templates

You use disk quota templates to define quota properties, including the limit, quota type, and notification thresholds. In File Server Resource Manager, you can view the currently defined disk quota templates by expanding the Quota Management node, and then selecting Quota Templates. Variables that can be used for automatically generated messages and events include:



- **[Admin Email]** Inserts the email addresses of the administrators defined under the global options
- **[File Screen Path]** Inserts the local file path, such as C:\Data
- **[File Screen Remote Path]** Inserts the remote path, such as \\server\share
- **[File Screen System Path]** Inserts the canonical file path, such as
- **\\?\VolumeGUID [Server Domain]** Inserts the domain of the server on which the notification occurred
- **[Server]** Inserts the server on which the notification occurred
- **[Source File Owner]** Inserts the user name of the owner of the file/folder
- **[Source File Owner Email]** Inserts the email address of the owner of the file/folder
- **[Source File Path]** Inserts the source path of the file/folder

You can modify existing disk quota templates by following these steps:

1. In File Server Resource Manager, expand the Quota Management node, and then select Quota Templates.

Currently defined disk quota templates are listed by name, limit, and quota type.

2. To modify disk quota template properties, double-click the disk quota template name. This displays a related Properties dialog box.





Quota Template Properties for 250 MB Extended Limit

Copy properties from quota template (optional):  
 250 MB Extended Limit Copy

Settings

Template name:  
 250 MB Extended Limit

Description (optional):

Space limit

Limit:  
 250.000 MB

☒ Hard quota: Do not allow users to exceed limit  
☐ Soft quota: Allow users to exceed limit (use for monitoring)

Notification thresholds

| Threshold      | E-mail | Event Log | Command | Report |
|----------------|--------|-----------|---------|--------|
| Warning (85%)  | ✓      |           |         |        |
| Warning (95%)  | ✓      | ✓         |         |        |
| Warning (100%) | ✓      | ✓         |         |        |

Add... Edit... Remove

3. On the Settings tab, you can set the template name, limit, and quota type. Current notification thresholds are listed. To modify an existing threshold, select it and then click Edit. To define a new threshold, click Add.
4. When you have finished modifying the quota template, click OK to save the changes.

You can create a new disk quota template by following these steps:

1. In File Server Resource Manager, expand the Quota Management node, and then select Quota Templates.
2. On the Action menu or in the Actions pane, click Create Quota Template to display the Create Quota Template dialog box.
3. On the Settings tab, set the template name, limit, and quota type. You should create a limit threshold first, and then create additional warning thresholds as necessary. In the Limit list, enter the limit value and specify whether you are setting the limit in kilobytes, megabytes, gigabytes, or terabytes.

4. Click Add to add warning thresholds. In the Add Threshold dialog box, enter a percentage value under Generate Notifications When Usage Reaches (%).

Warning thresholds are considered to be any percentage of the limit that is less than 100 percent. Limit thresholds occur when the limit reached is 100 percent.

5. On the E-Mail Message tab, you can configure notification as follows:
  - To notify an administrator when the disk quota is triggered, select the Send E-Mail To The Following Administrators check box, and then enter the email address or addresses to use. Be sure to separate multiple email addresses with a semicolon. Use the value [Admin Email] to specify the default administrator as configured previously under the global options.
  - To notify users, select the Send E-Mail To The User Who Exceeded The Threshold check box.
  - Specify the contents of the notification message in the Subject and Message Body text boxes.
6. On the Event Log tab, you can configure event logging. Select the Send Warning To Event Log check box to enable logging, and then specify the text of the log entry in the Log Entry text box.
7. On the Command tab, you can optionally specify a command or script to run, arguments to pass in to the command or script, and a working directory. The default the security context for commands is Local Service, which grants standard user access to local resources but denies access to network resources. If the command requires access to both local and network resources, you can run the command as Network Service.
8. On the Report tab, select the Generate Reports check box to enable incident reporting, and then select the types of reports to generate. Incident reports are stored under %SystemDrive%\StorageReports\Incident by default, and they can also be sent to designated administrators. Use the value [Admin Email ] to specify the default administrator as configured previously under the global options.
9. Repeat steps 5–7 to define additional notification thresholds.
10. Click OK when you have finished creating the template.

## Creating Resource Manager Disk Quotas

You use disk quotas to designate file paths that have specific usage limits. In File Server

Resource Manager, you can view current disk quotas by expanding the Quota



Management node, and then selecting Quotas. Before you define disk quotas, you should specify screening file groups and disk quota templates that you will use, as discussed in “Managing Disk Quota Templates” earlier in this chapter.

After you define the necessary file groups and disk quota templates, you can create a disk quota by following these steps:

1. In File Server Resource Manager, expand the Quota Management node, and then select Quotas.
2. Click Create Quota on the Action menu or in the Actions pane.
3. In the Create Quota dialog box, set the local computer path for the quota by clicking Browse and then by using the Browse For Folder dialog box to select the path, such as C:\Data. Click OK.
4. In the Derive Properties From This Quota Template list, choose the disk quota template that defines the quota properties you want to use.
5. Click Create.





# Chapter 29. Using Group Policy for Administration

You can use Group Policy to manage users and computers in many different ways. In this chapter, I'll describe some specific management areas, including the following:

- [Folder redirection](#)
- [Computer and user scripts](#)
- [Software deployment](#)
- [Work Folders options](#)
- [Computer and user certificate enrollment](#)
- [Automatic update settings](#)





# Managing Automatic Updates in Group Policy

Automatic Updates help you keep the operating system up to date. Although you can configure Automatic Updates on a per-computer basis, you'll typically want to configure this feature for all users and computers that process a GPO—this is a much more efficient management technique.

Note that by default, Windows 10 and Windows Server 2016 use Windows Update to download Windows Components in addition to binaries for roles, role services, and features. If the Windows diagnostics framework detects that a Windows component needs to be repaired, Windows uses Windows Update to download the component. If an administrator is trying to install a role, role service, or feature and the payload is missing, Windows uses Windows Update to download the related binaries.





## Configuring Automatic Updates

When you manage Automatic Updates through Group Policy, you can set the update configuration to any of the following options:

- **Auto Download And Schedule The Install** Updates are automatically downloaded and installed according to a schedule you specify. When updates have been downloaded, the operating system notifies the user so that she can review the updates that are scheduled to be installed. The user can install the updates at that time or wait for the scheduled installation time.
- **Auto Download And Notify For Install** The operating system retrieves all updates as they become available, and then prompts the user when they're ready to be installed. The user can then accept or reject the updates. Accepted updates are installed. Rejected updates aren't installed but remain on the system, where they can be installed at a later date.
- **Notify For Download And Notify For Install** The operating system notifies the user before retrieving any updates. If a user elects to download the updates, the user still has the opportunity to accept or reject them. Accepted updates are installed. Rejected updates aren't installed but remain on the system, where they can be installed at a later date.
- **Allow Local Admin To Choose Setting** Allows the local administrator to configure Automatic Updates on a per-computer basis. Note that if you use any other setting, local users and administrators are unable to change settings for Automatic Updates.

You can configure Automatic Updates in Group Policy by following these steps:

1. In the GPMC, right-click the GPO with which you want to work, and then click Edit.
  2. In the policy editor, access Windows Components\Windows Update in the Administrative Templates policies for Computer Configuration.
  3. Double-click Configure Automatic Updates. In the Properties dialog box, you can now enable or disable Group Policy management of Automatic Updates. To enable management of Automatic Updates, select Enabled. To disable management of Automatic Updates, select Disabled, click OK, and then skip the remaining steps.
  4. Choose an update configuration from the options in the Configure Automatic Updating list. On Windows 8 and later as well as Windows Server 2012 and later, updates can be automatically installed during the scheduled maintenance window by selecting the Install During Automatic Maintenance check box.
- 

5. If you select Auto Download And Schedule The Install, you can schedule the installation day and time by using the lists provided. Click OK to save your settings.

Configure Automatic Updates

Previous Setting Next Setting

☐ Not Configured Comment:

☒ Enabled

☐ Disabled

Supported on: Windows XP Professional Service Pack 1 or At least Windows 2000 Service Pack 3

Options:

Configure automatic updating:

4 - Auto download and schedule the install

The following settings are only required and applicable if 4 is selected.

☒ Install during automatic maintenance

Scheduled install day:

0 - Every day

Scheduled install time: 03:00

Help:

Specifies whether this computer will receive security updates and other important downloads through the Windows automatic updating service.

Note: This policy does not apply to Windows RT.

This setting lets you specify whether automatic updates are enabled on this computer. If the service is enabled, you must select one of the four options in the Group Policy Setting:

2 = Notify before downloading and installing any updates.

When Windows finds updates that apply to this computer, users will be notified that updates are ready to be downloaded. After going to Windows Update, users can download and install any available updates.

3 = (Default setting) Download the updates automatically and notify when they are ready to be installed

Windows finds updates that apply to the computer and

By default, Windows Update runs daily at 3:00 a.m. as part of other automatic maintenance and uses the computer's power management features to wake the computer from hibernation or sleep at the scheduled update time and then install updates. Generally, this wake-up-and-install process will occur whether the computer is on battery or AC power.

If a restart is required to finalize updates applied as part of automatic maintenance and there is an active user session, Windows caches the credentials of the user currently logged on to the console and then restarts the computer automatically. After the restart, Windows uses the cached credentials to sign in as this user. Next, Windows restarts applications that were running previously, and then locks the session using the Secure Desktop. If BitLocker is enabled, the entire process is protected by BitLocker encryption as well.

The maintenance process does not need a user to be logged on. The maintenance process runs whether a user is logged on or not. If no user is logged on when scheduled maintenance begins and a restart is required, Windows restarts the computer without



caching credentials or storing information about running applications. When Windows restarts, Windows does not log on as any user.

Because Windows automatically wakes computers to perform automatic maintenance and updates, you'll also want to carefully consider the power options that are applied. Unless a power plan is configured to turn off the display and put the computer to sleep, the computer may remain powered on for many hours after automatic maintenance and updates.






# Optimizing Automatic Updates

By default, current releases of Windows and Windows Server schedule a restart if required to complete updates. This restart occurs at the scheduled install time during the maintenance window. If you want important updates to be installed quicker, follow these steps:

1. In the GPMC, right-click the GPO with which you want to work, and then click Edit.
2. In the policy editor, access Windows Components\Windows Update in the Administrative Templates policies for Computer Configuration.

The screenshot shows the 'Always automatically restart at the scheduled time' policy settings dialog box. At the top, there are 'Previous Setting' and 'Next Setting' buttons. Below them, there are three radio buttons: 'Not Configured', 'Enabled' (which is selected), and 'Disabled'. To the right of these is a 'Comment:' text box. Below the radio buttons is a 'Supported on:' section with a dropdown menu showing 'At least Windows Server 2012, Windows 8 or Windows RT'. At the bottom, there is an 'Options:' section with a dropdown menu set to '60' minutes, and a 'Help:' section with a text box containing instructions: 'If you enable this policy, a restart timer will always begin immediately after Windows Update installs important updates, instead of first notifying users on the login screen for at least two days. The restart timer can be configured to start with any value from 15 to 180 minutes. When the timer runs out, the restart will proceed even if the PC has signed-in users.'

3. Double-click Always Automatically Restart At The Scheduled Time. In the Properties dialog box, select Enabled to allow a restart timer to be used to ensure important updates are installed.
4. Set a restart timer value. The restart timer always begins immediately after Windows Update installs important updates and can be configured to start with any value from 15 to 180 minutes. Then when the timer runs out, the computer is restarted.
5. Click OK.

Updates for Windows 10 and Windows Server 2016 include changes that upgrade the operating system as well as update components and features. On critical systems you might not want these types of updates or any others to be installed immediately. Instead, you may want to delay updates or defer upgrades to ensure they can be tested first by following these steps:

1. In the GPMC, right-click the GPO with which you want to work, and then click Edit.

2. In the policy editor, access Windows Components\Windows Update in the Administrative Templates policies for Computer Configuration.
3. Double-click **Defer Upgrades And Updates**. In the Properties dialog box, select **Enabled**.
4. Specify how long upgrades should be deferred with a value between 1 and 8 months. If you enter a value of 0, upgrades will not be deferred.
5. Specify how long updates should be deferred with a value between 1 and 4 weeks. If you enter a value of 0, updates will not be deferred.
6. Click **OK**.

The screenshot shows the 'Defer Upgrades and Updates' policy settings dialog box. At the top, there are 'Previous Setting' and 'Next Setting' buttons. Below them are three radio buttons: 'Not Configured', 'Enabled' (which is selected), and 'Disabled'. To the right of these is a 'Comment:' text box. Below the radio buttons is a 'Supported on:' dropdown menu showing 'At least Windows 10 Server or Windows 10'. Under the 'Options:' section, there are two spinners: 'Defer upgrades for the following duration (months):' set to 8, and 'Defer updates for the following duration (weeks):' set to 4. There is also a checkbox labeled 'Pause Upgrades and Updates' which is currently unchecked. On the right side, under the 'Help:' section, there is explanatory text: 'Enable this policy in order to defer upgrades for up to eight months. You can also choose to delay updates for up to one month. If you do not delay updates, your PC will remain up to date with security updates as they become available. If an issue arises with an update or upgrade, select "Pause Upgrades and Updates". This will delay updates and upgrades until the next monthly update or upgrade becomes available. Once a new update or upgrade is available, the value will go back to the previously selected option, re-enabling your validation groups.'

**TIP** Defer Upgrades And Updates policy can also be used to pause upgrades and updates until the next monthly update or upgrade becomes available. Use this option if a problem occurs and you need to temporarily delay applying updates or upgrades. The Pause option will clear automatically when a new update or upgrade becomes available. Note also that the settings of this policy are ignored if Specify Intranet Microsoft Update Service Location policy is enabled.

## Using Intranet Update Service Locations

Computers running Windows 10 get updates from other PCs on the local network or the Internet and send parts of previously downloaded updates to other PCs. Computers running other versions of Windows and all versions of Windows Server get their updates through the Automatic Update process. On networks with hundreds or thousands of computers, either approach can use a considerable amount of network bandwidth, and having all the computers check for updates and install them over the Internet doesn't make sense.

Instead, computers running Windows 10 can be configured so that they don't share updates or get updates from PCs on the Internet and like other computers, they also can be configured to get updates from a dedicated update server on your network. The designated update server must run Windows Server Update Services (WSUS), be configured as a web server running IIS, and be able to handle the additional workload, which might be considerable on a large network during peak usage times.

Additionally, the update server must have access to the external network on port 80. The use of a firewall or proxy server on this port shouldn't present any problems.

The update process also tracks configuration information and statistics for each computer. This information is necessary for the update process to work properly, and it can be stored on a separate statistics server (an internal server running IIS) or on the update server itself.

To specify an internal update server, follow these steps:

1. After you install and configure an update server, open the GPO with which you want to work for editing. In the policy editor, access Windows Components\Windows Update in the Administrative Templates policies for Computer Configuration.
2. Double-click Specify Intranet Microsoft Update Service Location. In the Properties dialog box, select Enabled.

Specify intranet Microsoft update service location

Previous Setting Next Setting

☐ Not Configured Comment:

☒ Enabled

☐ Disabled

Supported on: At least Windows XP Professional Service Pack 1 or Windows 2000 Service Pack 3, excluding Windows RT

Options:

Set the intranet update service for detecting updates:  
http://CorpUpdateServer01

Set the intranet statistics server:  
http://CorpUpdateServer01

Help:

Specifies an intranet server to host updates from Microsoft Update. You can then use this update service to automatically update computers on your network.

This setting lets you specify a server on your network to function as an internal update service. The Automatic Updates client will search this service for updates that apply to the computers on your network.

3. In the Set The Intranet Update Service For Detecting Updates text box, enter the URL of the update server. In most cases, this is http://servername, such as http://CorpUpdateServer01.
4. Enter the URL of the statistics server in the Set The Intranet Statistics Server text box. This doesn't have to be a separate server; you can specify the update server in this text box.



**NOTE** If you want a single server to handle both updates and statistics, enter the same URL in both boxes. Otherwise, if you want a different server for updates and statistics, enter the URL for each server in the appropriate box.

5. Click OK. After the applicable GPO is refreshed, systems running appropriate versions of Windows will look to the update server for updates. You'll want to monitor the update and statistics servers closely for several days or weeks to ensure that everything is working properly. Directories and files will be created on the update and statistics servers.

## User and Computer Script Management

You can configure four types of scripts with Windows Server:

- **Computer Startup** Executed during startup
- **Computer Shutdown** Executed prior to shutdown
- **User Logon** Executed when a user logs on
- **User Logoff** Executed when a user logs off

Windows supports scripts written as command-shell batch scripts ending with the .bat or .cmd extension or scripts that use the Windows Script Host (WSH). WSH is a feature of Windows Server that enables you to use scripts written in a scripting language, such as VBScript, without needing to insert the script into a webpage. To provide a multipurpose scripting environment, WSH relies on scripting engines. A scripting engine is the component that defines the core syntax and structure of a particular scripting language. Windows Server ships with scripting engines for VBScript and JScript. Other scripting engines are also available.

Current Windows operating systems also support Windows PowerShell scripts. When Windows PowerShell is installed on computers that process a particular GPO, you can use Windows PowerShell scripts in much the same way as you use other scripts. You have the option of running Windows PowerShell scripts before or after other types of scripts.

## Assigning Computer Startup and Shutdown Scripts


Computer startup and shutdown scripts are assigned as part of a GPO; as a result, all computers that are members of the site, domain, or organizational unit—or all three—execute scripts automatically when they're started or shut down.

To assign a computer startup or shutdown script, follow these steps:

1. In File Explorer, open the folder containing the script or scripts you want to use.



2. In the GPMC, right-click the GPO for the site, domain, or organizational unit with which you want to work, and then click Edit to open the policy editor for the GPO.
3. In the Computer Configuration\Policies node, double-click the Windows Settings folder, and then click Scripts.
4. To work with startup scripts, right-click Startup, and then click Properties. To work with shutdown scripts, right-click Shutdown and select Properties.



5. On the Scripts tab, you can manage command-shell batch scripts ending with the .bat or .cmd extension and scripts that use the Windows Script Host. On the PowerShell Scripts tab, you can manage Windows PowerShell scripts. When working with either tab, click Show Files.
6. Copy the files in the open File Explorer window, and then paste them into the window that opened when you clicked Show Files.
7. Click Add to open the Add A Script dialog box and assign a script. In the Script Name text box, enter the name of the script you copied to the -



Machine\Scripts\Startup or the Machine\Scripts\Shutdown folder for the related policy. In the Script Parameters text box, enter any parameters to pass to the script. Repeat this step to add other scripts.

8. During startup or shutdown, scripts are executed in the order in which they're listed in the Properties dialog box. On the Scripts tab, use the Up and Down buttons to reorder scripts as necessary. Do the same on the PowerShell Scripts tab. On the PowerShell Scripts tab, you can also use the selection list to specify whether Windows PowerShell scripts should run before or after other types of scripts.
9. If you want to edit the script name or parameters later, select the script in the Script For list, and then click Edit. To delete a script, select the script in the Script For list, and click Remove.
10. To save your changes, click OK.





# Assigning User Logon and Logoff Scripts

You can assign user scripts in one of the three following ways:

- You can assign logon and logoff scripts as part of a GPO; and as a result, all users who are members of the site, domain, or organizational unit—or all three—execute scripts automatically when they log on or log off.
- You can also assign logon scripts individually through the Active Directory Users And Computers console, which enables you to assign a separate logon script to each user or group.
- You can also assign individual logon scripts as scheduled tasks by using the Scheduled Task Wizard.

To assign a logon or logoff script in a GPO, follow these steps:

1. In File Explorer, open the folder containing the script or scripts you want to use.
2. In the GPMC, right-click the GPO for the site, domain, or organizational unit with which you want to work, and then click Edit to open the policy editor for the GPO.
3. Double-click the Windows Settings folder in the User Configuration\Policies node, and then click Scripts.
4. To work with logon scripts, right-click Logon, and then click Properties. To work with logoff scripts, right-click Logoff, and then click Properties.
5. On the Scripts tab, you can manage command-shell batch scripts ending with the .bat or .cmd extension and scripts that use the Windows Script Host. On the PowerShell Scripts tab, you can manage Windows PowerShell scripts. When working with either tab, click Show Files.
6. Copy the files in the open File Explorer window, and then paste them into the window that opened when you clicked Show Files.
7. Click Add to open the Add A Script dialog box and assign a script. In the Script Name text box, enter the name of the script you copied to the User\Scripts\Logon or the User\Scripts\Logoff folder for the related policy. In the Script Parameter text box, enter any parameters to pass to the script. Repeat this step to add other scripts.



8. During logon or logoff, scripts are executed in the order in which they're listed in the Properties dialog box. On the Scripts tab, use the Up and Down buttons to reorder scripts as necessary. Do the same on the PowerShell Scripts tab, on which you can also use the selection list to specify whether Windows PowerShell scripts should run before or after other types of scripts.
9. If you want to edit the script name or parameters later, select the script in the Script For list, and then click Edit. To delete a script, select the script in the Script For list, and then click Remove.
10. To save your changes, click OK.

## Centrally Managing Special Folders

You can centrally manage special folders used by Windows Server through folder redirection. You do this by redirecting special folders to a central network location instead of using multiple default locations on each computer. For current releases of Windows, the special folders you can manage are AppData (Roaming), Desktop, Start Menu, Documents, Pictures, Music, Videos, Favorites, Contacts, Downloads, Links, Searches, and Saved Games.



Note that even though current releases of Windows store personal folders in slightly different ways, you manage the folders in the same way within Group Policy.

You have two general options for redirection. You can redirect a special folder to the same network location for all users, or you can designate locations based on user membership in security groups. In either case, you should make sure that the network location you plan to use is available as a network share. See Chapter 25, “Implementing File Sharing,” for details on sharing data on a network.

By default, users can redirect folders no matter which computer they’re using within the domain. Windows 10 and Windows Server 2016 enable you to modify this behavior by specifying from which computers a user can access roaming profiles and redirected folders. You do this by designating certain computers as primary computers, and then configuring domain policy to restrict the downloading of profiles, redirected folders, or both to primary computers. For more information, see Chapter 11, “Managing Computers, Users and Groups.”

## Redirecting a Special Folder to a Single Location


You can redirect a special folder to a single location by following these steps:

1. In the Group Policy Management Console (GPMC), right-click the Group Policy object (GPO) for the site, domain, or organizational unit with which you want to work, and then click Edit to open the policy editor for the GPO.

**NOTE** If you’d rather create a new GPO, right-click the site, domain or organizational unit and then select Create A GPO... And Link It Here. In the New GPO dialog box, enter a name for the GPO and then select OK.

2. In the policy editor, expand the following nodes: User Configuration, Policies, Windows Settings, and Folder Redirection.
  3. Under Folder Redirection, right-click the special folder with which you want to work, such as AppData(Roaming), and then click Properties.
  4. In the Setting list on the Target tab, choose Basic—Redirect Everyone’s Folder To The Same Location.
  5. Under Target Folder Location, you have several options depending on the folder with which you’re working, and those options include the following:
- **Redirect To The User’s Home Directory** If you select this option, the folder is redirected to a subdirectory within the user’s home directory. You set the location of the user’s home directory with the %HomeDrive% and %HomePath% environment variables.

- **Create A Folder For Each User Under The Root Path** If you select this option, a folder is created for each user at the location you enter in the Root Path text box. The folder name is the user account name as specified by %UserName%. Thus, if you enter the root path value \\Zeta\UserDocuments, the folder for Williams will be located at \\Zeta\UserDocuments\Williams.



- **Redirect To The Following Location** If you select this option, the folder is redirected to the location you enter in the Root Path text box. Here, you typically want to use an environment variable to customize the folder location for each user. For example, you could use the root path value \\Zeta\UserData\%UserName%\docs. • **Redirect To The Local Userprofile Location** If you select this option, the folder is redirected to a subdirectory within the user profile directory. You set the location of the user profile with the %UserProfile% variable.

**IMPORTANT** When specifying the root path, be sure to specify the UNC path for the server and not a local path. The basic syntax for a UNC path is \\ServerName\ShareName, such as \\CorpServer38\CorpData.

6. Click the Settings tab, configure the following additional options, and then click OK to complete the process:




- **Grant The User Exclusive Rights To** Gives users full rights to access their data in the special folder.
- **Move The Contents Of FolderNameTo The New Location** Moves the data in the special folders from the individual systems on the network to the central folder or folders.
- **Also Apply Redirection Policy To** Applies the redirection policy to previous releases of Windows as well.



# Redirecting a Special Folder Based on Group Membership

You can redirect a special folder based on group membership by following these steps:

1. In the GPMC, right-click the GPO for the site, domain, or organizational unit with which you want to work, and then click Edit to open the policy editor for the GPO.
2. In the policy editor, expand the following nodes: User Configuration, Policies, Windows Settings, and Folder Redirection.
3. Under Folder Redirection, right-click the special folder with which you want to work, such as AppData(Roaming), and then click Properties.
4. In the Setting list on the Target tab, choose Advanced – Specify Locations For Various User Groups. A Security Group Membership panel is added to the Properties dialog box.



5. Click Add to open the Specify Group And Location dialog box. Or select a group entry, and then click Edit to modify its settings.



6. In the Security Group Membership text box, enter the name of the security group for which you want to configure redirection, or click Browse to find a security group to add.
7. As with basic redirection, the options available depend on the folder with which you're working and include the following:
  - **Redirect To The User's Home Directory** If you select this option, the folder is redirected to a subdirectory within the user's home directory. You set the location of the user's home directory with the %HomeDrive% and %HomePath% environment variables.
  - **Create A Folder For Each User Under The Root Path** If you select this option, a folder is created for each user at the location you enter in the Root Path text box. The folder name is the user account name as specified by %UserName%. Thus, if you enter the root path value \\Zeta\UserDocuments, the folder for Williams will be located at \\Zeta\UserDocuments\Williams.
  - **Redirect To The Following Location** If you select this option, the folder is redirected to the location you enter in the Root Path text box. Here, you typically want to use an environment variable to customize the folder location for each user. For example, you could use the root path value \\Zeta\UserData\%UserName%\docs.
  - **Redirect To The Local Userprofile Location** If you select this option, the folder is redirected to a subdirectory within the user profile directory. You set the location of the user profile with the %UserProfile% variable.
8. Click OK. Repeat steps 5–7 for other groups you want to configure.
9. When you're done creating group entries, click the Settings tab, configure the following additional options, and then click OK to complete the process:
  - **Grant The User Exclusive Rights To** Gives users full rights to access their data in the special folder.
  - **Move The Contents Of FolderNameTo The New Location** Moves the data in the special folders from the individual systems on the network to the central folder or folders.
  - **Also Apply Redirection Policy To** Applies the redirection policy to early releases of Windows as well.

## Removing Redirection

Sometimes you might want to remove redirection from a particular special folder. You remove redirection by following these steps:





1. In the GPMC, right-click the GPO for the site, domain, or organizational unit with which you want to work, and then click Edit to open the policy editor for the GPO.
2. In the policy editor, expand the following nodes: User Configuration, Policies, Windows Settings, and Folder Redirection.
3. Under Folder Redirection, right-click the special folder with which you want to work, and then click Properties.
4. Click the Settings tab, and then make sure that an appropriate Policy Removal option is selected. The following two options are available:
  - **Leave The Folder In The New Location When Policy Is Removed** When you select this option, the folder and its contents remain at the redirected location and current users are still permitted to access the folder and its contents at this location.
  - **Redirect The Folder Back To The Local Userprofile Location When Policy Is Removed** When you select this option, the folder and its contents are copied back to the original location; however, the contents aren't deleted from the previous location.
5. If you changed the Policy Removal option, click Apply, and then click the Target tab. Otherwise, just click the Target tab.
6. To remove all redirection definitions for the special folder, choose Not Configured in the Setting list.
7. To remove redirection for a particular security group, select the security group in the Security Group Membership panel, and then click Remove. Click OK.

## Deploying Software Through Group Policy

Group Policy includes basic functionality, called Software Installation policy, for deploying software. Although Software Installation policy is not designed to replace enterprise solutions such as System Center, you can use it to automate the deployment and maintenance of software in just about any size organization if your computers are running the Windows operating system.



# Getting to Know Software Installation Policy

In Group Policy, you can deploy software on a per-computer or per-user basis. Percomputer applications are available to all users of a computer and configured under Computer Configuration\Policies\Software Settings\Software Installation. Per-user applications are available to individual users and configured under User Configuration\Policies\Software Settings\Software Installation.

You deploy software in three key ways:

- **Computer assignment** Assigns the software to client computers so that it is installed when the computer starts. This technique requires no user intervention, but it does require a restart to install the software. Installed software is then available to all users on the computer.
- **User assignment** Assigns the software to users so that it is installed when a user logs on. This technique requires no user intervention, but it does require the user to log on to install or advertise the software. The software is associated with the user only and not the computer.
- **User publishing** Publishes the software so that users can install it manually through Programs And Features. This technique requires the user to explicitly install software or activate installation. The software is associated with the user only.

When you use user assignment or user publishing, you can advertise the software so that a computer can install the software when it is first used. With advertisements, the software can be installed automatically in the following situations:

- When a user accesses a document that requires the software
- When a user opens a shortcut to the application
- When another application requires a component of the software

When you configure Software Installation policy, you should generally not use existing GPOs. Instead, you should create GPOs that configure software installation, and then link those GPOs to the appropriate containers in Group Policy. When you use this approach, it's much easier to redeploy software and apply updates.

After you create a GPO for your software deployment, you should set up a distribution point. A distribution point is a shared folder that is available to the computers and users to which you are deploying software. With basic applications, you prepare the distribution point by copying the installer package file and all required application files to the share and configuring permissions so that these files can be accessed. With other applications, you may need to prepare the distribution point by performing an administrative installation to the share. The advantage of an administrative installation



is that the software can be updated and redeployed through Software Installation policy.

You can update applications deployed through Software Installation policy by using an update or service pack or by deploying a new version of the application. Each task is performed in a slightly different way.



# Deploying Software Throughout Your Organization

Software Installation policy uses either Windows Installer Packages (.msi) or ZAW Down-Level Application Packages (.zap). When you use computer assignment, user assignment, or user publishing, you can deploy software by using Windows Installer Packages. When you use user publishing, you can deploy software by using either Windows Installer Packages or ZAW Down-Level Application Packages. With either technique, you must set file permissions on the installer package so that the appropriate computer and user accounts have read access.

Because Software Installation policy is applied only during foreground processing of policy settings, per-computer application deployments are processed at startup and peruser application deployments are processed at logon. You can customize installation by using transform (.mst) files. Transform files modify the installation process according to the settings you defined for specific computers and users.

You can deploy software by following these steps:

1. In the GPMC, right-click the GPO you want to use for the deployment, and then click Edit.
2. In the policy editor, open Computer Configuration\Policies\Software Settings\Software Installation or User Configuration\Policies\Software Settings\Software Installation as appropriate for the type of software deployment.
3. Right-click Software Installation. On the shortcut menu, click New, and then click Package.
4. In the Open dialog box, go to the network share where your package is located, click the package to select it, and then click Open.

**NOTE** Windows Installer Packages (.msi) is selected by default in the Files Of Type list. If you are performing a user publishing deployment, you can also choose ZAW Down-Level Application Packages (.zap) as the file type.

5. In the Deploy Software dialog box, select one of the following deployment methods, and then click OK:
  - **Published** To publish the application without modifications
  - **Assigned** To assign the application without modifications
  - **Advanced** To deploy the application by using advanced configuration options

Deploy Software

Select deployment method:

☒ Published

☐ Assigned

☐ Advanced

---

Select this option to Assign the application without modifications.

OK Cancel

# Configuring Software Deployment Options

You can view and set general options for a software package by following these steps:

1. In the GPMC, right-click the GPO you want to use for the deployment, and then click Edit.
2. In the policy editor, access Computer Configuration\Policies\Software Settings\Software Installation or User Configuration\Policies\Software Settings\Software Installation as appropriate for the type of software deployment.
3. Double-click the Software Installation package. In the Properties dialog box, review or modify software deployment options.
4. On the Deployment tab, you can change the deployment type and configure the following deployment and installation options:
  - **Auto-Install This Application By File Extension Activation** Advertises any file extensions associated with this package for install-on-first-use deployment. This option is selected by default.
  - **Uninstall This Application When It Falls Out Of The Scope Of Management** Removes the application if it no longer applies to the user.
  - **Do Not Display This Package In The Add/Remove Programs Control Panel** Prevents the application from appearing in Add/Remove Programs, which prevents a user from uninstalling an application.
  - **Install This Application At Logon** Configures full installation—rather than advertisement—of an application when the user logs on. This option cannot be set when you publish a package for users.
  - **Installation User Interface Options** Controls how the installation is performed. With the default setting, Maximum, the user gets all setup screens and messages during installation. With the Basic option, the user gets only error and completion messages during installation.
5. Click OK.

## Updating Deployed Software

When an application uses a Windows Installer package, you can apply an update or service pack to a deployed application by following these steps:

1. After you obtain an .msi file or .msp (security update) file containing the update or service pack to be applied, copy the .msi or .msp file and any new



installation files to the folder containing the original .msi file. Overwrite any duplicate files as necessary.

2. In the GPMC, right-click the GPO you want to use for the deployment, and then click Edit.
3. In the policy editor, access Computer Configuration\Policies\Software Settings\Software Installation or User Configuration\Policies\Software Settings\Software Installation as appropriate for the type of software deployment.
4. Right-click the package with which you want to work. On the shortcut menu, click All Tasks, and then click Redeploy Application.
5. When prompted to confirm the action, click Yes. The application is then redeployed to all users and computers as appropriate for the GPO with which you are working.

When an application uses an Installer package that is not Windows-based, you can update a deployed application or apply a service pack by following these steps:

1. In the GPMC, right-click the GPO you want to use for the deployment, and then click Edit.
2. In the policy editor, access Computer Configuration\Policies\Software Settings\Software Installation or User Configuration\Policies\Software Settings\Software Installation as appropriate for the type of software deployment.
3. Right-click the package. On the shortcut menu, click All Tasks, and then click Remove. Click OK to accept the default option of immediate removal.
4. Copy the new .zap file and all related files to a network share, and redeploy the application.

## Upgrading Deployed Software

You can upgrade a previously deployed application to a new version by following these steps:

1. Obtain a Windows Installer file for the new software version, and copy it along with all required files to a network share. Alternatively, you can perform an administrative installation to the network share.
2. In the GPMC, right-click the GPO you want to use for the deployment, and then click Edit.
3. In the policy editor, access Computer Configuration\Policies\Software Settings\Software Installation or User Configuration\Policies\Software



Settings\Software Installation as appropriate for the type of software deployment.

4. Right-click Software Installation. On the shortcut menu, click New, and then click Package. Create an assigned or published application by using the Windows Installer file for the new software version.
5. Right-click the upgrade package, and then click Properties. On the Upgrades tab, click Add. In the Add Upgrade Package dialog box, do one of the following:
  - If the original application and the upgrade are in the current GPO, select Current Group Policy Object, and then select the previously deployed application in the Package To Upgrade list.
  - If the original application and the upgrade are in different GPOs, select A Specific GPO, click Browse, and then select the GPO from the Browse For A Group Policy Object dialog box. Select the previously deployed application in the Package To Upgrade list.
6. Choose an upgrade option. If you want to replace the application with the new version, select Uninstall The Existing Package, Then Install The Upgrade Package. If you want to perform an in-place upgrade over the existing installation, select Package Can Upgrade Over The Existing Package.
7. Click OK to close the Add Upgrade Package dialog box. If you want to make this a required upgrade, select the Required Upgrade For Existing Packages check box, and then click OK to close the upgrade package's Properties dialog box.

## Automatically Configuring Work Folders

Computers that are members of a workplace can access internal network resources, such as internal websites and business applications. Work Folders enable users to synchronize their corporate data to their devices and vice versa. Those devices can be joined to the corporate domain or a workplace. Devices access Work Folders via a remote web gateway running on Microsoft Internet Information Services (IIS).

To deploy Work Folders, you add the File And Storage Services \ Work Folders role to a file server, and then configure Work Folders by using Server Manager. Afterward, you can use policy settings to control related options, such as the server to which users can connect remotely and access Work Folders. You control the connection server in one of two ways:

- By specifying the exact URL of a file server hosting the Work Folders for the user, such as <https://server12.imagedlands.com>.
- By specifying the URL used within your organization for Work Folders discovery, such as <https://workfolders.imagedlands.com>.

**REAL WORLD** Clients use secure encrypted communications to connect to work folders as long as the file servers hosting the Work Folders have valid SSL certificates. When a device initiates an SSL connection, the server sends the certificate to the client. The client evaluates the certificate and continues only if the certificate is valid and can be trusted. If you configure a connection to an exact URL, the client can connect directly to the specified sever and synchronize data in Work Folders. The server's certificate must have a Common Name (CN) or a Subject Alternative Name (SAN) that matches the host header in the request. For example, if the client makes a request to <https://server18.imagedlands.com>, the CN or SAN must be [server18.imagedlands.com](https://server18.imagedlands.com).

In Group Policy, you specify the URL used within your organization for Work Folders discovery by using the Specify Work Folders Settings policy found under Administrative Templates policies for User Configuration\Windows Components\Work Folders. Any server configured with Work Folders acts as a discovery server by default. If you configure a discovery URL, a client connects to one of several servers, and the email address of the user is used to discover which specific server hosts the Work Folders for the client. The client is then connected to this server. Each discovery server will need to have a certificate with multiple Subject Alternative Names, which includes the server name and the discovery name. For example, if a client makes a request to <https://workfolders.imagedlands.com> and connects to [FileServer87.imagedlands.com](https://FileServer87.imagedlands.com), the server's certificate must have a CN or SAN of [fileserver87.imagedlands.com](https://fileserver87.imagedlands.com) and a SAN of [workfolders.imagedlands.com](https://workfolders.imagedlands.com).

If you want to configure Work Folders in Group Policy, use the following technique:

1. Access Group Policy for the system, site, domain, or OU with which you want to work. Next, access the Work Folders node by using the Administrative Templates policies for User Configuration under Windows Components\Work Folders.
2. Double-click Specify Work Folders Settings, and then select Enabled.

Specify Work Folders settings

Previous Setting Next Setting

☐ Not Configured Comment:

☒ Enabled

☐ Disabled

Supported on: At least Windows 8.1 or Windows RT 8.1

Options:

Work Folders URL:

☒ Force automatic setup

Help:

This policy setting specifies the Work Folders server for affected users, as well as whether or not users are allowed to change settings when setting up Work Folders on a domain-joined computer.

If you enable this policy setting, affected users user receive Work Folders settings when they sign in to a domain-joined PC.

3. In the World Folders URL text box, enter the URL of the file server that hosts the Work Folders for the user or the URL used within your organization for Work Folders discovery.
4. If you want to prevent users from changing settings when setting up Work Folders, select Force Automatic Setup.
5. Click OK.

## Automatically Enrolling Computer and User Certificates

A server designated as a certificate authority (CA) is responsible for issuing digital certificates and managing certificate revocation lists (CRLs). Servers running Windows Server can be configured as certificate authorities by installing Active Directory Certificate Services. Computers and users can use certificates for authentication and encryption.

In an enterprise configuration, enterprise CAs are used for automatic enrollment. This means authorized users and computers can request a certificate, and the certificate authority can automatically process the certificate request so that the users and computers can immediately install the certificate.

Group Policy controls the way automatic enrollment works. When you install enterprise CAs, automatic enrollment policies for users and computers are enabled automatically.

The policy for computer certificate enrollment is Certificate Services Client—AutoEnrollment Settings under Computer Configuration\Policies\Windows Settings\Security






Settings\Public Key Policies. The policy for user certificate enrollment is Certificate Services Client—Auto-Enrollment under User Configuration\Policies\Windows Settings\Security Settings\Public Key Policies.

You can configure automatic enrollment by following these steps:

1. In the GPMC, right-click the GPO with which you want to work, and then click Edit.
  2. In the policy editor, access User Configuration\Policies\Windows Settings\Security Settings\Public Key Policies or Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies as appropriate for the type of policy you want to review.
  3. Double-click Certificate Services Client—Auto-Enrollment. To disable automatic enrollment, select Disabled from the Configuration Model list, click OK, and then skip the remaining steps in this procedure. To enable automatic enrollment, select Enabled from the Configuration Model list.
  4. To automatically renew expired certificates, update pending certificates, and remove revoked certificates, select the related check box.
  5. To ensure that the latest version of certificate templates are requested and used, select the Update Certificates That Use Certificate Templates check box.
  6. To notify users when a certificate is about to expire, specify when notifications are sent using the box provided. By default, notifications are sent when 10 percent of the certificate lifetime remains.
  7. Click OK to save your settings.
- 





## Chapter 30. Implementing Print Services

As an administrator, you need to take two main steps to enable users throughout a network to access print devices connected to Windows Server 2016. First, you need to set up a print server, and then you need to use the print server to share print devices on the network.

This chapter covers the essentials of setting up shared printing and describes how users access shared printing on the network. You'll also find advice on administering printers and troubleshooting printer problems, which is where we'll begin.



# Managing the Print and Document Services Role

A print server provides a central location for sharing printers on a network. When many users require access to the same printers, you should configure print servers in the domain. In earlier releases of the Windows Server operating system, all servers were installed with basic print services. With Windows Server 2016, you must specifically configure a server to be a print server.





## Using Print Devices

Two types of print devices are used on a network:

- **Local print device** A print device that's physically attached to a user's computer and employed only by the user who's logged on to that computer.
- **Network print device** A print device that's set up for remote access over the network. This can be a print device attached directly to a print server or a print device attached directly to the network through a network interface card (NIC).

**NOTE** The key difference between a local printer and a network printer is that a local printer isn't shared. You can easily make a local printer a network printer. To learn how to do this, see "Starting and Stopping Printer Sharing" in Chapter 31.

You install new network printers on print servers or as separate print devices attached to the network. A print server is a workstation or server configured to share one or more printers. These printers can be physically attached to the computer or connected to the network. The disadvantage of running a print server with a workstation operating system instead of a server operating system is the limited number of allowed connections. With Windows Server 2016, you don't have to worry about operating system-enforced connection limits.

You can configure any Windows Server 2016 system as a print server. The print server's primary job is to share print devices on the network and to handle print spooling. The main advantages of print servers are that printers have a centrally managed print queue and you don't have to install printer drivers on client systems.

You don't have to use a print server, however. You can connect users directly to a network-attached printer. When you do this, the network printer is handled much like a local printer attached directly to the user's computer. The key differences are that multiple users can connect to the printer and that each user has a different print queue. Each individual print queue is managed separately, which can make administration and problem resolution difficult.

## Printing Essentials

An understanding of how printing works can go a long way when you're trying to troubleshoot printer problems. When you print documents, many processes, drivers, and devices work together to print the documents. If you use a printer connected to a print server, the key operations are as follows:

- **Printer driver** When you print a document in an application, your computer uses a printer driver to handle the printing process. If the print device is attached to your





computer physically, the printer driver is accessed from a local disk drive. If the print device is located on a remote computer, the printer driver might be downloaded from the remote computer. The availability of printer drivers on the remote computer is configurable by operating system and chip architecture. If the computer can't obtain the latest printer driver, an administrator probably hasn't enabled the driver for the computer's operating system. For more information, see "Managing Printer Drivers" in Chapter 31.

- **Local print spool and print processor** The application from which you print uses the printer driver to translate the document into a file format that the selected print device can understand. Your computer then passes the document to the local print spooler. The local spooler in turn passes the document to a print processor, which creates the raw print data necessary for printing on the print device.
- **Print router and print spooler on the print server** The raw data is passed back to the local print spooler. If you're printing to a remote printer, the raw data is then routed to the print spooler on the print server. On Windows Server 2016 systems, the printer router, Winspool.drv, handles the tasks of locating the remote printer, routing print jobs, and downloading printer drivers to the local system if necessary. If any of these tasks fails, the print router is usually the culprit. See "Solving spooling problems" and "Setting Printer Access Permissions" in Chapter 31 to learn about possible fixes for this problem. If these procedures don't work, you might want to replace or restore Winspool.drv.

The main reason for downloading printer drivers to clients is to maintain a single location for installing driver updates. This way, instead of having to install a new driver on all the client systems, you install the driver on the print server and enable clients to download the new driver. For more information on working with printer drivers, see "Managing Printer Drivers" in Chapter 31.

- **Printer (print queue)** The document goes from the print spooler to the printer stack—which in some operating systems is called the print queue—for the selected print device. When in the queue, the document is referred to as a print job—a task for the print spooler to handle. The length of time the document waits in the printer stack is based on its priority and position within the printer stack. For more information, see "Scheduling and Prioritizing Print Jobs" in Chapter 31.
- **Print monitor** When the document reaches the top of the printer stack, the print monitor sends the document to the print device, where it's actually printed. If the printer is configured to notify users that the document has been printed, you get a confirmation message.





The specific print monitor used by Windows Server 2016 depends on the print device configuration and type. You might also find monitors from the print device manufacturer. This dynamic-link library (DLL) is required to print to the print device. If it's corrupted or missing, you might need to reinstall it.

- **Print device** The print device is the physical device that prints documents on paper. Common print device problems and display errors include Insert Paper Into Tray X, Low Toner, Out Of Paper, Out Of Toner or Out Of Ink, Paper Jam, and Printer Offline.

Group Policy can affect your ability to install and manage printers. If you're having problems and believe they're related to Group Policy, you should examine the key policies in the following locations:

- Computer Configuration\Administrative Templates\Printers
- User Configuration\Administrative Templates\Control Panel\Printers
- User Configuration\Administrative Templates\Start Menu And Taskbar

**REAL WORLD** Windows Server 2016 is a 64-bit operating system and a print queue running on this operating system cannot function without a native 64-bit printer driver. Because of this, you should verify the availability of required 64-bit print drivers before migrating print queues from servers running 32-bit operating systems to Windows Server 2016. Keep in mind that if your organization has older printers still in use, those printers might have third-party 32-bit drivers and there might not be 64-bit equivalents.

## Configuring Print Servers

You configure a server as a print server by adding the Print and Document Services role and configuring this role to use one or more of the following role services:

- **Print Server** Configures the server as a print server and installs the Print Management console. You can use the Print Management console to manage multiple printers and print servers, to migrate printers to and from other print servers, and to manage print jobs.
- **Line Printer Daemon (LPD) Service** Enables UNIX-based computers or other computers using the Line Printer Remote (LPR) service to print to shared printers on the server.
- **Internet Printing** Creates a website where authorized users can manage print jobs on the server. It also lets users who have Internet Printing Client installed to connect and print to shared printers on the server by using the Internet Printing Protocol




(IPP). The default Internet address for Internet Printing is `http:// ServerName /Printers`, where `ServerName` is a placeholder for the internal or external server name, such as `http://PrintServer15/Printers` or `http://www.imaginedlands.com/Printers`.

- **Distributed Scan Server** Establishes the server as a scan server, which is used to run scan processes. Scan processes are rules that define scan settings and control delivery of scanned documents on your network. The Scan Management snap-in is installed when you install this role service. Scan Management enables you to manage Web Services on Devices (WSD)–enabled scanners, scan servers, and scan processes.

You can add the Print and Document Services role to a server by following these steps:

1. In Server Manager, click Manage, and then click Add Roles And Features, or select Add Roles And Features in the Quick Start pane. This starts the Add Roles And Features Wizard. If the wizard displays the Before You Begin page, read the Welcome text, and then select Next.
2. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Select Next.




3. On the Server Selection page, you can choose to install roles and features on running servers or virtual hard disks. Either select a server from the server pool or select a server from the server pool on which to mount a virtual hard disk (VHD). If you are adding roles and features to a VHD, click Browse, and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, select Next.








4. On the Select Roles page, select Print And Document Services. When you are ready to continue, select Next three times.



5. On the Select Role Services page, select one or more role services to install. To allow for interoperability with UNIX, be sure to add LPD Service. Click Next.
6. When you install Internet Printing, you must also install Web Server (IIS) and some related components. You are prompted to automatically add the required role services. Select Add Features to continue.



**NOTE** If the server on which you want to install the Print And Document Services role doesn't have all the required binary source files, the server gets the files via Windows Update by default or from a location specified in Group Policy. You also can specify an alternate path for the required source files by clicking the Specify An Alternate Source Path link, entering that alternate path in the box provided, and then clicking OK. For network shares, enter the UNC path to the share, such as \\CorpServer82\\WinServer2012\\. For mounted Windows images, enter the WIM

path prefixed with WIM: and including the index of the image to use, such as WIM:\\CorpServer82\\WinServer2012\\install.wim:4.

7. After you review the installation options and save them as necessary, select Install to begin the installation process. The Installation Progress page tracks the progress of the installation. If you close the wizard, select the Notifications icon in Server Manager, and then select the link provided to reopen the wizard.
8. When Setup finishes installing the Print And Document Services role, the Installation Progress page will be updated to reflect this. Review the installation details to ensure that all phases of the installation were completed successfully.
9. When you install Distributed Scan Server, you'll get a notification that additional configuration is required. Select the link provided to open Scan Management, which you can use to specify scanners to manage. In Scan Management, expand the Scan Management node in the left pane, right-click the Managed Scanners node, and then select Manage to open the Add Or Remove Scanners dialog box. Use the Add Or Remove Scanners dialog box to identify the distributed scanners in your organization.

When you install Distributed Scan Server, a security group called Scan Operators is added to the Users container in Active Directory Domain Services (AD DS) for the current logon domain. Any users who need to manage the Distributed Scan service should be added to this group.

## Enabling and Disabling File and Printer Sharing


File and printer sharing settings control access to file shares and printers that are attached to a computer. You can manage a computer's file and printer sharing configuration by following these steps:

1. In File Explorer, select Network in the left pane. On the toolbar in Network Explorer, select Network, and then select Network And Sharing Center.






2. In Network And Sharing Center, select Change Advanced Sharing Settings in the left pane. Select the network profile for the network for which file and printer sharing should be enabled. Typically, this is the Domain profile.



3. Standard file and printer sharing controls network access to shared resources. To configure standard file sharing, do one of the following:


- Select Turn On File And Printer Sharing to enable file and printer sharing.
- Select Turn Off File And Printer Sharing to disable file and printer sharing.



## Getting Started with Print Management


Print Management should be your tool of choice for working with printers and print servers. After you install the Print and Document Services role, Print Management is available on the Tools menu in Server Manager as a stand-alone console. You can also add Print Management as a snap-in to any custom console you create.

By using Print Management, you can install, view, and manage the printers and Windows print servers in your organization. Print Management also displays the status of printers and print servers. When you expand a server-level node and select the Printers node, you get a list of printers the server is hosting. If you are accessing the selected print server by using a Remote Desktop connection, you might also find entries for printers being redirected from your logon computer. Redirected printers are listed clearly as such with a (redirected) suffix.



By default, Print Management enables you to manage local print servers. You can manage and monitor other print servers in the organization by adding them to the console. Additionally, to manage a remote print server, you must be a member of the local Administrators group on the print server or a member of the Domain Admins group in the domain of which the print server is a member.

When you select a print server's Printers node, the main pane lists the associated printer queues by printer name, queue status, number of jobs in a queue, and server name. If you right-click Printers, and then click Show Extended View, you can turn on Extended view. Extended view makes it easy to track the status of both printers and print jobs by displaying information about the status of a print job, its owner, the number of pages, the size of the job, when it was submitted, its port, and its priority.




In addition, when a printer has a webpage, Extended view displays a Printer Web Page tab that lets you directly access the printer's webpage. This webpage provides details about the printer's status, its physical properties, and its configuration, and it sometimes allows remote administration.



You can add print servers to Print Management by following these steps:

1. In Print Management, right-click the Print Servers node in the left pane, and then click Add/Remove Servers.
2. In the Add/Remove Servers dialog box, you'll find a list of the print servers you previously added.



Do one of the following, and then click Add To List:

- In the Add Servers list, enter or paste the names of the print servers you want to add. Use commas to separate computer names.
  - Click Browse to display the Select Print Server dialog box. Click the print server you want to use, and then click Select Server.
3. Repeat the previous step as necessary, and then click OK.



You can remove print servers from Print Management by following these steps:

4. In Print Management, right-click the Print Servers node in the left pane, and then click Add/Remove Servers.
5. In the Add/Remove Servers dialog box, you'll find a list of the print servers that are being monitored. Under Print Servers, select one or more servers, and then click Remove.





# Installing Printers


The following sections examine techniques you can use to install printers. Windows Server 2016 makes it possible for you to install and manage printers anywhere on the network. To install or configure a new printer on Windows Server 2016, you must be a member of the Administrators, Print Operators, or Server Operators group. To connect to a printer and print documents, you must have the appropriate access permissions. See “Setting Printer Access Permissions” in Chapter 31 for details.



## Using Autoinstall


Print Management can automatically detect all network printers located on the same subnet as the computer on which the console is running. After detection, Print Management can automatically install the appropriate printer drivers, set up print queues, and share the printers. To automatically install network printers and configure a print server, follow these steps:

1. Start Print Management by clicking Tools in Server Manager, and then clicking Print Management.
2. In Print Management, expand the Print Servers node by double-clicking it.
3. Right-click the entry for the local or remote server with which you want to work, and then click Add Printer to start the Network Printer Installation Wizard.
4. On the Print Installation page, select Search The Network For Printers, and then click Next.



5. The wizard searches the local subnet for network printers. If printers are found, you'll get a list of printers by name and IP address. Click a printer to install, click Next, and then continue this procedure to install the automatically detected printer.

If a printer you want to use is not listed, you should ensure that the printer is powered on and is online and then repeat this procedure. If a printer you want to use is powered on and online but is not listed, see "Installing Network-Attached Print Devices" to complete the installation.



6. The wizard automatically detects the TCP/IP port configuration of the selected printer and then communicates with the print to obtain the information required to configure the printer. Afterward, the wizard sets the default name and share name of the printer. The printer is configured as a shared resource as well.


The printer name is the name you encounter when you work with the printer in Print Management. The share name is the name users view when they work with the printer.

Optionally, enter information in the Location and Comment text boxes that will help users locate and identify the printer. For example, you might want to specify the printer location as: Room 314 in Building 7.



**NOTE** The printer name and share name can be up to 256 characters and can include spaces. In a large organization, you'll want the share name to be logical and helpful in locating the printer. For example, you might want to give the name Twelfth Floor NE to the printer that points to a print device in the northeast corner of the twelfth floor.

7. The next page lets you review the settings. When you're ready to complete the installation, click Next.
8. When you share a printer, Windows Server automatically makes drivers available so that users can download them when they first connect to the printer. The status page should confirm that printer driver and printer installation were successful. If there was a problem with the installation, note the errors provided. For example, someone might have powered off the printer while you were trying to configure it. If so, you'll need to power the printer back on and repeat this procedure.
9. If you'd like to print a test page on the printer, select the Print Test Page check box, and then click Finish. Otherwise, just click Finish.
10. By default, the printer share is not listed in Active Directory. Listing the printer share in Active Directory enables users to search for and find the printer more easily. If you want the printer share to be listed in Active Directory, select the Printers node in the left pane, right-click the printer in the main window, and then select List In Directory.



11. By default, print jobs are sent to the print server where they are rendered and then sent to the printer. You can change this behavior by using Branch Office Direct Printing. With Branch Office Direct Printing, print jobs are rendered on client computers and then sent directly to the printer. If you want to enable



direct printing, select the Printers node in the left pane, right-click the printer in the main window, and then select Enable Branch Office Direct Printing.

## Installing Physically-Attached Print Devices

Most physically attached print devices are connected to a computer directly through a USB cable. You can configure physically attached printers as local print devices or as network print devices. The key difference is that a local device is accessible only to users logged on to the computer and a network device is accessible to network users as a shared print device. Remember that the workstation or server you're logged on to becomes the print server for the device you're configuring. If the computer is sleeping or turned off, the printer will not be available.

You can install physically attached print devices locally by logging on to the print server you want to configure; you can install the print devices remotely through Remote Desktop. If you're configuring a local Plug and Play printer and are logged on to the print server, installing a print device is a snap. After the printer is installed, you need to configure it for use.

You can install and configure a print device by following these steps:

1. Power on the printer, and then connect the print device to the server by using the appropriate cable.
2. If Windows Server automatically detects the print device, Windows begins installing the device and the necessary drivers. If the necessary drivers aren't found, you might need to insert the printer's driver disc into the CD/DVD drive.
3. If Windows Server doesn't detect the print device automatically, you need to install the print device manually as described in the next set of instructions.
4. After you install the printer, you can configure the printer. In Print Management, expand the Print Servers node and the node for the server with which you want to work. When you select the Printers node for the server you are configuring, you'll find a list of available printers in the main pane. Right-click the printer you want to configure, and then click Manage Sharing. This displays the printer's Properties dialog box with the Sharing tab selected.
5. When you select the Share This Printer check box, Windows Server sets the default share name to the name of the printer. You can enter a different name for the printer share in the Share Name text box.
6. By default, the Render Print Jobs On Client Computers check box is selected, which configures the printer for Branch Office Direct Printing. With Branch



Office Direct Printing, print jobs are rendered on client computers and then sent directly to the printer. If you want print jobs to be sent to the print server for rendering and then sent to the printer, clear the Render Print Jobs On Client Computers check box.

7. Listing the printer share in Active Directory enables users to search for and find the printer more easily. If you want the printer share to be listed in Active Directory, select the List In The Directory check box.
8. Click OK.

Sometimes Windows Server won't detect your printer. In this case, follow these steps to install the print device:

1. In Print Management, expand the Print Servers node and the node for the server with which you want to work.
2. Right-click the server's Printers node, and then click Add Printer to start the Network Printer Installation Wizard.
3. On the Printer Installation page, select Add A New Printer Using An Existing Port, and then choose the appropriate LPT, COM, or USB port. You can also print to a file. If you do, Windows Server 2016 prompts users for a file name each time they print. Click Next.
4. On the Printer Driver page, choose one of the following options:
  - If Windows detected the printer type on the selected port and a compatible driver was found automatically, a printer driver is listed that reflects the printer manufacturer and model and the Use The Printer Driver That The Wizard Selected option is selected by default. To accept this setting, click Next.
  - If a compatible driver is not available and you want to choose an existing driver installed on the computer, select the Use An Existing Driver On The Computer option. After you choose the appropriate driver from the selection list, click Next. • If a compatible driver is not available and you want to install a new driver, select Install A New Driver, and then click Next. You must now specify the print device manufacturer and model. This enables Windows Server 2016 to assign a printer driver to the print device. After you choose a print device manufacturer, choose a printer model.

**REAL WORLD** If the device manufacturer and model you're using aren't displayed in the list, click Windows Update. Windows will then connect to the Windows Update Web site to update the list of printers to show additional models. This automatic driver provisioning process can take several minutes. When the update process is complete, you should then be able to select your printer manufacturer and model. If





you can't, download the driver from the manufacturer's website, and then extract the driver files. Click Have Disk. In the Install From Disk dialog box, click Browse. In the Locate File dialog box, locate the .inf driver file for the device, and then click Open.

**NOTE** If a driver for the specific printer model you're using isn't available, you often can select a generic driver or a driver for a similar print device. Consult the print device documentation for pointers.

5. Assign a name to the printer. This is the name that will be listed in Print Management.
6. Specify whether the printer is available to remote users. To create a printer accessible to remote users, select the Share This Printer option, and then enter a name for the shared resource. In a large organization you'll want the share name to be logical and helpful in locating the printer. For example, you could give the name Twelfth Floor NE to the printer that points to the print device in the northeast corner of the twelfth floor.
7. If you like, you can enter a location description and comment. This information can help users find a printer and determine its capabilities. Click Next.
8. The final page lets you review the settings. When you're ready to complete the installation, click Next.
9. After Windows installs the printer driver and configures the printer, you'll get a status page. Ensure that the driver and printer installation succeeded before continuing. If there were errors, you need to correct any problems and repeat this process. To test the printer, select Print Test Page, and then click Finish.

When the Network Printer Installation Wizard finishes installing the new printer, the Printers folder will have an additional icon with the name set the way you specified. You can change the printer properties and check printer status at any time. For more information, see "Configuring Printer Properties" in Chapter 31.

**TIP** If you repeat this process, you can create additional printers for the same print device. All you need to do is change the printer name and share name. Having additional printers for a single print device makes it possible for you to set different properties to serve different needs. For example, you could have a highpriority printer for print jobs that need to be printed immediately and a lowpriority printer for print jobs that aren't as urgent.

# Installing Network-Attached Print Devices

A network-attached print device is attached directly to the network through a network adapter card or a wireless network card. Network-attached printers are configured as network print devices so that they're accessible to network users as shared print devices. Remember that the server on which you configure the print device becomes the print server for the device you're configuring.

Install a network-attached print device by following these steps:

1. In Print Management, expand the Print Servers node and the node for the server with which you want to work.
2. Right-click the server's Printers node, and then click Add Printer to start the Network Printer Installation Wizard.
3. On the Printer Installation page, select Add A TCP/IP Or Web Services Printer By IP Address Or Hostname, and then click Next.
4. On the Printer Address page, choose one of the following options in the Type Of Device list:
  - **Autodetect** Choose this option if you are unsure of the printer device type. Windows Server will try to detect the type of device automatically.
  - **TCP/IP Device** Choose this option if you are sure the printer is a TCP/IP device.
  - **Web Services Printer** Choose this option if you are sure the printer is a Web Services for Devices (WSD) capable printer.
  - **Web Services Secure Printer** Choose this option if you are sure the printer supports WSD Secure Printing.

**REAL WORLD** With WSD Secure Printing, print servers can create a private secure channel to the print device on the network without the need for additional security technologies such as IPsec. However, users and computers that work with the secure printer must be members of an Active Directory domain. You use the domain settings to manage printer permissions and Active Directory Domain Services acts as the trust arbitrator between the print server and the printer.

5. Enter the host name or IP address for the printer, such as 192.168.1.90. With the Autodetect and TCP/IP Device options, the wizard sets the port name to the same value, but you can also choose to use a different value.

**TIP** The port name doesn't matter as long as it's unique on the server. If you're configuring multiple printers on the print server, be sure to record the port-to-printer mapping.



6. Click Next, and the wizard attempts to contact the printer and automatically configure the print device. If the wizard is unable to detect the print device, be sure that the following are true:
    - You selected the correct type of print device.
    - The print device is turned on and connected to the network.
    - The printer is configured properly.
    - You entered the correct IP address or printer name.
  7. If the device type, IP address, or printer name is incorrect, click OK to close the warning dialog box. Next, click Back, and then reenter the required printer information.
  8. On the Printer Driver page, choose one of the following options:
    - If Windows detected the printer type on the selected port and a compatible driver was found automatically, a printer driver is listed that reflects the printer manufacturer and model and the Use The Printer Driver That The Wizard Selected option is selected by default. To accept this setting, click Next.
    - If a compatible driver is not available and you want to choose an existing driver installed on the computer, select the Use An Existing Driver option. After you use the selection list to choose the appropriate driver, click Next.
    - If a compatible driver is not available and you want to install a new driver, select Install A New Driver, and then click Next. Specify the print device manufacturer to enable Windows Server to assign a printer driver to the print device. After you choose a print device manufacturer, choose a printer model.
- REAL WORLD** If the device manufacturer and model you're using aren't displayed in the list, click Windows Update. Windows will then connect to the Windows Update Web site to update the list of printers to show additional models. This feature is part of the automatic driver provisioning feature. It can take several minutes to retrieve the updated list. You should then be able to select your printer manufacturer and model. If you can't, download the driver from the manufacturer's website, and then extract the driver files. Click Have Disk. In the Install From Disk dialog box, click Browse. In the Locate File dialog box, locate the .inf driver file for the device, and then click Open.
9. Assign a name to the printer. This is the name that will be listed in Print Management.
  10. Specify whether the printer is available to remote users. To create a printer accessible to remote users, select the Share Name option, and then enter a





name for the shared resource. In a large organization, you should use a share name that is logical and helpful in locating the printer. For example, Twelfth Floor NE would be a good name for the printer that points to the print device in the

northeast corner of the twelfth floor.

11. If you like, you can enter a location description and comment. This information can help users find a printer and determine its capabilities. Click Next.
12. The final page lets you review the settings. When you're ready to complete the installation, click Next.
13. After Windows installs the printer driver and configures the printer, you'll get a status page. Ensure that the driver and printer installation succeeded before continuing. If there were errors, you need to correct any problems and repeat this process. To test the printer, select Print Test Page, and then click Finish. To install another printer, select Add Another Printer, and then click Finish.
14. By default, the printer share is not listed in Active Directory. Listing the printer share in Active Directory makes it possible for users to search for and find the printer more easily. If you want the printer share to be listed in Active Directory, select the Printers node in the left pane, right-click the printer in the main window, and then select List In Directory.
15. By default, print jobs are sent to the print server where they are rendered and then sent to the printer. You can change this behavior by using Branch Office Direct Printing. With Branch Office Direct Printing, print jobs are rendered on client computers and then sent directly to the printer. If you want to enable direct printing, select the Printers node in the left pane, right-click the printer in the main window, and then select Enable Branch Office Direct Printing.

When the Network Printer Installation Wizard finishes installing the new printer, the Printers folder will have an additional icon with the name set the way you specified. You can change the printer properties and check printer status at any time. For more information, see “Configuring Printer Properties” in Chapter 31.

**TIP** If you repeat this process, you can create additional printers for the same print device. All you need to do is change the printer name and share name. Having additional printers for a single print device makes it possible for you to set different properties to serve different needs. For example, you could have a highpriority printer for print jobs that need to be printed immediately and a lowpriority printer for print jobs that aren't as urgent.







## Connecting to Network Printers

After you create a network printer, remote users can connect to it and use it much as they do any other printer. You need to set up a connection on a user-by-user basis or have users do this themselves. To create a connection to the printer on a Windows 7 system, follow these steps:

1. With the user logged on, click Start, and then click Devices And Printers. In Devices And Printers, click Add A Printer to start the Add Printer Wizard.
2. Select Add A Network, Wireless Or Bluetooth Printer. The wizard searches for available devices.
3. If the printer you want is listed in the Select A Printer list, click it, and then click Next.
4. If the printer you want is not listed in the Select A Printer list, click The Printer That I Want Isn't Listed. On the Find A Printer By Name Or TCP/IP Address page, do one of the following:
  - To browse the network for shared printers, choose Find A Printer In The Directory, Based On Location Or Feature, and then click Next. Click the printer to use, and then click Select.
  - To specify the printer to use by its share path, choose Select A Shared Printer By Name. Enter the UNC path to the shared printer, such as **\\PrintServer12\\Twelfth Floor NE** , or the Web path to an Internet Printer, such as **http://PrintServer12/Printers/IPrinter52/.printer** .
  - To specify a printer to use by TCP/IP address or host name, select Add A Printer Using A TCP/IP Address Or Hostname, and then click Next. Choose a device type, and then enter the host name or IP address for the printer, such as **192.168.1.90** . If you select the Autodetect or TCP/IP Device options, the wizard will set the port name to the same value, but you can also choose a different value. Click Next.
5. On the Type A Printer Name page, the printer name is set for you. You can accept the default name or enter a new name. Click Next to install the printer, and then click Finish. The user can now print to the network printer by selecting the printer in an application. The Device And Printers folder on the user's computer shows the new network printer as well.

To create a connection to the printer on a Windows 10 system, follow these steps:


1. Open the Settings app. On the home page, click Devices.




2. On the Printers & Scanners page, click Add A Printer Or Scanner. Settings will then attempt to detect the printer automatically.

If the wizard finds the printer with which you want to work, click it in the list provided, follow the prompts, and skip the rest of the steps in this procedure.

If the wizard doesn't find the printer, click The Printer That I Want Isn't Listed, and then complete the rest of this procedure.



3. In the Add Printer Wizard, click Add Bluetooth, Wireless Or Network Discoverable Printer.



4. In the list of available printers, select the printer you want to use, and then click Next.
5. If prompted, install the printer driver on your computer. Complete the additional steps in the wizard, and then click Finish. You can confirm the printer is working by printing a test page.

If you have trouble connecting to the printer, try the following as part of troubleshooting:

- Be sure that a firewall isn't blocking connectivity to the printer. You might need to open a firewall port to enable access between the computer and the printer.
- Be sure the printer is turned on and connected to the same network as the computer. If your network consists of multiple subnets connected together, try to connect the





printer to the same network subnet. You can determine the subnet by looking at the computer's IP address.

- Be sure the printer is configured to broadcast its presence on the network. Most network printers automatically do this.
- Be sure the printer has an IP address and proper network settings. With DHCP, network routers assign IP addresses automatically as printers connect to the network.





## Managing Installed Printers

Once you've installed network printers, you'll need to perform a number of management tasks. This section looks at routine tasks, including those for deployment printer connections, configuring restrictions, moving printers, monitoring print services and resolving spooling problems.






# Deploying Printer Connections

Connecting to printers is fairly easy, but you can make the process even easier by deploying printer connections through Group Policy. You can deploy printer connections to computers or users via the Group Policy objects (GPOs) that Windows applies. Deploy the connections to groups of users when you want users to be able to access the printers from any computer they log on to. Deploy the connections to groups of computers when you want all users of the computers to access the printers. For percomputer connections, Windows adds or removes printer connections when the computer starts. For per-user connections, Windows adds or removes printer connections when the user logs on.

To deploy printer connections to computers, you must follow these steps:

1. In Print Management, expand the Print Servers node and the node for the server with which you want to work.
2. Select the server's Printers node. In the main pane, right-click the printer you want to deploy, and then click Deploy With Group Policy. This displays the Deploy With Group Policy dialog box.



3. Click Browse. In the Browse For Group Policy Object dialog box, select the GPO to use, and then click OK.
4. Do one or both of the following:
  - To deploy to the printer connection on a per-user basis, under Deploy This Printer Connection To The Following, select The Users That This GPO Applies To.
  - To deploy to the printer connection on a per-computer basis, under Deploy This





Printer Connection To The Following, select The Computers That This GPO Applies To.

5. Click Add to create a print connection entry.
6. Repeat steps 3–5 to deploy the printer connection to other GPOs.
7. Click OK to save the changes to the GPO. In the confirmation dialog box, ensure that all operations were completed successfully. If an error occurred, click Details to get more information about the error. The most common errors involve editing permissions for the GPO with which you are working. If the account you are using doesn't have appropriate permissions, you need to use an account with additional privileges. Click OK.





# Configuring Point and Print Restrictions

In Group Policy, the Point And Print Restrictions setting controls security warnings and elevation prompts when users point and print and when drivers for printer connections need to be configured. This setting is found in the Administrative Templates for Computer Configuration under the Printers node.

The Point And Print Restrictions setting is used as follows:

- When enabled, clients can point and print to any server. You can configure Clients to show or hide warning and elevation prompts when users point and print and when a driver for an existing printer connection needs to be updated.
- When not configured, clients can point and print to any server in the forest. Clients also will not show a warning and elevation prompt when users point and print or when a driver for an existing printer connection needs to be updated.
- When disabled, clients can point and print to any server. Clients also will not show a warning and elevation prompt when users point and print or when a driver for an existing printer connection needs to be updated.

By default, Windows allows a user who is not a member of the local Administrators group to install only trustworthy printer drivers, such as those provided by Windows or in digitally signed printer driver packages. When you enable the Point And Print Restrictions setting, you also make it possible for users who are not members of the local Administrators group to install printer connections deployed in Group Policy that include additional or updated printer drivers that are not in the form of digitally signed printer driver packages. If you do not enable this setting, users might need to provide the credentials of a user account that belongs to the local Administrators group.

You can enable and configure the Point And Print Restrictions setting in Group Policy by following these steps:

1. In the Group Policy Management Console, right-click the GPO for the site, domain, or organizational unit with which you want to work, and then click Edit. This opens the policy editor for the GPO.
2. In the Group Policy Management Editor, expand the Administrative Templates for Computer Configuration, and then select the Printers node.
3. In the main pane, double-click Point And Print Restrictions.
4. In the Point And Print Restrictions dialog box, select Enabled.



Point and Print Restrictions

Previous Setting Next Setting

☐ Not Configured ☒ Enabled ☐ Disabled

Comment:

Supported on: At least Windows Vista

Options:

☒ Users can only point and print to these servers:  
Enter fully qualified server names separated by semicolons  
PrintServer43.tvpress.local;PrintServer18.i

☒ Users can only point and print to machines in their forest

Security Prompts:

When installing drivers for a new connection:  
Show warning and elevation prompt

When updating drivers for an existing connection:  
Show warning and elevation prompt

Help:

This policy setting controls the client Point and Print behavior, including the security prompts for Windows Vista computers. The policy setting applies only to non-Print Administrator clients, and only to computers that are members of a domain.

If you enable this policy setting:

- Windows XP and later clients will only download print driver components from a list of explicitly named servers. If a compatible print driver is available on the client, a printer connection will be made. If a compatible print driver is not available on the client, no connection will be made.
- You can configure Windows Vista clients so that security warnings and elevated command prompts do not appear when users Point and Print, or when printer connection drivers need to be updated.

If you do not configure this policy setting:

- Windows Vista client computers can point and print to any server.
- Windows Vista computers will show a warning and an elevated command prompt when users create a printer


5. When you enable point and print restrictions, you can configure policy so that users can point and print only to a named list of servers. To enforce this restriction, select the related check box and enter a list of fully qualified server names separated by semicolons. To remove this restriction, clear the related check box.
6. When you enable point and print restrictions, you can configure policy so that users can point and print only to servers in their forest. To enforce this restriction, select the related check box. To remove this restriction, clear the related check box.
7. When you install drivers for a new connection, clients can show or not show a warning or elevation prompt. Use the related selection list to choose the option you want to use.
8. When you update drivers for an existing connection, clients can show or not show a warning or elevation prompt. Use the related selection list to choose the option you want to use.
9. Click OK to apply the configuration.

## Moving Printers to a New Print Server


You can use the Printer Migration Wizard to move print queues, printer drivers, printer processors, and printer ports from one print server to another. This is an efficient way to consolidate multiple print servers or replace an older print server.

When you move printers, the server on which the printers are currently located is the source server, and the server to which you want to move the printers is the destination server. With this in mind, you can move printers to a new print server by following these steps:

1. In Print Management, right-click the source server, and then click Export Printers To A File. This starts the Printer Migration Wizard.
2. On the initial page, note the printer-related objects that will be exported, and then click Next.
3. On the Select The File Location page, click Browse. In the dialog box provided, select a save location for the printer migration file. After you enter a name for the file, click Save.



4. Printer migration files are saved with the .printerExport extension. Click Next to save the printer settings to this file.
5. If an error occurred during processing, click Open Event Viewer to review the events generated during the export process and then use the event entries to determine what happened and possible actions to take to resolve the problem. When you have finished, exit the Event Viewer.
6. On the Exporting page, click Finish to exit the Printer Migration Wizard.
7. In Print Management, right-click the destination server, and then click Import Printers From A File. This launches the Printer Migration Wizard.
8. On the Select The File Location page, click Browse. In the dialog box provided select the printer migration file you created in steps 3 and 4, and then click Open.




9. Click Next. Note the objects that will be imported, and then click Next. On the Select Import Options page, choose one of the following options in the Import Mode list:

- **Keep Existing Printers** When you choose this option and existing printer queues have the same names as those you are importing, the wizard creates copies to ensure that the original printer queues and the imported printer queues are both available.
- **Overwrite Existing Printers** When you choose this option and existing printer queues have the same names as those you are importing, the wizard overwrites the existing printer queues with the information from the printer queues you are importing.

10. On the Select Import Options page, choose one of the following options in the List In The Directory list:

- **List Printers That Were Previously Listed** Choose this option to specify that only printers that were previously listed are listed in Active Directory.
- **List All Printers** Choose this option to specify that all printers are listed in Active Directory.
- **Don't List Any Printers** Choose this option to specify that no printers are listed in Active Directory.







11. Click Next to begin the import process. If an error occurred during processing, click Open Event Viewer to review the events generated during the export process and then use the event entries to determine what happened and possible actions to take to resolve the problem. When you have finished, exit the Event Viewer.
12. On the Importing page, click Finish to exit the Printer Migration Wizard.





## Monitoring Printers and Printer Queues

Printer filters display only the printers, printer queues, and printer drivers that meet specific criteria. Through automated notification, you can use printer filters to automate monitoring of printers.

In Print Management, you can view existing filters by expanding the Custom Filters node. If you expand the Custom Filters node and then select a filter, the main pane shows all printers or print drivers that match the filter criteria. Print Management includes the following default printer filters:

- **All Printers** Lists all printers associated with print servers that have been added to the console
- **All Drivers** Lists all printer drivers associated with print servers that have been added to the console
- **Printers Not Ready** Lists all printers that are not in a Ready state, such as those with errors
- **Printers With Jobs** Lists all printers associated with print servers that have active or pending print jobs

You can create a new custom filter by follow these steps:

1. In Print Management, right-click the Custom Filters node, and then click Add New Printer Filter to start the New Printer Filter Wizard.
2. On the Printer Filter Name And Description page, enter a filter name and description. If you'd like the number of matching items to be displayed after the filter name, select the Display The Total Number Of Printers check box. Click Next.
3. On the Define A Filter page, define the filter by specifying the text box, condition, and value to match in the first row. If you want to further narrow the possible matches, define additional criteria in the second, third, and subsequent rows. Click Next when you are ready to continue.

**NOTE** When you use filters for monitoring and notification, you use the Queue Status text box the most. This text box enables you to receive a notification when a printer has a specific status. You can match the following status values: Attention Required, Busy, Deleting, Door Open, Error, Initializing, IO Active, Manual Feed Required, No Toner Ink, Not Available, Offline, Out of Memory, Out of Paper, Output Bin Full, Page Punt,



Paper Jam, Paper Problem, Paused, Printing, Processing, Ready, Toner Ink Low, Waiting, and Warming Up

**TIP** When you are matching conditions in a filter, you can match an exact condition that does exist or one that does not exist. For example, if you want to be notified only of conditions that need attention, you can look for Queue Status conditions that are not exactly the following: Deleting, Initializing, IO Active, Printing, Processing, Waiting, Warming Up, and Ready.

4. On the Set Notifications (Optional) page, you can specify whether to send an email message, run a script, or take both actions when the specified criteria are met. Click Finish to complete the configuration.

You can modify an existing custom filter by following these steps:

5. In Print Management, expand the Custom Filters node. Select and then right-click the filter with which you want to work. On the shortcut menu, click Properties.
  6. In the filter's Properties dialog box, use the options provided to manage the filter settings. This dialog box has the following three tabs:
- **General** Shows the name and description of the printer filter. Enter a new name and description as necessary.
  - **Filter Criteria** Shows the filter criteria. Enter new filter criteria as necessary.
  - **Notification** Shows the email and script options. Enter new email and script options as necessary.



## Solving Spooling Problems

Windows Server 2016 uses the Print Spooler service to control the spooling of print jobs. If this service isn't running, print jobs can't be spooled. Use the Services console to check the status of the Print Spooler. Follow these steps to check and restart the Print Spooler service:

1. On the Tools menu in Server Manager, click Computer Management.
2. If you want to connect to a remote computer, right-click the Computer Management entry in the console tree, and then click Connect To Another Computer. You can now choose the system whose services you want to manage.
3. Expand the Services And Applications node, and then click Services.
4. Select the Print Spooler service. The Status should be Started. If it isn't, rightclick Print Spooler, and then click Start. The Startup Type should be Automatic. If it isn't, double-click Print Spooler and set Startup Type to Automatic.

**TIP** Spoolers can become corrupted. Symptoms of a corrupted spooler include a frozen printer or one that doesn't send jobs to the print device. Sometimes the print device might print pages of garbled data. In most of these cases, you can resolve the problem by stopping and starting the Print Spooler service. Other spooling problems might be related to permissions. See "Setting Printer Access Permissions" in Chapter 31.





# Chapter 31. Configuring and Maintaining Print Services

Few things are as frustrating for users as when printing doesn't work properly. Users want to be able to click Print in an application and have the document print on a printer. They don't care how it all works; they just want it to work when they need it to. For this reason and others, you need to work behind the scenes to ensure print services operate smoothly. In this chapter, you'll learn how to work with print servers, printers, print queues and print jobs. Print server properties are managed separately from printer properties as are print queues and print jobs.





# Configuring Print Server Properties

In Windows Server 2016, you control global settings for print servers by using print server properties. To access these properties, complete the following steps:

1. **Open Print Management.** In Server Manager, click Tools and then select Print Management.
2. **Open the Print Server Properties dialog box.** In Print Management, right-click the server entry for the print server with which you want to work, and then click Properties.

**NOTE** If the print server isn't listed, you can add it in the Add/Remove Servers dialog box. To open the dialog box, right-click Print Servers, and then click Add/Remove Servers.

The sections that follow examine some of the print server properties that you can configure.



# Locating the Spool Folder and Enabling Printing


The Spool folder holds a copy of all documents in the printer spool. By default, this folder is located at %SystemRoot%\System32\Spool\Printers. On the NTFS file system, all users who access a printer must have the following special permissions on this directory:

- Read Attributes
- Read Extended Attributes
- Create Files / Write Data
- Create Folders / Append Data


If users don't have these permissions, they won't be able to print documents.

If you're experiencing problems, check the permission on this directory by following these steps:


1. Access the Print Server Properties dialog box.
2. Click the Advanced tab. The location of the Spool folder is shown in the Spool Folder text box. Note this location.




3. In File Explorer, open the properties dialog box for the Spool folder. For example, if the Spool folder is C:\Windows\System32\Spool\PRINTERS. Access C:\Windows\System32\Spool, right-click PRINTERS and then select Properties.



4. In the Properties dialog box, select the Security tab, and then click Advanced.
5. In the Advanced Security Settings dialog box, click the user or group whose permissions you want to verify and then click Edit.



6. In the Permission Entry For dialog box, select Show Advanced Permissions. You'll then be able to verify that the permissions are set appropriately.



## Managing High-Volume Printing

Printers used in corporate environments can print hundreds of documents daily or more. This high volume puts a heavy burden on print servers and can cause printing delays, document corruption, and other problems. To alleviate some of this burden, you should do the following:

- Use network-attached printers rather than directly attached printers. Network-attached printers use fewer system resources (namely CPU time) than do other printers.
- Use Branch Office Direct Printing to render print jobs on client computers, and then send the print jobs directly to a printer. Otherwise, print jobs are sent to the print server where they are rendered, and then sent to the printer.
- Dedicate the print server to handle print services only, such as with a virtual server. If the print server is handling other network duties, it might not be very responsive to print requests and management. To increase responsiveness, you can move other network duties to other servers.






- Move the Spool folder to a drive dedicated to printing. By default, the Spool folder is on the same file system as the operating system. To further improve disk input/output (I/O), use a drive that has a separate controller.



# Enabling Print Job Error Notification

Print servers can beep to notify users when a remote document fails to print. By default, this feature is turned off because it can be annoying. If you want to activate or remove notification, access the Advanced tab of the Print Server Properties dialog box, and then select or clear the check box labeled Beep On Errors Of Remote Documents.





# Managing Printer Drivers


In a Windows Server 2016 domain you should configure and update printer drivers only on your print servers. You don't need to update printer drivers on Windows clients. Instead, you configure the network printer to provide the drivers to client systems as necessary.




# Updating a Printer Driver

You can update a printer's driver by following these steps:


1. In Print Management, double-click the Print Servers node, and then double-click the entry for the print server itself.
2. Select Printers. Right-click the printer with which you want to work and then select Properties.



3. In the printer's Properties dialog box, click the Advanced tab.
4. In the Driver list, you can select the driver from a list of currently installed drivers. Use the Driver list to select a driver from a list of known drivers.



5. If the driver you need isn't listed or if you obtained a new driver, click New Driver to start the Add Printer Driver Wizard. Click Next.
6. Select the manufacturer and model of the print driver to install. Optionally, click Windows Update to see whether there are more drivers for the previously selected printer and processor. Here, Windows connects to the Windows Update Web site to update the list of printers to show additional models. This feature is part of the automatic driver provisioning feature. It can take several minutes to retrieve the updated list. You should then be able to select your printer manufacturer and model.



7. If the device manufacturer and model you're using aren't displayed in the list, choose Have Disk to install a new driver from a file or disc. In the Install From Disk dialog box, enter the folder path to the printer driver file or click Browse to find the printer driver file by using the Locate File dialog box. Click OK.
8. Click Next, and then click Finish.

## Configuring Drivers for Clients

After you install a printer or change drivers, you might want to select the operating systems that should download the driver from the print server. By enabling clients to download the printer driver, you provide a single location for installing driver updates. This way, instead of having to install a new driver on all the client systems, you install the driver on the print server and enable clients to download the new driver.

You can enable clients to download the new driver by following these steps:

1. In Print Management, double-click the Print Servers node, and then double-click the entry for the print server itself.
2. Select Printers. Right-click the icon of the printer you want to configure, and then click Printer Properties.
3. Click the Sharing tab, and then click Additional Drivers.
4. In the Additional Drivers dialog box, select operating systems that can download the printer driver. As necessary, insert the Windows Server 2016 installation media, printer driver discs, or both for the selected operating



systems. The Windows Server 2016 installation media has drivers for most Windows operating systems.





# Configuring Printer Properties


In the sections that follow, I explain how to set commonly used printer properties. To open a printer's properties dialog box on a printer server, follow these steps:

1. In Print Management, expand the Print Servers node and the node for the server with which you want to work.
2. Select the server's Printers node. In the main pane, right-click the printer with which you want to work, and then click Properties. You can now set printer properties.



## Adding Comments and Locations

To make it easier to determine which printer to use and when, you can add comments and location information for printers. Comments provide general information about the printer, such as the type of print device and who is responsible for it. Location describes the actual site of the print device. After set, applications can display these text boxes. For example, Microsoft Office Word displays this information in the Comment and Where text boxes when you select Print from the File menu.




You can add comments and location information to a printer by using the text boxes on the General tab of the printer's Properties dialog box. Enter your comments in the Comment text box. Enter the printer location in the Location text box.

## Listing Printers in Active Directory

Listing printers in Active Directory makes it easier for users to locate and install printers. You can list a printer in Active Directory by doing one of the following:

- Right-click the printer's name, and then click **List In Directory**.
- Open the printer's Properties dialog box, and then click the **Sharing** tab. Select the **List In Directory** check box, and then click **OK**.






# Setting Separator Pages and Print Device Modes

Separator pages have two uses on Windows Server 2016 systems:

- You can use them at the beginning of a print job to make it easier for users to find a document on a busy print device.
- You can use them to change the print device mode, such as whether the print device uses PostScript or Printer Control Language (PCL).

To set a separator page for a print device, follow these steps:

1. On the Advanced tab of the printer's Properties dialog box, click Separator Page. This opens the Separator Page dialog box.
2. When you work with a local server, click the Browse button in the Separator Page dialog box to open the %SystemRoot%\Windows\System32 folder for browsing. In this case, you can browse and select the separator page to use. In contrast, when you work with a remote server, the Browse button is typically not available, so you must enter the exact file name for the separator page.



Typically, you'll want to use one of these available separator pages:

- **Pcl.sep** Switches the print device to PCL mode and prints a separator page before each document
- **Pscript.sep** Sets the print device to PostScript mode but doesn't print a separator page
- **Sysprint.sep** Sets the print device to PostScript mode and prints a separator page before each document

NOTE Sysprintj.sep is an alternate version of Sysprint.sep. If fonts for Japanese are available and you want to use them, you can use Sysprintj.sep.




To stop using the separator page, open the Separator Page dialog box and remove the file name.



## Changing the Printer Port

You can change the port used by a print device at any time by using the Properties dialog box for the printer you're configuring. Open the Properties dialog box, and then click the Ports tab.




You can now add a port for printing by selecting its check box or remove a port by clearing its check box. To add a new port type, click Add Port. In the Printer Ports dialog box, select the port type, and then click New Port. Enter a valid port name, and then click OK. To remove a port permanently, select it, and then click Delete Port.

## Starting and Stopping Printer Sharing

You set printer sharing in the Properties dialog box of the printer you're configuring. Right-click the icon of the printer you want to configure, and then click Manage Sharing. This opens the printer's Properties dialog box with the Sharing tab selected. You can use this tab to change the name of a network printer and to start sharing or stop sharing a printer. Printer sharing tasks that you can perform include the following:

- **Sharing a local printer (thus making it a network printer)** To share a printer, select Share This Printer, and then specify a name for the shared resource in the Share Name text box. Click OK when you have finished.
- **Changing the shared name of a printer** To change the shared name, just enter a new name in the Share Name text box, and then click OK.
- **Stopping the sharing of a printer** To quit sharing a printer, clear the Share This Printer check box, and then click OK.



## Setting Printer Access Permissions

Network printers are shared resources. As such, you can set access permissions for them. You set access permissions in the Properties dialog box of the printer you're configuring. Open the printer's Properties dialog box, and then click the Security tab. Permissions that you can grant or deny for printers are:

- **Print** Allows you to print documents; pause, restart, resume, and cancel your own documents; and connect to printers.
- **Manage Documents** Allows you to perform all Print permission tasks. Also allows you to control settings for print jobs and to pause, restart, and delete print jobs.
- **Manage This Printer** Allows you to perform all Manage Documents tasks. Also allows you to share printers, change printer properties, change printer permissions, and delete printers.

The default permissions are used for any new network printer you create. These settings are as follows:

- Members of the Administrators, Print Operators, and Server Operators groups have full control over printers by default. This makes it possible for you to administer a printer and its print jobs.
- The creator or owner of the document can manage his or her own document. This enables the person who printed a document to change its settings and to delete it.
- Everyone can print to the printer. This makes the printer accessible to all users on the network.

As with other permission sets, the basic permissions for printers are created by combining special permissions into logical groups. By using Advanced permission settings, you can assign these special permissions individually if necessary.

HP Color LaserJet CP3505 Properties

General

Sharing

Ports

Advanced

Color Management

Security

Device Settings

Group or user names:

Everyone

ALL APPLICATION PACKAGES

CREATOR OWNER

William R. Stanek (vrstanek@imaginedlands.local)

Administrators (CORPSERVER54\Administrators)

Add...

Remove

Permissions for Everyone

|                     | Allow                               | Deny                     |
|---------------------|-------------------------------------|--------------------------|
| Print               | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| Manage this printer | <input type="checkbox"/>            | <input type="checkbox"/> |
| Manage documents    | <input type="checkbox"/>            | <input type="checkbox"/> |
| Special permissions | <input type="checkbox"/>            | <input type="checkbox"/> |


For special permissions or advanced settings, click Advanced.

Advanced

# Setting Document Defaults


Document default settings are used only when you print from applications that are not based on Windows, such as when you print from the command prompt. You can set document defaults by following these steps:

1. Open the printer's Properties dialog box, and then click the General tab.
2. Click Preferences.
3. Use the text boxes on the tabs provided to configure the default settings for orientation, paper source, media, color and more.




# Scheduling and Prioritizing Print Jobs

You can use the Properties dialog box for the printer you're configuring to set default settings for print job priority and scheduling. Open the Properties dialog box, and then click the Advanced tab. You can now set the default schedule and priority settings.



## Scheduling Printer Availability and Priority

Printers are always available or available only during the hours specified. You set printer availability on the Advanced tab. Select Always Available to make the printer available at all times. Select Available From to set specific hours of operation.








Use the Priority box on the Advanced tab to set the default priority for print jobs that are printed using the selected printer. Print jobs always print in order of priority. Jobs with higher priority print before jobs with lower priority. As a single physical printer could have multiple print queues, each with different priorities, you can use prioritization to ensure important documents are always printed before other documents.



# Configuring Print Spooling

For print devices attached to the network, you usually want the printer to spool files rather than print files directly. Print spooling makes it possible to use a printer to manage print jobs.

A screenshot of the Windows Print Spooler configuration window, specifically the 'Advanced' tab. The window has a light gray background and a dark gray border. At the top, there are three radio button options: 'Spool print documents so program finishes printing faster' (selected), 'Start printing after last page is spooled', and 'Start printing immediately'. Below these is a horizontal line, followed by a radio button option 'Print directly to the printer'. Underneath the line are four checkbox options: 'Hold mismatched documents' (unchecked), 'Print spooled documents first' (checked), 'Keep printed documents' (unchecked), and 'Enable advanced printing features' (checked). At the bottom, there are three buttons: 'Printing Defaults...', 'Print Processor...', and 'Separator Page...'.

To enable spooling, use one of the following options on the Advanced tab:

- **Spool Print Documents So Program Finishes Printing Faster** Select this option to spool print jobs.
- **Start Printing After Last Page Is Spooled** Select this option if you want the entire document to be spooled before printing begins. This option ensures that the entire document makes it into the print queue before printing. If for some reason printing is canceled or not completed, the job won't be printed.
- **Start Printing Immediately** Select this option if you want printing to begin immediately when the print device isn't already in use. This option is preferable when you want print jobs to be completed more quickly or when you want to ensure that the application returns control to users as soon as possible.

You can disable spooling by selecting the Print Directly To The Printer option. The following additional check boxes let you configure other spooling options:

- **Hold Mismatched Documents** If you select this option, the spooler holds print jobs that don't match the setup for the print device. Selecting this option is a good idea if you frequently have to change printer form or tray assignments.
- **Print Spooled Documents First** If you select this option, jobs that have completed spooling will print before jobs in the process of spooling, regardless of whether the spooling jobs have higher priority.




- **Keep Printed Documents** Typically, documents are deleted from the queue after they're printed. To keep a copy of documents in the printer, select this option. Use this option if you're printing files that can't easily be re-created. In this way, you can reprint the document without having to re-create it. For details, see "Pausing, Resuming, and Restarting Document Printing" later in this chapter.
- **Enable Advanced Printing Features** When you enable this option, you can use advanced printing options (if available), such as Page Order and Pages Per Sheet. If you note compatibility problems when using advanced options, you should disable the advanced printing features by clearing this check box.




# Managing Print Jobs

You manage print jobs and printers by using the print queue window. If the printer is configured on your system, you can access the print queue window by using one of the following techniques:


- If the printer is configured on your system, you can manage the printer via Settings. On the Settings home page, select Devices. Under Printers & Scanners, click the icon of the printer with which you want to work and then select Open Queue.



- In Print Management, double-click the Print Servers node, and then double-click the entry for the print server itself. Select Printers. Right-click the printer with which you want to work, and then click Open Printer Queue.



- In Print Management, right-click the Printers node, and then click Show Extended View. With the printer selected in the upper pane, the upper and lower panes in the main window provide functionality similar to that of the print management window.





## Viewing Printer Queues and Print Jobs

You can now manage print jobs and printers by using the print management window. The print management window shows information about documents in the printers. This information tells you the following:

- **Document Name** The document file name, which can include the name of the application that printed it.
- **Status** The status of the print job, which can include the document's status and the printer's status. Document status entries you'll find include Printing, Spooling, Paused, Deleting, and Restarting. Document status can be preceded by the printer status, such as Printer Off-Line.
- **Owner** The document's owner.
- **Pages** The number of pages in the document.
- **Size** The document size in kilobytes or megabytes.
- **Submitted** The time and date the print job was submitted.
- **Port** The port used for printing, such as LPT1, COM3, File, or IP address (if applicable).





## **Pausing the Printer and Resuming Printing**

Sometimes you need to pause a printer. You do this by using the print management window and selecting the Pause Printing option from the Printer menu. (A check mark indicates that the option is selected.) When you pause printing, the printer completes the current job and then puts all other jobs on hold.

To resume printing, select the Pause Printing option a second time. This should remove the check mark next to the option.





## Emptying the Print Queue

You can use the print management window to empty the print queue and delete all its contents. To do this, choose Cancel All Documents from the Printer menu.





## Pausing, Resuming, and Restarting Document Printing

You set the status of individual documents by using the Document menu in the print management window. To change the status of a document, right-click the document, and then use one of the following options on the shortcut menu to change the status of the print job.

- **Pause** Puts the document on hold and lets other documents print
- **Resume** Tells the printer to resume printing the document from where it left off
- **Restart** Tells the printer to start printing the document again from the beginning







## Removing Documents and Canceling Print Jobs

To remove a document from the printer or cancel a print job, select the document in the print management window. Right-click the document and click Cancel or press Delete.

**NOTE** When you cancel a print job that's currently printing, the print device might continue to print part or all of the document. This occurs because most print devices cache documents in an internal buffer and the print device might continue to print the contents of this cache.





## Checking Document Properties

Document properties can tell you many things about documents that are in the printer, such as the page source, orientation, and size. You can check the properties of a document in the printer by doing either of the following:

- Right-click the document in the print management window, and then click
- Properties.

Double-click the document name in the print management window.





## Setting Document Priority

Scheduling priority determines when documents print. Documents with higher priority print before documents with lower priority. You can set the priority of individual documents in the printer by following these steps:

1. Right-click the document in the print management window, and then click Properties.
2. On the General tab, use the Priority slider to change the document's priority. The lowest priority is 1 and the highest is 99.





## Scheduling Document Printing

In a busy printing environment, you might need to schedule the printing of documents in the printer. For example, you might want large print jobs of low priority to print at night. To set the printing schedule, follow these steps:

1. Right-click the document in the print management window, and then click Properties.
2. On the General tab, select the Only From option, and then specify a time interval. The time interval you set determines when the job can print. For example, you can specify that the job can print only between midnight and 5:00 A.M.





# Auditing Print Jobs

Windows Server 2016 lets you audit common printer tasks by following these steps:

1. Open the printer's Properties dialog box, and then click the Security tab.
2. Click Advanced to open the Advanced Security Settings dialog box.

NOTE Actions aren't audited by default. You must first enable auditing by establishing a group policy to audit the printer.

3. On the Auditing tab, add the names of users or groups you want to audit by using the Add button and remove names of users or groups by using the Remove button.
4. Select the events you want to audit by selecting the check boxes under the Successful and Failed headings, as appropriate.
5. Click OK.







## Chapter 32. Implementing DHCP

You can use Dynamic Host Configuration Protocol (DHCP) to simplify administration of Active Directory domains, and in this chapter you'll learn how to do that. You use DHCP to dynamically assign TCP/IP configuration information to network clients. This not only saves time during system configuration, but also provides a centralized mechanism for updating the configuration. To enable DHCP on the network, you need to install and configure a DHCP server. This server is responsible for assigning the necessary network information.





# Understanding DHCP

DHCP gives you centralized control over IP addressing and more. After DHCP is installed, you rely on the DHCP server to supply the basic information necessary for TCP/IP networking. This basic information can include the following: IP address, subnet mask, and default gateway; primary and secondary Domain Name System (DNS) servers; primary and secondary Windows Internet Name Service (WINS) servers; and the DNS domain name. DHCP servers can assign a dynamic IP version 4 (IPv4) address, a dynamic IP version 6 (IPv6) address, or both addresses to any of the network interface cards (NICs) on a computer.







# Using Dynamic IPv4 Addressing and Configuration

A computer that uses dynamic IPv4 addressing and configuration is called a DHCPv4 client . When you start a DHCPv4 client, a 32-bit IPv4 address can be retrieved from a pool of IPv4 addresses defined for the network's DHCP server. The address is assigned to the client for a specified time period known as a lease . When the lease is approximately 50 percent expired, the client tries to renew it. If the client can't renew the lease at that time, it tries again before the lease expires. If this attempt fails, the client tries to contact an alternate DHCP server. IPv4 addresses that aren't renewed are returned to the address pool. If the client is able to contact the DHCP server but the current IP address can't be reassigned, the DHCP server assigns a new IPv4 address to the client.

The availability of a DHCP server doesn't affect startup or logon (in most cases). DHCPv4 clients can start and users can log on to the local computer even if a DHCP server isn't available. During startup, the DHCPv4 client looks for a DHCP server. If a DHCP server is available, the client gets its configuration information from the server. If a DHCP server isn't available and the client's previous lease is still valid, the client pings the default gateway listed in the lease. A successful ping tells the client that it's probably on the same network it was on when it was issued the lease, and the client continues to use the lease as described previously. A failed ping tells the client that it might be on a different network. In this case, the client uses IPv4 autoconfiguration. The client also uses IPv4 autoconfiguration if a DHCP server isn't available and the previous lease has expired.

IPv4 autoconfiguration works like this:

1. The client computer selects an IP address from the Microsoft-reserved class B subnet 169.254.0.0 and uses the subnet mask 255.255.0.0. Before using the IPv4 address, the client performs an Address Resolution Protocol (ARP) test to be sure that no other client is using this IPv4 address.
2. If the IPv4 address is in use, the client repeats step 1, testing up to 10 IPv4 addresses before reporting failure. When a client is disconnected from the network, the ARP test always succeeds. As a result, the client uses the first IPv4 address it selects.
3. If the IPv4 address is available, the client configures the NIC with this address. The client then attempts to contact a DHCP server, sending out a broadcast every five minutes to the network. When the client successfully contacts a server, the client obtains a lease and reconfigures the network interface.





As part of your planning, you need to consider how many DHCP servers should be installed on the network. Typically, you'll want to install at least two DHCP servers on each physical network segment. Windows Server 2016 includes DHCP failover for IPv4. DHCP failover enables high availability of DHCP services by synchronizing IPv4 address leases between two DHCP servers in one of two modes:

- **Load Balance** When you load balance the servers, you specify the percentage of the load each server should handle. Typically, you use a 50/50 approach to make each server equally share the load. You also could use other approaches, such as 60/40 to make one server carry 60 percent of the load and the other 40 percent of the load.
- **Hot Standby** With hot standby, one of the servers acts as the primary server and handles the DHCP services. The other acts as a standby server in case the primary fails or runs out of addresses to lease. A specific percentage of available IP addresses are reserved for the hot standby—5 percent by default.

The configuration of DHCP failover is simple and straightforward, and it does not require clustering or any advanced configuration. To configure DHCP failover, all you need to do is complete the following steps:

4. Install and configure two DHCP servers. The servers should be on the same physical network.
5. Create a DHCPv4 scope on one of the servers. Scopes are pools of IPv4 or IPv6 addresses you can assign to clients through leases.
6. When you establish the other server as a failover partner for the DHCPv4 scope, the scope is replicated to the partner.

## Using Dynamic IPv6 Addressing and Configuration

Both IPv4 and IPv6 are enabled by default when networking hardware is detected during installation. As discussed in Chapter 18, “Implementing TCP/IP networking,” IPv4 is the primary version of IP used on most networks, and IPv6 is the next generation version of IP. IPv6 uses 128-bit addresses. In a standard configuration, the first 64 bits represent the network ID, and the last 64 bits represent the network interface on the client computer.

You can use DHCP to configure IPv6 addressing in two key ways:

- **DHCPv6 stateful mode** In DHCPv6 stateful mode, a client acquires its IPv6 address in addition to its network configuration parameters through DHCPv6.
- **DHCPv6 stateless mode** In DHCPv6 stateless mode, a client uses autoconfiguration to acquire its IP address and acquires its network configuration parameters through DHCPv6.



A computer that uses dynamic IPv6 addressing, configuration, or both mechanisms is called a DHCPv6 client . As with DHCPv4, the components of the DHCPv6 infrastructure consist of DHCPv6 clients that request configuration, DHCPv6 servers that provide configuration, and DHCPv6 relay agents that convey messages between clients and servers when clients are on subnets that do not have a DHCPv6 server.

Unlike in DHCPv4, you must also configure your IPv6 routers to support DHCPv6. A DHCPv6 client performs autoconfiguration based on the following flags in the Router Advertisement message sent by a neighboring router:

- **Managed Address Configuration flag, which is also known as the M flag .** When set to 1, this flag instructs the client to use a configuration protocol to obtain stateful addresses.
- **Other Stateful Configuration flag, which is also known as the O flag .** When set to 1, this flag instructs the client to use a configuration protocol to obtain other configuration settings.

Windows includes a DHCPv6 client. The DHCPv6 client attempts DHCPv6-based configuration depending on the values of the M and O flags in the Router Advertisement messages it receives. If there is more than one advertising router for a given subnet, the additional router or routers should be configured to advertise the same stateless address prefixes and the same values for the M and O flags. All current Windows desktop and server operating systems include IPv6 clients and, therefore, accept the values of the M and O flags in router advertisements they receive.

You can configure an IPv6 router to set the M flag to 1 in router advertisements by entering the following command at an administrator command prompt, where InterfaceName is the actual name of the interface:

```
netsh interface ipv6 set interface InterfaceName managedaddress=enabled
```

Similarly, you can set the O flag to 1 in router advertisements by entering the following command at an administrator command prompt:

```
netsh interface ipv6 set interface InterfaceName otherstateful=enabled
```

If the interface name contains spaces, enclose the related value in quotation marks, as shown in the following example:

```
netsh interface ipv6 set interface "Wired Ethernet Connection 2"  
managedaddress=enabled
```

Keep the following in mind when you are working with the M and O flags:

- If the M and O flags are both set to 0, the network is considered not to have DHCPv6 infrastructure. Clients use router advertisements for non-link-local addresses and manual configuration to configure other settings.
- If the M and O flags are both set to 1, DHCPv6 is used for both IP addressing and other configuration settings. This combination is known as DHCPv6 stateful mode , in which DHCPv6 assigns stateful addresses to IPv6 clients.
- If the M flag is set to 0 and the O flag is set to 1, DHCPv6 is used only to assign other configuration settings. Neighboring routers are configured to advertise non-link-local address prefixes from which IPv6 clients derive stateless addresses. This combination is known as DHCPv6 stateless mode .
- If the M flag is set to 1 and the O flag is set to 0, DHCPv6 is used for IP address configuration but not for other settings. Because IPv6 clients typically need to be configured with other settings, such as the IPv6 addresses of DNS servers, this combination typically is not used.

Windows obtains dynamic IPv6 addresses by using a process similar to obtaining dynamic IPv4 addresses. Typically, IPv6 autoconfiguration for DHCPv6 clients in stateful mode works like this:

1. The client computer selects a link-local unicast IPv6 address. Before using the IPv6 address, the client performs an ARP test to make sure that no other client is using this IPv6 address.
2. If the IPv6 address is in use, the client repeats step 1. Keep in mind that when a client is disconnected from the network, the ARP test always succeeds. As a result, the client uses the first IPv6 address it selects.
3. If the IPv6 address is available, the client configures the NIC with this address. The client then attempts to contact a DHCP server, sending out a broadcast every five minutes to the network. When the client successfully contacts a server, the client obtains a lease and reconfigures the network interface.

This is not how IPv6 autoconfiguration works for DHCPv6 clients in stateless mode. In stateless mode, DHCPv6 clients configure both link-local addresses and additional nonlink-local addresses by exchanging Router Solicitation and Router Advertisement messages with neighboring routers.

Like DHCPv4, DHCPv6 uses User Datagram Protocol (UDP) messages. DHCPv6 clients listen for DHCP messages on UDP port 546. DHCPv6 servers and relay agents listen for DHCPv6 messages on UDP port 547. The structure for DHCPv6 messages is much simpler than for DHCPv4, which had its origins in Bootstrap Protocol (BOOTP) to support diskless workstations.



DHCPv6 messages start with a 1-byte Msg-Type field that indicates the type of DHCPv6 message. This is followed by a 3-byte Transaction-ID field determined by a client and used to group together the messages of a DHCPv6 message exchange. Following the Transaction-ID field, DHCPv6 options are used to indicate client and server identification, addresses, and other configuration settings.

Three fields are associated with each DHCPv6 option. A 2-byte Option-Code field indicates a specific option. A 2-byte Option-Len field indicates the length of the OptionData field in bytes. The Option-Data field contains the data for the option.

Messages exchanged between relay agents and servers use a different message structure to transfer additional information. A 1-byte Hop-Count field indicates the number of relay agents that have received the message. A receiving relay agent can discard the message if the message exceeds a configured maximum hop count. A 15-byte LinkAddress field contains a non-link-local address that is assigned to an interface connected to the subnet on which the client is located. Based on the Link-Address field, the server can determine the correct address scope from which to assign an address. A 15-byte Peer-Address field contains the IPv6 address of the client that originally sent the message or the previous relay agent that relayed the message. Following the PeerAddress field are DHCPv6 options. A key option is the Relay Message option. This option provides an encapsulation of the messages being exchanged between the client and the server.

IPv6 does not have broadcast addresses. The use of the limited broadcast address for some DHCPv4 messages has been replaced with the use of the All\_DHCP\_Relay\_Agents\_and\_Servers address of FF02::1:2 for DHCPv6. A DHCPv6 client attempting to discover the location of the DHCPv6 server on the network sends a Solicit message from its link-local address to FF02::1:2. If there is a DHCPv6 server on the client's subnet, it receives the Solicit message and sends an appropriate reply. If the client and server are on different subnets, a DHCPv6 relay agent on the client's subnet that receives the Solicit message forwards it to a DHCPv6 server.

## Checking IP Address Assignment

You can use Ipconfig to check the currently assigned IP address and other configuration information. To obtain information for all network adapters on the computer, enter the command **ipconfig /all** at the command prompt. If the IP address has been assigned automatically, you'll notice an entry for Autoconfiguration IP Address. In the following example, the autoconfiguration IPv4 address is 169.254.98.59:



## Windows IP Configuration

Host Name .....: FRODO  
Primary DNS Suffix .....: tvpress.local  
Node Type .....: Hybrid  
IP Routing Enabled.....: No  
WINS Proxy Enabled.....: No  
DNS Suffix Search List.....: tvpress.local  
Ethernet adapter Ethernet:  
    Connection-specific DNS Suffix...:  
    Description .....: Intel Pro/1000 Network Connection  
    Physical Address.....: 23-15-C6-F8-FD-67  
    DHCP Enabled.....: Yes  
    Autoconfiguration Enabled...: Yes  
    Autoconfiguration IP Address:  
169.254.98.59 Subnet Mask .....:  
255.255.0.0 Default Gateway .....:  
DNS Servers .....:

## Understanding Scopes

Scopes are pools of IPv4 or IPv6 addresses you can assign to clients through leases. DHCP also provides a way to permanently assign a lease on an address. To do this, you need to create a reservation by specifying the IPv4 address to reserve and the media access control (MAC) address of the computer that will hold the IPv4 address. The reservation thereafter ensures that the client computer with the specified MAC address always gets the designated IPv4 address. With IPv6, you can specify that a lease is temporary or nontemporary. A nontemporary lease is similar to a reservation.

You create scopes to specify IP address ranges that are available for DHCP clients. For example, you could assign the IP address range 192.168.12.2 to 192.168.12.250 to a scope called Enterprise Primary. Scopes can use public or private IPv4 addresses on the following networks:

- **Class A networks** IP addresses from 1.0.0.0 to 126.255.255.255
- **Class B networks** IP addresses from 128.0.0.0 to 191.255.255.255
- **Class C networks** IP addresses from 192.0.0.0 to 223.255.255.255
- **Class D networks** IP addresses from 224.0.0.0 to 239.255.255.255

NOTE The IP address 127.0.0.1 is used for local loopback (and so are any other IP addresses in the 127.x.y.z address range).

Scopes can also use link-local unicast, global unicast, and multicast IPv6 addresses. Link-local unicast addresses begin with FE80. Multicast IPv6 addresses begin with FF00. Global (site-local) unicast addresses include all other addresses except :: (unspecified) and ::1 (loopback) addresses.





A single DHCP server can manage multiple scopes. With IPv4 addresses, four types of scopes are available:

- **Normal scopes** Used to assign IPv4 address pools for class A, B, and C networks.
- **Multicast scopes** Used to assign IP address pools for IPv4 class D networks. Computers use multicast IP addresses as secondary IP addresses in addition to a standard IP address.
- **Superscopes** Containers for other scopes that are used to simplify management of multiple scopes and also support DHCP clients on a single physical network where multiple logical IP networks are used.
- **Failover scopes** Scopes split between two DHCP servers to increase fault tolerance, provide redundancy, and enable load balancing.

With IPv6, only normal scopes are available. Although you can create scopes on multiple network segments, you'll usually want these segments to be in the same network class, such as all IP addresses that are class C.

**TIP** Don't forget that you must configure DHCPv4 and DHCPv6 relays to relay DHCPv4 and DHCPv6 broadcast requests between network segments. You can configure relay agents with the Routing and Remote Access Service (RRAS) and the DHCP Relay Agent Service. You can also configure some routers as relay agents. These services can be installed as part of the Remote Access role. On a server with no other policy and access role services configured, you can install the Remote Access role by using the Add Roles And Features Wizard.





# Installing a DHCP Server

Dynamic IP addressing is available only if a DHCP server is installed on the network. By using the Add Roles And Features Wizard, you install the DHCP server as a role service, configure its initial settings, and authorize the server in Active Directory Domain Services (AD DS). Only authorized DHCP servers can provide dynamic IP addresses to clients.







# Installing DHCP Components

On a server running Windows Server 2016, follow these steps to enable the server to function as a DHCP server:

1. DHCP servers should be assigned a static IPv4 and IPv6 address on each subnet to which they are connected and will service. Be sure that the server has static IPv4 and IPv6 addresses.
2. In Server Manager, click Manage, and then click Add Roles And Features, or select Add Roles And Features in the Quick Start pane. This starts the Add Roles And Features Wizard. If the wizard displays the Before You Begin page, read the Welcome text, and then click Next.
3. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Click Next.
4. On the Server Selection page, you can choose to install roles and features on running servers or virtual hard disks. Either select a server from the server pool or select a server from the server pool on which to mount a virtual hard disk (VHD). If you are adding roles and features to a VHD, click Browse, and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, click Next.

NOTE Only servers running Windows Server 2016 and that have been added for management in Server Manager are listed.

5. On the Select Roles page, select DHCP Server. If additional features are required to install a role, you'll get an additional dialog box. Click Add Features to close the dialog box and add the required features to the server installation. When you are ready to continue, click Next three times.




6. If the server on which you want to install the DHCP Server role doesn't have all the required binary source files, the server gets the files via Windows Update by default or from a location specified in Group Policy.




**NOTE** You also can specify an alternate path for the required source files. To do this, click the Specify An Alternate Source Path link, enter that alternate path in the box provided, and then click OK. For network shares, enter the UNC path to the share, such as \\Server23\\WinServer2016\\. For mounted Windows images, enter the WIM path prefixed with WIM: and including the index of the image to use, such as WIM:\\Server23\\WS16\\install.wim:4.

7. After you review the installation options and save them as necessary, click **Install** to begin the installation process. The **Installation Progress** page tracks the progress of the installation. If you close the wizard, click the **Notifications** icon in **Server Manager**, and then click the link provided to reopen the wizard.
8. When Setup finishes installing the DHCP Server role, the **Installation Progress** page will be updated to reflect this. Review the installation details to ensure that all phases of the installation were completed successfully.
9. Additional configuration is required for the DHCP server. Click the **Complete DHCP Configuration** link to start the DHCP Post-Install Configuration Wizard. If you closed the wizard, click **Notifications** on **Server Manager**'s toolbar and then click the **Complete DHCP Configuration** link in the **Post-Deployment Configuration** task panel.



10. The **Description** page states that the DHCP Administrators and DHCP Users groups will be created in the domain for delegation of DHCP Server administration. Additionally, if the DHCP server is joined to a domain, the server will be authorized in Active Directory. Click **Next**.






11. On the Authorization page, do one of the following to specify the credentials to use to authorize the DHCP server in Active Directory:

- Your current user name is shown in the User Name text box. If you have administrator privileges in the domain that the DHCP server is a member of and you want to use your current credentials, click Commit to attempt to authorize the server by using these credentials.
- If you want to use alternate credentials or if you are unable to authorize the server by using your current credentials, select Use Alternate Credentials, and then click Specify. In the Windows Security dialog box, enter the user name and password for the authorized account, and then click OK. Click Commit to attempt to authorize the server by using these credentials.
- If you want to authorize the DHCP server later, select Skip AD Authorization, and then click Commit. Keep in mind that in domains, only authorized DHCP servers can provide dynamic IP addresses to clients.

The screenshot shows the 'Authorization' page of the DHCP Server Configuration Wizard. On the left, there is a navigation pane with 'Description', 'Authorization' (selected), and 'Summary'. The main area has the title 'Specify the credentials to be used to authorize this DHCP server in AD DS.' Below this, there are three radio button options: 'Use the following user's credentials' (selected), 'Use alternate credentials', and 'Skip AD authorization'. Under the selected option, there is a 'User Name' text box containing 'IMAGINEDLANDS\wvrstane'. Under the 'Use alternate credentials' option, there is a 'UserName' text box and a 'Specify...' button.

12. When the wizard finishes the post-install configuration, review the installation details to ensure that tasks were completed successfully, and then click Close.
13. Next, you need to restart the DHCP Server service on the DHCP server so that the DHCP Administrators and DHCP Users groups can be used. To do this, click DHCP in the left pane of Server Manager. Next, in the main pane, on the Servers panel, select the DHCP server. Finally, on the Services panel, right-click the entry for the DHCP Server, and then click Restart Service.





14. To complete the installation, you need to do the following:


- If the server has multiple network cards, review the server bindings and specify the connections that the DHCP server supports for servicing clients. See “Configuring Server Bindings” later in this chapter.
- Configure server options to assign common configuration settings for DHCPv4 and DHCPv6 clients, including 003 Router, 006 DNS Servers, 015 DNS Domain Name, and 044 WINS/NBNS Servers. See “Setting Scope Options” in Chapter 33. •

Create and activate any DHCP scopes that the server will use, as discussed in “Creating and Managing Scopes” in Chapter 33.

# Starting and Using the DHCP Console

After you install a DHCP server, you use the DHCP console to configure and manage dynamic IP addressing. In Server Manager, click Tools, and then click DHCP to open the DHCP console. Alternatively, enter **Dhcpmgmt.msc** in the Search box or at a prompt.

The main window for the DHCP console is divided into two panes. The left pane lists the DHCP servers in the domain according to their fully qualified domain name (FQDN). You can expand a server listing to show subnodes for IPv4 and IPv6. If you expand the IP nodes, you'll find the scopes and options defined for the related IP version. The right pane shows the expanded view of the current selection.



Icons on the various nodes show the current status of the nodes. For server and IP nodes, you might find the following icons:

- A server icon with a green circle with a check mark indicates that the DHCP service is running and the server is active.
- A server icon with red circle with an X through it indicates that the console can't connect to the server. The DHCP service has been stopped or the server is inaccessible.
- A red down arrow indicates that the DHCP server hasn't been authorized. A blue warning icon indicates that the server's state has changed or a warning has been issued.

For scopes, you might find the following icons:


- A red down arrow indicates that the scope hasn't been activated.
- A blue warning icon indicates that the scope's state has changed or a warning has been issued.



# Connecting to Remote DHCP Servers

When you start the DHCP console, you are connected directly to a local DHCP server, but you won't find entries for remote DHCP servers. You can connect to remote servers by following these steps:

- 1. Right-click DHCP in the console tree, and then click Add Server.
- 2. Select This Server, and then enter the IP address or computer name of the DHCP server you want to manage.
- 3. Click OK. An entry for the DHCP server is added to the console tree.



**TIP** When you work with remote servers, you might find that you can't select certain options. A simple refresh of the server information might resolve this problem. To refresh the server information right-click the server node, and then select Refresh.

# Starting and Stopping a DHCP Server

You manage DHCP servers through the DHCP Server service. As with any other service, you can start, stop, pause, and resume the DHCP Server service in the Services node of Computer Management or from the command line. You can also manage the DHCP Server service in the DHCP console. Right-click the server you want to manage in the DHCP console, point to All Tasks, and then click Start, Stop, Pause, Resume, or Restart, as appropriate.

**NOTE** You also can use Server Manager to start and stop a DHCP server. Click DHCP in the left pane of Server Manager. Next, in the main pane, on the Servers panel, select the DHCP server. Finally, on the Services panel, right-click the entry for the DHCP Server, and then click Start Service, Stop Service, Pause Service, Resume Service, or Restart Service, as appropriate.





## Authorizing a DHCP Server in Active Directory

Before you can use a DHCP server in the domain, you must authorize it in Active Directory. By authorizing the server, you specify that the server is authorized to provide dynamic IP addressing in the domain. Windows Server 2016 requires authorization to prevent unauthorized DHCP servers from serving domain clients. This in turn ensures that network operations can run smoothly.

Only Enterprise Admins can authorize DHCP servers. In the DHCP console, you authorize a DHCP server by right-clicking the server entry in the tree view, and then selecting Authorize. To remove the authorization, right-click the server, and then select Unauthorize.

At an administrator PowerShell prompt, you can use `Add-DhcpServerInDC` to authorize DHCP servers. Use the `-DnsName` parameter to specify the name of the server to authorize or the `-IpAddress` to specify the IP address of the server to authorize as shown in the following examples:

```
Add-DhcpServerInDC -DnsName CorpSvr03.imaginedlands.com Add-  
DhcpServerInDC -IpAddress 192.168.1.1
```

Use `Remove-DhcpServerInDC` to remove the authorization. The basic syntax is the same.

## Configuring DHCP Servers

After you install a new DHCP server, you need to configure and optimize the server for the network environment. A separate set of options are provided for IPv4 and IPv6.






# Configuring Server Bindings

A server with multiple NICs has multiple local area network connections and can provide DHCP services on any of these network connections. However, you might not want DHCP to be served over all available connections. For example, if the server has both a 100-megabits per second (Mbps) connection and a 1-gigabit per second (Gbps) connection, you might want all DHCP traffic to go over the 1-Gbps connection.

To bind DHCP to a specific network connection, follow these steps:

1. In the DHCP console, right-click the server with which you want to work, and then click Add/Remove Bindings.
2. Select the IPv4 or IPv6 tab as appropriate for the type of binding with which you want to work.



3. The Bindings dialog box displays a list of available network connections for the DHCP server.
- If you want the DHCP Server service to use a connection to service clients, select the check box for the connection.
  - If you don't want the service to use a connection, clear the related check box.
  - If there are no network connections listed for the protocol with which you are working, ensure that the server has a static address for that protocol.
4. Click OK when you have finished.

## Updating DHCP Statistics


The DHCP console provides statistics concerning IPv4 and IPv6 address availability and usage. In the DHCP console, you can view these statistics by expanding the node for the server with which you want to work, right-clicking IPv4 or IPv6 as appropriate for the type of address with which you want to work, and then clicking Display Statistics.

By default, these statistics are updated only when you start the DHCP console or when you select the server and then click the Refresh button on the toolbar. If you monitor



DHCP routinely, you might want these statistics to be updated automatically, which you can do by following these steps:

1. In the DHCP console, expand the node for the server with which you want to work, right-click IPv4 or IPv6 as appropriate for the type of address with which you want to work, and then click Properties.
2. On the General tab, select Automatically Update Statistics Every and enter an update interval in hours and minutes. Click OK.



## Auditing and Troubleshooting DHCP

Windows Server 2016 is configured to audit DHCP processes by default. Auditing tracks DHCP processes and requests in log files.

You can use audit logs to help you troubleshoot problems with a DHCP server. Just as you enable and configure logging separately for IPv4 and IPv6, the two protocols use different log files. %SystemRoot%\System32\DHCP is the default location for DHCP logs. In this directory, you'll find a different log file for each day of the week. The IPv4 log file for Monday is named DhcpSrvLog-Mon.log, the log file for Tuesday is named DhcpSrvLog-Tue.log, and so on. The IPv6 log file for Monday is named DhcpV6SrvLog-Mon.log, the log file for Tuesday is named DhcpV6SrvLog-Tue.log, and so on.

When you start the DHCP server or a new day arrives, a header message is written to the log file. This header provides a summary of DHCP events and their meanings. Stopping and starting the DHCP Server service doesn't clear a log file. Log data is kept for a week. For example, the DHCP Server service clears and starts over Monday's log the following Monday. You don't have to monitor space usage by the DHCP Server service. The service is configured to monitor itself and restricts disk space usage by default.

You can enable or disable DHCP auditing by following these steps:



1. In the DHCP console, expand the node for the server with which you want to work, right-click IPv4 or IPv6 as appropriate for the type of address with which you want to work, and then click Properties.
2. On the General tab, select or clear the Enable DHCP Audit Logging check box, and then click OK.

By default, DHCP logs are stored in %SystemRoot%\System32\DHCP. You can change the location of DHCP logs by following these steps:

1. In the DHCP console, expand the node for the server with which you want to work, right-click IPv4 or IPv6 as appropriate for the type of address with which you want to work, and then click Properties.
2. Click the Advanced tab. Audit Log File Path shows the current folder location for log files. Enter a new folder location, or click Browse to select a new location.
3. Click OK. Windows Server 2016 now needs to restart the DHCP Server service. When prompted to restart the service, click Yes. The service will be stopped and then started again.

The DHCP server has a self-monitoring system that checks disk space usage. By default, the maximum size of all DHCP server logs is 70 megabytes (MB), with each individual log being limited to one-seventh of this space. If the server reaches the 70-MB limit or an individual log grows beyond the allocated space, logging of DHCP activity stops until log files are cleared or space is otherwise made available. Typically, this happens at the beginning of a new day when the server clears the previous week's log file for that day.

Registry keys that control log usage and other DHCP settings are located under HKEY\_LOCAL\_MACHINE in:

SYSTEM\CurrentControlSet\Services\DHCP\Parameters

The following keys control the logging:

- **DhcpLogFilesMaxSize** Sets the maximum file size for all logs. The default is 70 MB.
- **DhcpLogDiskSpaceCleanupInterval** Determines how often DHCP checks disk space usage and cleans up as necessary. The default interval is 60 minutes.
- **DhcpLogMinSpaceOnDisk** Sets the free space threshold for writing to the log. If the disk has less free space than the value specified, logging is temporarily disabled. The default value is 20 MB.





DhcpLogMinSpaceOnDisk is considered an optional key and is not created automatically. You need to create this key as necessary and set appropriate values for your network.

## Integrating DHCP and DNS

DNS is used to resolve computer names in Active Directory domains and on the Internet. Thanks to the DNS dynamic update protocol, you don't need to manually register DHCP clients in DNS. The protocol allows the client or the DHCP server to register the forward-lookup and reverse-lookup records in DNS as necessary. When configured by using the default setup for DHCP, current DHCP clients automatically update their own DNS records after receiving an IP address lease. You can modify this behavior globally for each DHCP server or on a per-scope basis.

Name protection is an additional feature in Windows Server 2016. With name protection, the DHCP server registers records on behalf of the client only if no other client with this DNS information is already registered. You can configure name protection for IPv4 and IPv6 at the network adapter level or at the scope level. Name protection settings configured at the scope level take precedence over the setting at the IPv4 or IPv6 level.

Name protection is designed to prevent name squatting. Name squatting occurs when a computer not based on the Windows operating system registers a name in DNS that is already registered to a computer running a Windows operating system. By enabling name protection, you can prevent name squatting by computers not based on the Windows operating system. Although name squatting generally does not present a problem when you use Active Directory to reserve a name for a single user or computer, it usually is a good idea to enable name protection on all Windows networks.


Name protection is based on the Dynamic Host Configuration Identifier (DHCID) and support for the DHCID RR (resource record) in DNS. The DHCID is a resource record stored in DNS that maps names to prevent duplicate registration. DHCP uses the DHCID resource record to store an identifier for a computer along with related information for the name, such as the A and AAAA records of the computer. The DHCP server can request a DHCID record match and then refuse the registration of a computer with a different address attempting to register a name with an existing DHCID record.

You can view and change the global DNS integration settings by following these steps:

1. In the DHCP console, expand the node for the server with which you want to work, right-click IPv4 or IPv6, and then click Properties.
2. Click the DNS tab. The default DNS integration settings for DHCP are as shown in the dialog box that follows. Because these settings are configured by default, you usually don't need to modify the configuration. However, if you only want host (A) records to be dynamically updated instead of both host (A) and pointer (PTR) records, select the Disable Dynamic Updates For DNS PTR Records check box.

**IMPORTANT** The default configuration, which registers and maintains both A and PTR records, assumes that you've configured reverse lookup zones for your organization. If you haven't, attempts to register and update PTR records will fail. You can prevent repeated failed attempts to register and update PTR records by disabling dynamic updates for PTR records. If you disable this option in the IPv4 properties, you are disabling the option for all IPv4 scopes. Alternatively, you can use scope properties to disable the option on a per scope basis.

3. Optionally, you can enable or disable the name protection feature. With name protection, the DHCP server registers records on behalf of the client only if no other client with this DNS information is already registered. To enable or disable name protection, click Configure. In the Name Protection dialog box, select or clear Enable Name Protection, and then click OK.



You can view and change the per-scope DNS integration settings by following these steps:



1. In the DHCP console, expand the node for the server with which you want to work, and then expand IPv4 or IPv6.
2. Right-click the scope with which you want to work, and then click Properties.
3. Click the DNS tab. The options available are the same as those shown in the previous figure. Because these settings are configured by default, you usually don't need to modify the configuration.
4. Optionally, you can enable or disable the name-protection feature. Click Configure. In the Name Protection dialog box, select or clear Enable Name Protection, and then click OK.



## Avoiding IP Address Conflicts

IPv4 address conflicts are a common cause of problems with DHCP. No two computers on the network can have the same unicast IP address. If a computer is assigned the same unicast IPv4 address as another, one or both of the computers might become disconnected from the network. More specifically, the computer already using the IPv4 address can continue using the address and any other computer that tries to use that IPv4 address is blocked from using it.

To better detect and avoid potential conflicts, you can enable IPv4 address conflict detection by following these steps:

1. In the DHCP console, expand the node for the server with which you want to work, right-click IPv4, and then click Properties.
2. On the Advanced tab, set Conflict Detection Attempts to a value other than 0. The value you enter determines the number of times the DHCP server checks an IP address before leasing it to a client. The DHCP server checks IP addresses by sending a ping request over the network.

**REAL WORLD** A unicast IPv4 address is a standard IP address for class A, B, and C networks. When a DHCP client requests a lease, a DHCP server checks its pool of available addresses and assigns the client a lease on an available IPv4 address. By default, the server checks only the list of current leases to determine whether an address is available. It doesn't actually query the network to determine whether an address is in use. Unfortunately, in a busy network environment, an administrator might have assigned this IPv4 address to another computer or an offline computer might have been brought online with a lease that it believes hasn't expired, even though the DHCP server believes the lease has expired. Either way, you have an address conflict that will cause problems on the network. To reduce these types of conflicts, set the conflict detection to a value greater than 0.

## Saving and Restoring the DHCP Configuration

After you configure all the necessary DHCP settings, you might want to save the DHCP configuration so that you can restore it on the DHCP server. To save the configuration, enter the following command at the command prompt:

```
netsh dump DHCP > dhcpconfig.dmp
```

In this example, dhcpconfig.dmp is the name of the configuration script you want to create. By default, the script is created in the current working directory. Alternatively, you can specify the full file path in which to save the script. After you create this script,





you can restore the configuration by executing the script. If you saved the script in the current working directory, you can enter the following command at the command prompt:

```
netsh exec dhcpconfig.dmp
```

If you saved the script in another directory, you can specify the full path to the script, such as:

```
netsh exec d:\dhcp\scripts\dhcpconfig.dmp
```

**TIP** You can also use this technique to set up another DHCP server with the same configuration. Just copy the configuration script to a folder on the destination computer, and then execute it.

You can save or restore the DHCP configuration by using the DHCP console as well. To save the configuration, right-click the DHCP server entry, click Backup, use the dialog box provided to select the folder for the backup, and then click OK. To restore the configuration, right-click the DHCP server entry, click Restore, use the dialog box provided to select the backup folder, and then click OK. When prompted to confirm, click Yes.

At an administrator PowerShell prompt, you use Export-DhcpServer to save the configuration settings. The basic syntax is:

```
Export-DhcpServer -ComputerName ServerID -File SavePath
```

Here, ServerID is the DNS name or IP address of the DHCP server and SavePath is the path and name of the file in which you want to store the configuration settings. If you omit the name of the server to work with, the local server is used. If you don't specify a save path along with the file name, the configuration file is created in the current working directory. In the following example, you store the configuration settings in the d:\dhcp\scripts directory with the name dhcpconfig.dmp:

```
Export-DhcpServer -File d:\dhcp\scripts\dhcpconfig.dmp
```

You can restore the configuration using Import-DhcpServer. The basic syntax is:

```
Import-DhcpServer -ComputerName ServerID -BackupPath CurrentConfigSavePath -  
File SavePath
```

Here, SavePath is the path and name of the file in which you stored the configuration settings and CurrentConfigSavePath specifies the path where the current configuration should be saved prior to importing and overwriting existing settings. In the following example, you back up the settings to d:\dhcp\backup\origconfig.dmp and then apply the saved configuration from d:\dhcp\scripts\dhcpconfig.dmp:



```
Import-DhcpServer-BackupPath d:\dhcp\backup\origconfig.dmp  
-File d:\dhcp\scripts\dhcpconfig.dmp
```







## Chapter 33. Managing and Maintaining DHCP

After you install a DHCP server, you need to configure the scopes that the DHCP server will use, manage addressing, and perform basic maintenance tasks. Scopes are pools of IP addresses you can lease to clients. For each scope you define, you can add reservations, exclusions and filters to control which IP addresses are available and used. You can create superscopes, normal scopes, multicast scopes, and failover scopes with IPv4 addresses, but you can create only normal scopes with IPv6 addresses.





# Creating and Managing Scopes


Scopes provide a pool of IP addresses for DHCP clients. A normal scope is a scope with class A, B, or C network addresses. A multicast scope is a scope with class D network addresses. Although you create normal scopes and multicast scopes differently, you manage them in much the same way. The key differences are that multicast scopes can't use reservations, and you can't set additional options for WINS, DNS, routing, and so forth.



# Creating Normal Scopes for IPv4 Addresses

You can create a normal scope for IPv4 addresses by following these steps:


1. In the DHCP console, expand the node for the server with which you want to work, and then right-click IPv4. If you want to add the new scope to a superscope automatically, right-click the superscope instead.
2. On the shortcut menu, click New Scope to start the New Scope Wizard. Click Next.
3. Enter a name and description for the scope, and then click Next.



4. The Start IP Address and End IP Address boxes define the valid IP address range for the scope. On the IP Address Range page, enter a start address and an end address in these boxes.

**NOTE** Generally, the scope doesn't include the x.x.x.0 and x.x.x.255 addresses, which are usually reserved for network addresses and broadcast messages, respectively. Accordingly, you would use a range such as 192.168.10.1 to 192.168.10.254 rather than 192.168.10.0 to 192.168.10.255.

5. When you enter an IP address range, the bit length and subnet mask are filled in for you automatically. Unless you use subnets, you should use the default values.



The image shows a screenshot of the 'New Scope Wizard' window, specifically the 'IP Address Range' step. The window has a title bar 'New Scope Wizard' and a subtitle 'IP Address Range'. Below the subtitle is a description: 'You define the scope address range by identifying a set of consecutive IP addresses.' There is a folder icon in the top right corner. The main content area is divided into two sections. The first section is titled 'Configuration settings for DHCP Server' and contains the instruction 'Enter the range of addresses that the scope distributes.' Below this are two input fields: 'Start IP address:' with the value '10 . 0 . 10 . 1' and 'End IP address:' with the value '10 . 0 . 10 . 254'. The second section is titled 'Configuration settings that propagate to DHCP Client' and contains two input fields: 'Length:' with the value '8' and 'Subnet mask:' with the value '255 . 0 . 0 . 0'.

New Scope Wizard

**IP Address Range**  
You define the scope address range by identifying a set of consecutive IP addresses.

Configuration settings for DHCP Server

Enter the range of addresses that the scope distributes.

Start IP address: 10 . 0 . 10 . 1

End IP address: 10 . 0 . 10 . 254

Configuration settings that propagate to DHCP Client

Length: 8

Subnet mask: 255 . 0 . 0 . 0

6. Click Next. If the IP address range you entered is too large for a single scope, you're given the opportunity to create a superscope that contains separate scopes for each network and, in this case, select the Yes option button to continue, and then move on to step 8. If you make a mistake, click Back, and then modify the IP address range you entered.
7. Use the Start IP Address and End IP Address boxes on the Add Exclusions And Delay page to define IP address ranges that are to be excluded from the scope. You can exclude multiple address ranges as follows:
  - To define an exclusion range, enter a start address and an end address in the Exclusion Range's Start IP Address and End IP Address boxes, and then click Add. To exclude a single IP address, use that address as both the start IP address and the end IP address.
  - To track which address ranges are excluded, use the Excluded Address Range list.
  - To delete an exclusion range, select the range in the Excluded Address Range list, and then click Remove.

**Add Exclusions and Delay**

Exclusions are addresses or a range of addresses that are not distributed by the server. A delay is the time duration by which the server will delay the transmission of a DHCP OFFER message.

Type the IP address range that you want to exclude. If you want to exclude a single address, type an address in Start IP address only.

Start IP address:  End IP address:

Excluded address range:

|                            |                                       |
|----------------------------|---------------------------------------|
| 10.0.10.90 to 10.0.10.99   | <input type="button" value="Remove"/> |
| 10.0.10.170 to 10.0.10.189 |                                       |

Subnet delay in milli second:

8. Click Next. Specify the duration of leases for the scope by using the Day(s), Hour(s), and Minutes boxes. The default duration is eight days. Click Next.

Set the duration for scope leases when distributed by this server.

Limited to:

Days:  Hours:  Minutes:

**NOTE** A lease duration that's set too long can reduce the effectiveness of DHCP and might eventually cause you to run out of available IP addresses, especially on networks with mobile users or other types of computers that aren't fixed members of the network. A good lease duration for most networks is from one to three days.

9. You have the opportunity to configure common DHCP options for DNS, WINS, gateways, and more. If you want to set these options now, select Yes, I Want To Configure These Options Now; otherwise, select No, I Will Configure These Options Later and skip steps 10–15.

Do you want to configure the DHCP options for this scope now?

☒ Yes, I want to configure these options now

☐ No, I will configure these options later

10. Click Next. The first option you can configure is the default gateway. In the IP Address box, enter the IP address of the primary default gateway, and then click Add. Repeat this process for other default gateways.
11. The first gateway listed is the one clients try to use first. If the gateway isn't available, clients try to use the next gateway, and so on. Use the Up and Down buttons to change the order of the gateways, as necessary.

12. Click Next. Configure default DNS settings for DHCP clients. Enter the name of the parent domain to use for DNS resolution of computer names that aren't fully qualified.



**Router (Default Gateway)**  
You can specify the routers, or default gateways, to be distributed by this scope.


To add an IP address for a router used by clients, enter the address below.

IP address:

|           |
|-----------|
| 10.0.10.1 |
|-----------|

Buttons: Add, Remove, Up, Down

13. In the IP Address box, enter the IP address of the primary DNS server, and then click Add. Repeat this process to specify additional DNS servers. Again, the order of the entries determines which IP address is used first. Change the order as necessary by using the Up and Down buttons. Click Next.



**Domain Name and DNS Servers**  
The Domain Name System (DNS) maps and translates domain names used by clients on your network.

You can specify the parent domain you want the client computers on your network to use for DNS name resolution.

Parent domain:


To configure scope clients to use DNS servers on your network, enter the IP addresses for those servers.

| Server name:         | IP address: |
|----------------------|-------------|
| <input type="text"/> | 10.0.0.171  |
| <input type="text"/> | 10.0.10.175 |

Buttons: Resolve, Add, Remove, Up, Down

**TIP** If you know the name of a server instead of its IP address, enter the name in the Server Name box, and then click Resolve. The IP address is then entered in the IP Address box, if possible. Add the server by clicking Add. Using this approach helps to ensure you've entered the correct information.

14. Configure default WINS settings for the DHCP clients. The techniques you use are the same as those previously described. Click Next.



**WINS Servers**  
Computers running Windows can use WINS servers to convert NetBIOS computer names to IP addresses.


Entering server IP addresses here enables Windows clients to query WINS before they use broadcasts to register and resolve NetBIOS names.

Server name:  IP address:

10.0.10.179  
10.0.10.185

To change this behavior for Windows DHCP clients modify option 046, WINS/NBT Node Type, in Scope Options.

15. If you want to activate the scope, select Yes, I Want To Activate This Scope Now, and then click Next. Otherwise, select No, I Will Activate This Scope Later and then click Next.



**Activate Scope**  
Clients can obtain address leases only if a scope is activated.

Do you want to activate this scope now?

☒ Yes, I want to activate this scope now

☐ No, I will activate this scope later

16. Click Finish to complete the process.

## Creating Normal Scopes for IPv6 Addresses


You create normal scopes for IPv6 addresses by using the New Scope Wizard. When you are configuring DHCP for IPv6 addresses, you must enter the network ID and a preference value. Typically, the first 64 bits of an IPv6 address identify the network, and a 64-bit value is what the New Scope Wizard expects you to enter. The preference value sets the priority of the scope relative to other scopes. The scope with the lowest preference value will be used first. The scope with the second-lowest preference will be used second, and so on.

You can create a normal scope for IPv6 addresses by following these steps:

1. In the DHCP console, expand the node for the server with which you want to work and then expand the IPv6 node.
2. Right-click IPv6. On the shortcut menu, click New Scope to start the New Scope Wizard. Click Next.




3. Enter a name and description for the scope, and then click Next.



The screenshot shows the 'Scope Name' configuration page. At the top, the title 'Scope Name' is followed by the instruction: 'You have to provide an identifying scope name. You also have the option of providing a description.' Below this, a text box prompts the user: 'Type a name and description for this scope. This information helps you quickly identify how the scope is to be used on your network.' There are two input fields: 'Name:' with the value 'Primary IPv6 Scope' and 'Description:' which is empty. A folder icon is visible in the top right corner.

4. On the Scope Prefix page, enter the 64-bit network prefix, and then set a preference value. Click Next.



The screenshot shows the 'Scope Prefix' configuration page. At the top, the title 'Scope Prefix' is followed by the instruction: 'You have to provide a prefix to create the scope. You also have the option of providing a preference value for the scope.' Below this, a text box prompts the user: 'Enter the IPv6 Prefix for the addresses that the scope distributes and the preference value for the scope.' There are two input fields: 'Prefix' with the value 'FD00::' and a '/64' suffix, and 'Preference' with a value of '0' and a small icon to its right. A folder icon is visible in the top right corner.

5. Use the Start IPv6 Address and End IPv6 Address boxes on the Add Exclusions page to define IPv6 address ranges that are to be excluded from the scope. You can exclude multiple address ranges as follows:

- To define an exclusion range, enter a start address and an end address in the Exclusion Range's Start IPv6 Address and End IPv6 Address boxes, and then click Add. To exclude a single IPv6 address, use that address as the start IPv6 address, and then click Add.
- To track which address ranges are excluded, use the Excluded Address Range list.
- To delete an exclusion range, select the range in the Excluded Address Range list, and then click Remove.



**Add Exclusions**  
Exclusions are addresses or a range of addresses that are not distributed by the server.

Type the IPv6 address range that you want to exclude for the given scope. If you want to exclude a single address, type an identifier in Start IPv6 Address only.

Start IPv6 Address: fd00::

End IPv6 Address: fd00::

Excluded address range:

|                        |
|------------------------|
| fd00::100 to fd00::199 |
|------------------------|

- Click Next. Dynamic IPv6 addresses can be temporary or nontemporary. A nontemporary address is similar to a reservation. On the Scope Lease page, specify the duration of leases for nontemporary addresses by using the Days, Hours, and Minutes boxes under Preferred Life Time and Valid Life Time. The preferred lifetime is the preferred amount of time the lease should be valid. The valid lifetime is the maximum amount of time the lease is valid. Click Next.

Non Temporary Address(IANA)

Preferred Life Time

|                                |                                |                                |
|--------------------------------|--------------------------------|--------------------------------|
| Days:                          | Hours:                         | Minutes:                       |
| <input type="text" value="3"/> | <input type="text" value="0"/> | <input type="text" value="0"/> |

Valid Life Time

|                                 |                                |                                |
|---------------------------------|--------------------------------|--------------------------------|
| Days:                           | Hours:                         | Minutes:                       |
| <input type="text" value="12"/> | <input type="text" value="0"/> | <input type="text" value="0"/> |


**NOTE** A lease lifetime that's set too long can reduce the effectiveness of DHCP. A good lease duration for nontemporary leases is from 8 to 30 days.

- If you want to activate the scope, select Yes under Activate Scope Now, and then click Finish. Otherwise, select No under Activate Scope Now, and then click Finish.

# Creating Multicast Scopes

To create a multicast scope, follow these steps:

1. In the DHCP console, expand the node for the server with which you want to work. Select and then right-click IPv4. If you want to add the new scope to a superscope, select and then right-click the superscope instead.
2. On the shortcut menu, click New Multicast Scope to start the New Multicast Scope Wizard. Click Next.
3. Enter a name and description for the scope, and then click Next.




**Multicast Scope Name**  
You have to provide an identifying multicast scope name. You also have the option of providing a description.

Type a name and description for this multicast scope. This information helps you quickly identify how the multicast scope is to be used on your network.

Name:

Description:

4. The Start IP Address and End IP Address boxes define the valid IP address range for the scope. Enter a start address and an end address in these boxes. You must define multicast scopes by using Class D IP addresses. This means the valid IP address range is 224.0.0.0 to 239.255.255.255.



**IP Address Range**  
You set the range of IP addresses that define this multicast scope.

The valid IP address range is 224.0.0.0 to 239.255.255.255.

Start IP address:

End IP address:


Time to Live (TTL) is the number of routers that multicast traffic passes through on your network.

TTL:


5. Messages sent by computers using multicast IP addresses have a specific time-to-live (TTL) value. The TTL value specifies the maximum number of routers the message can go through. The default value is 32, which is sufficient on most networks. If you have a large network, you might need to increase this value to reflect the actual number of routers that might be used.
6. Click Next. If you make a mistake, click Back, and then modify the IP address range you entered.

7. Use the exclusion range to define IP address ranges that are to be excluded from the scope. You can exclude multiple address ranges as follows:

- To define an exclusion range, enter a start address and an end address in the Start IP Address and End IP Address boxes, and then click Add.
- To track which address ranges are excluded, use the Excluded Addresses list. To
- delete an exclusion range, select the range in the Excluded Addresses list, and then click Remove.



8. Click Next. Specify the duration of leases for the scope by using the Day(s), Hour(s), and Minutes boxes. The default duration is 30 days. Click Next.



**TIP** If you haven't worked a lot with multicast, you shouldn't change the default value. Multicast leases aren't used in the same way as normal leases. Multiple computers can use a multicast IP address, and all of these computers can have a lease on the IP address. A good multicast lease duration for most networks is from 30 to 60 days.

9. If you want to activate the scope, select Yes, and then click Next. Otherwise, select No, and then click Next.
10. Click Finish to complete the process.

## Setting Scope Options

Scope options enable you to precisely control a scope's functioning and to set default TCP/IP settings for clients that use the scope. For example, you can use scope options to enable clients to automatically find DNS servers on the network. You can also



define settings for default gateways, WINS, and more. Scope options apply only to normal scopes, not to multicast scopes.

You can set scope options in any of the following ways:

- Globally for all scopes by setting default server options
- On a per-scope basis by setting scope options
- On a per-client basis by setting reservation options
- On a client-class basis by configuring user-specific or vendor-specific classes

IPv4 and IPv6 have different scope options. Scope options use a hierarchy to determine when certain options apply. The previous list shows the hierarchy. Basically, this means the following:

- Per-scope options override global options
- Per-client options override per-scope and global options
- Client-class options override all other options

## Assigning Server Options

Server options are applied to all scopes configured on a particular DHCP server. You can view and assign server options by following these steps:

1. In the DHCP console, double-click the server with which you want to work, and then expand its IPv4 and IPv6 folders in the tree view.
2. To view current settings, select the Server Options node under IPv4 or IPv6, depending on the type of address with which you want to work. Currently configured options are displayed in the right pane.
3. To assign new settings, right-click Server Options, and then click Configure Options to open the Server Options dialog box. Under Available Options, select the check box for the first option you want to configure. Then, with the option selected, enter any required information in the Data Entry panel. Repeat this step to configure other options.
4. Click OK to save your changes.


## Assigning Scope Options

Scope options are specific to an individual scope and override the default server options. You can view and assign scope options by following these steps:

1. In the DHCP console, expand the entry for the scope with which you want to work.



2. To view current settings, select Scope Options. Currently configured options are displayed in the right pane.
3. To assign new settings, right-click Scope Options, and then click Configure Options. This opens the Scope Options dialog box. Under Available Options, select the check box for the first option you want to configure. Then, with the option selected, enter any required information in the Data Entry panel. Repeat this step to configure other options.
4. Click OK.





## Assigning Reservation Options

You can assign reservation options to a client that has a reserved IPv4 or IPv6 address. These options are specific to an individual client and override server-specific and scope-specific options. To view and assign reservation options, follow these steps:

1. In the DHCP console, expand the entry for the scope with which you want to work.
2. Double-click the Reservations folder for the scope.
3. To view current settings, click the reservation you want to examine. Currently configured options are displayed in the right pane.

4. To assign new settings, right-click the reservation, and then click Configure Options to open the Reservation Options dialog box. Under Available Options, select the check box for the first option you want to configure. Then, with the option selected, enter any required information in the Data Entry panel. Repeat this step to configure other options.



## Modifying Scopes

You can modify an existing scope by following these steps:

1. In the DHCP console, double-click the server with which you want to work, and then expand its IPv4 and IPv6 folders in the tree view. This should display the currently configured scopes for the server.
2. Right-click the scope you want to modify, and then click Properties.
3. You can now modify the scope properties. Keep the following in mind:





- When you modify normal IPv4 scopes, you have the option of setting an unlimited lease expiration time. If you do, you create permanent leases that reduce the effectiveness of pooling IP addresses with DHCP. Permanent leases aren't released unless you physically release them or deactivate the scope. As a result, you might eventually run out of addresses, especially as your network grows. A better alternative to unlimited leases is to use address reservations, and then only for specific clients that need fixed IP addresses.
- When you modify multicast scopes, you have the option of setting a lifetime for the scope. The scope lifetime determines the amount of time the scope is valid. By default, multicast scopes are valid as long as they're activated. To change this setting, click the Lifetime tab, select Multicast Scope Expires On, and then set an expiration date.

## Activating and Deactivating Scopes

In the DHCP console, inactive scopes are displayed with an icon showing a red arrow pointing down. Active scopes display the standard folder icon.

You can activate an inactive scope by right-clicking it in the DHCP console, and then selecting Activate. You can deactivate an active scope by right-clicking it in the DHCP console, and then selecting Deactivate.

**IMPORTANT** Deactivating turns off a scope but doesn't terminate current client leases. If you want to terminate leases, follow the instructions in "Deleting Leases and Reservations" later in this chapter.





## Enabling the Bootstrap Protocol

Bootstrap Protocol (BOOTP) is a dynamic IPv4 addressing protocol that predates DHCP. Normal scopes don't support BOOTP. To enable a scope to support BOOTP, follow these steps:

1. Right-click the normal scope for IPv4 addresses that you want to modify, and then click Properties.
2. On the Advanced tab, click Both to support DHCP and BOOTP clients.
3. As necessary, set a lease duration for BOOTP clients, and then click OK.

**NOTE** The typical lease duration for a BOOTP address is much longer than for a DHCP address. For BOOTP, the default of 30 days is a good compromise, though some scenarios may lend themselves to an unlimited lease duration.







## Removing a Scope

Removing a scope permanently deletes the scope from the DHCP server. To remove a scope, follow these steps:

1. In the DHCP console, right-click the scope you want to remove, and then click Delete.
2. When prompted to confirm that you want to delete the scope, click Yes.





## Configuring Multiple Scopes on a Network

You can configure multiple scopes on a single network. A single DHCP server or multiple DHCP servers can serve these scopes. However, any time you work with multiple scopes, it's extremely important that the address ranges used by different scopes not overlap. Each scope must have a unique address range. If it doesn't, the same IP address might be assigned to different DHCP clients, which can cause severe problems on the network.

To understand how you can use multiple scopes, consider the following scenario, in which each server has its respective DHCP scope IP address range on the same subnet:

- **Server A**    192.168.10.1 to 192.168.10.99
- **Server B**    192.168.10.100 to 192.168.10.199
- **Server C**    192.168.10.200 to 192.168.10.254

Each of these servers responds to DHCP discovery messages, and any of them can assign IP addresses to clients. If one of the servers fails, the other servers can continue to provide DHCP services to the network. To introduce fault tolerance and provide redundancy, you can use failover scopes as discussed in the next section.





# Creating and Managing Failover Scopes

Failover scopes are split between two DHCP servers to increase fault tolerance, provide redundancy over using a single DHCP server, and enable load balancing. With a failover scope, you identify the two DHCP servers that split the scope. If one of the servers becomes unavailable or overloaded, the other server can take its place by continuing to lease new IP addresses and renew existing leases. A failover scope can also help to balance server loads.



# Creating Failover Scopes

Failover scopes apply only to IPv4 addresses. You can split a single scope or a superscope that contains multiple scopes.

You create a failover scope on the DHCP server that you want to designate as the primary server by splitting an existing scope or superscope. During the failover-scope creation process, you need to specify the partner server with which you want to split the primary server's scope. This additional server acts as the secondary server for the scope. Because failover scopes are a server-side enhancement, no additional configuration is required for DHCP clients.

The way scope splitting works depends on the failover scope configuration settings. You do one of the following:

- **Optimize for load balancing** A failover scope optimized for load balancing has little or no time delay configured in its scope properties. With no time delay, both the primary and the secondary servers can respond to DHCP DISCOVER requests from DHCP clients. This enables the fastest server to respond to and accept a DHCPOFFER first. Fault tolerance continues to be a part of the scope. If one of the servers becomes unavailable or overloaded and is unable to respond to requests, the other server handles requests and continues distributing addresses until the normal process is restored. For load balancing, set Load Balance as the failover mode. •
- **Optimize for fault tolerance** A failover scope optimized for fault tolerance has an extended time delay configured in its scope properties. The time delay on the secondary DHCP server causes the server to respond with a delay to DHCP DISCOVER requests from DHCP clients. The delay on the secondary server enables the primary DHCP server to respond to and accept the DHCPOFFER first. However, if the primary server becomes unavailable or overloaded and is unable to respond to requests, the secondary server handles requests and continues distributing addresses until the primary server is available to service clients again. For fault tolerance, set Hot Standby as the failover mode.

You can create a failover scope by completing the following steps:

1. In the DHCP console, connect to the primary DHCP server for the failover scope. Double-click the entry for the primary server, and then expand its IPv4 folder in the tree view.

2. The scope with which you want to work must already be defined. Right-click the scope or superscope that you want to configure for failover, and then click Configure Failover to start the Configure Failover Wizard.
3. The scope you right-clicked previously is selected by default. Click Next.



4. Next, you need to specify the partner server to use for failover. Click Add Server. Use the options in the Add Server dialog box to select the secondary DHCP server for the failover scope, and then click OK. Clear the Reuse Existing Failover Relationships check box, and then click Next to continue.

**NOTE** If the scopes you are configuring already exists on the partner server, you'll see a warning that the scopes will need to be deleted on the partner server before you can continue. You'll need to open the DHCP console on the partner server and then delete the duplicated scope.

5. You'll next need to create a new failover relationship. Several parameters are set for you automatically, including:

- **Relationship name** A descriptor used to differentiate failover relationships. The default name is based on the name and domain of the servers in the partnership. Although you can change the name if you want, the default value usually is sufficient.

**Configure Failover**

**Specify the partner server to use for failover**

Provide the host name or IP address of the partner DHCP server with which failover should be configured.

You can select from the list of servers with an existing failover configuration or you can browse and select from the list of authorized DHCP servers.

Alternatively, you can type the host name or IP address of the partner server.

Partner Server:

☐ Reuse existing failover relationships configured with this server (if any exist).

**Add Server**

Select a server you want to add to your console.

☐ This server:


☒ This authorized DHCP server:

| Name                             | IP Address |
|----------------------------------|------------|
| corpserver12.imaginedlands.local | 10.0.0.171 |
| corpserver54.imaginedlands.local | 10.0.0.172 |

- Maximum Client Lead Time** Sets the lead time for assigning leases to clients. When a client is first assigned a lease, the lease is valid only for the interval defined by the maximum lead time and must be validated before the lead time expires. The lease is validated when the client contacts the originating DHCP server within the lead time interval. Once validated, the client is assigned the lease for the normal lease interval. For example, if the maximum lead time is 10 minutes and the normal lease duration is 8 days, the client is initially assigned a lease for 10 minutes and must contact the server again before the lead time interval expires to receive a full lease of 8 days. However, if the originating server is unavailable for some reason, the client will send out a message to which any available DHCP server can respond to confirm the full 8-day lease.

**REAL WORLD** Wondering why this lead time process is used? This process is designed to ensure client leases aren't inadvertently duplicated or lost should failover partners lose communication with each other. Normally, if two DHCP servers share a failover partnership, partner A and B share information about leases with each other. However, if partner A assigns a lease and then goes offline before telling partner B about it, partner B will learn about the lease when the client tries to renew the lease before the lead time expires. Partner B will then register the lease and assign the client a full lease. As a best practice, you'll want

to use a maximum lead time of 5 to 60 minutes. A shorter duration ensures clients are more quickly assigned full leases.



Configure Failover


Create a new failover relationship

Create a new failover relationship with partner corpserver12.imagedlands.local

Relationship Name: corpserver54-corpserver12.imagedlands.local

Maximum Client Lead Time: 1 hours 0 minutes

6. Use the Mode list to set the failover mode as Load Balance or Hot Standby.
7. If you set the failover mode for Load Balance, use the Load Balance Percentage combo boxes to specify the relative percentage for how to allocate the IP addresses to each of the servers. Here are configuration examples:
  - An 80/20 split works best when you want one server to handle most of the workload and want another server to be available as needed.
  - An 60/40 split works best when you want one server to handle a little more of the workload than the other, but you want both servers to have regular workloads.
  - A 50/50 split works best when you want to evenly balance the load between two



Mode: Load balance

Load Balance Percentage

Local Server: 50%

Partner Server: 50%

servers.

8. If you set the failover mode to Hot Standby, set the role of the partner as either Active or Standby and then specify the relative percentage of IP addresses to reserve for the standby partner. By default, 5 percent of the IP addresses are reserved for the standby partner.

Mode: Hot standby

Hot Standby Configuration

Role of Partner Server: Standby

Addresses reserved for standby server: 5 %

9. The State Switchover Interval enables automatic switchover by setting the amount of time a DHCP server waits before assuming control of an address range in the event a failover partner becomes unavailable. For automatic transition to occur, you must enable and set the State Switchover Interval to a specific value. Otherwise, if a server is unable to communicate with a failover partner, the server remains in a communication interrupted state until an administrator manually changes the server to a partner down state, which allows the server to respond to all DHCP client requests. As a best practice, you'll want to use a switchover interval of 30 to 90 minutes. This interval ensures switchover is handled in a timely manner while allowing for intermittent communication outages that may occur.
10. Message authentication is enabled by default to ensure all DHCP messages are authenticated. Authentication prevents spoofing and other attacks against your servers.


☒ State Switchover Interval: 60 minutes

☒ Enable Message Authentication

Shared Secret:


11. Enter a shared secret for the partners. The shared secret is a password that the partners use when synchronizing the DHCP database and performing other tasks related to maintaining the DHCP failover partnership. Like other passwords, the shared secret should be relatively complex to prevent spoofing and other attacks against your servers. When you are ready to continue, click Next.
12. Click Next and then click Finish. Review the summary of the failover scope configuration. If any errors were encountered, you might need to take corrective action. Click Close.





## Modifying or Removing Failover Scopes

Failover scopes are not identified as such in the DHCP console. You can identify a failover scope by its network ID and IP address pool. Generally, you'll find a scope with the same network ID on two DHCP servers, and the scope properties will include information about the failover partnership. To view this information, right-click the scope, and then select Properties. In the Properties dialog box, select the Failover tab.



You can manage the partnership in several ways:

- If you suspect the configuration details related to the partnership are out of sync, rightclick the scope, and then select Replicate Partnership.
- If you suspect the DHCP database that the partners share is out of sync, right-click the scope, and then select Replicate Scope.
- If you no longer want the scope to fail over, you can deconfigure failover by rightclicking the scope, and then selecting Deconfigure Failover.



You can't modify the failover settings after the partnership is established; however, you can deconfigure failover and then reconfigure failover.



# Creating and Managing Superscopes


A superscope is a container for IPv4 scopes in much the same way that an organizational unit is a container for Active Directory objects. Superscopes help you manage scopes available on the network by grouping them into a single point of management. For example, with a superscope, you can activate or deactivate multiple scopes through a single action. You can also view statistics for all scopes in the superscope rather than having to check statistics for each scope. Superscopes also support DHCP clients on a single physical network where multiple logical IP networks are used, or put another way, you can create superscopes to distribute IP addresses from different logical networks to the same physical network segment.



# Creating Superscopes


After you create at least one normal or multicast IPv4 scope, you can create a superscope by following these steps:

1. In the DHCP console, expand the node for the server with which you want to work, right-click IPv4, and then click New Superscope to start the New Superscope Wizard. Click Next.
2. Enter a name for the superscope, and then click Next.



The screenshot shows the 'New Superscope Wizard' dialog box. The title bar says 'New Superscope Wizard'. Below the title bar, the text 'Superscope Name' is followed by the instruction 'You have to provide an identifying superscope name.' To the right of this text is a folder icon. At the bottom, there is a text box labeled 'Name:' containing the text 'Building 110 Scopes'.

3. Select scopes to add to the superscope. Select individual scopes by clicking their entry in the Available Scopes list. Select multiple scopes by clicking entries while holding down Shift or Ctrl.



The screenshot shows the 'New Superscope Wizard' dialog box at the 'Select Scopes' step. The title bar says 'New Superscope Wizard'. Below the title bar, the text 'Select Scopes' is followed by the instruction 'You create a superscope by building a collection of scopes.' To the right of this text is a folder icon. Below this, the text 'Select one or more scopes from the list to add to the superscope.' is followed by the label 'Available scopes:'. Below the label is a list box containing the following items: '[10.0.0.0] Primary Scope', '[192.168.0.0] Secondary IPv4', '[192.168.1.0] Secondary IPv4', '[192.168.2.0] Secondary IPv4', '[192.168.3.0] Secondary IPv4', '[192.168.4.0] Secondary IPv4', and '[192.168.5.0] Secondary IPv4'. The entire list box is highlighted with a blue background.

4. Click Next, and then click Finish.

## Adding Scopes to a Superscope

You can add scopes to a superscope when you create it, or you can add the scopes later. To add a scope to a superscope, follow these steps:

1. Right-click the scope you want to add to a superscope, and then click Add To Superscope.
2. In the Add Scope To A Superscope dialog box, select a superscope.
3. Click OK. The scope is then added to the superscope.



Add Scope 192.168.0.0 to a Superscope?×

Select a superscope for this scope.

Available superscopes:

Superscope 0

OK

Cancel



# Removing Scopes From a Superscope

To remove a scope from a superscope, follow these steps:

1. Right-click the scope you want to remove from a superscope, and then click Remove From Superscope.
2. Confirm the action by clicking Yes when prompted. If this is the last scope in the superscope, the superscope is deleted automatically.





## Activating and Deactivating a Superscope

When you activate or deactivate a superscope, you make all the scopes within the superscope active or inactive. To activate a superscope, right-click the superscope, and then select Activate. To deactivate a superscope, right-click the superscope, and then select Deactivate.





## Deleting a Superscope

Deleting a superscope removes the superscope container but doesn't delete the scopes it contains. If you want to delete the member scopes, you'll need to do that separately. To delete a superscope, right-click the superscope, and then select Delete. When prompted, click Yes to confirm the action.







# Managing Addressing

Scopes have separate folders for address pools, leases, and reservations. By accessing these folders, you can view current statistics for the related data and manage existing entries.





## Viewing Scope Statistics

Scope statistics provide summary information about the address pool for the current scope or superscope. To view statistics, right-click the scope or superscope, and then select Display Statistics.

The primary columns in the Scope Statistics dialog box are used as follows:


- **Total Scopes** Shows the number of scopes in a superscope.
- **Total Addresses** Shows the total number of IP addresses assigned to the scope.
- **In Use** Shows the total number (as a numerical value and as a percentage of the total available addresses) of addresses being used. If the total reaches 85 percent or more, you might want to consider assigning additional addresses or freeing up addresses for use.
- **Available** Shows the total number (as a numerical value and as a percentage of the total available addresses) of addresses available for use.



# Enabling and Configuring MAC Address Filtering

MAC address filtering (aka link-layer filtering ) is a feature for IPv4 addresses that enables you to include or exclude computers and devices based on their MAC address. When you configure MAC address filtering, you can specify the hardware types that are exempted from filtering. By default, all hardware types defined in RFC 1700 are exempted from filtering. To modify hardware type exemptions, follow these steps:

1. In the DHCP console, right-click the IPv4 node, and then click Properties.
2. On the Filters tab, click Advanced. In the Advanced Filter Properties dialog box, select the check box for hardware types to exempt from filtering. Clear the check box for hardware types to filter.
3. Click OK to save your changes.




Before you can configure MAC address filtering, you must do one of the following:

- Enable and define an explicit allow list. The DHCP server provides DHCP services only to clients whose MAC addresses are in the allow list. Any client that previously received IP addresses is denied address renewal if its MAC address isn't on the allow list.
- Enable and define an explicit deny list. The DHCP server denies DHCP services only to clients whose MAC addresses are in the deny list. Any client that previously received IP addresses is denied address renewal if its MAC address is on the deny list.

- Enable and define an allow list and a block list. The block list has precedence over the allow list. This means that the DHCP server provides DHCP services only to clients whose MAC addresses are in the allow list, if no corresponding matches are in the deny list. If a MAC address has been denied, the address is always blocked even if the address is on the allow list.

To enable an allow list, deny list, or both, follow these steps:

1. In the DHCP console, right-click the IPv4 node, and then click Properties.
2. On the Filters tab, you'll find the current filter configuration details. To use an allow list, select Enable Allow List. To use a deny list, select Enable Deny List.
3. Click OK to save your changes.



**NOTE** As an alternative, you can right-click the Allow or Deny node, under the Filters node, and then select Enable to enable allow or deny lists. If you right-click the Allow or Deny node and then select Disable, you disable allow or deny lists.

After you enable filtering, you define your filters by using the MAC address for the client computer or device's network adapter. On a client computer, you can obtain the MAC address by entering the command **ipconfig /all** at the command prompt. The Physical Address entry shows the client's MAC address. You must enter this value exactly for the address filter to work.

A MAC address is defined by eight pairings of two-digit hexadecimal numbers separated by a hyphen, as shown here:

FE-01-56-23-18-94-EB-F2

When you define a filter, you can specify the MAC address with or without the hyphens. This means that you could enter FE-01-56-23-18-94-EB-F2 or FE0156231894EBF2.

You also can use an asterisk (\*) as a wildcard for pattern matching. To allow any value to match a specific part of the MAC address, you can insert \* where the values usually would be, as shown here:

FE-01-56-23-18-94-\*-F2


FE-\*-56-23-18-94-\*- FE-01-56-23-18-\*-

\*\_\*

FE01\*

To configure a MAC address filter, follow these steps:


1. In the DHCP console, double-click the IPv4 node, and then double-click the Filters node.
2. Right-click Allow or Deny as appropriate for the type of filter you are creating, and then click New Filter.
3. Enter the MAC address to filter, and then if you want to you can enter a comment in the Description text box. Click Add. Repeat this step to add other filters.
4. Click Close when you have finished.



## Setting a New Exclusion Range

You can exclude IPv4 or IPv6 addresses from a scope by defining an exclusion range. Scopes can have multiple exclusion ranges. To define an exclusion range for a scope with IPv4 addresses, follow these steps:

1. In the DHCP console, expand the scope with which you want to work, and then right-click the Address Pool folder or Exclusions folder. On the shortcut menu, click New Exclusion Range.
2. Enter a start address and an end address in the Start IP Address and End IP Address boxes, and then click Add. The range specified must be a subset of the range set for the current scope and must not be currently in use. Repeat this step to add other exclusion ranges.
3. Click Close when you have finished.



**TIP** To exclude a single IP address, enter a start address and leave the end address blank.

To define an exclusion range for a scope with IPv6 addresses, follow these steps:

1. In the DHCP console, expand the scope with which you want to work, and then right-click the Exclusions folder. On the shortcut menu, click New Exclusion Range.
2. Enter a start address and an end address in the Start IPv6 Address and End IPv6 Address boxes, and then click Add. The range specified must be a subset of the range set for the current scope and must not be currently in use. Repeat this step to add other exclusion ranges.
3. Click Close when you have finished.

If you don't need an exclusion anymore, you can delete it. Select Address Pool or Exclusions as appropriate. In the main pane, right-click the exclusion, select Delete, and then click Yes in response to the confirmation message.

## Reserving DHCP Addresses

DHCP provides several ways to assign permanent addresses to clients. One way is to use the Unlimited setting in the Scope dialog box to assign permanent addresses to all clients that use the scope. Another way is to reserve DHCP addresses on a per-client basis. When you reserve a DHCP address, the DHCP server always assigns the client the same IP address, and you can do so without sacrificing the centralized management features that make DHCP so attractive.

If the client is on the network and has a current IPv4 or IPv6 lease, you can create a reservation by completing the following steps:


1. In the DHCP console, expand the scope with which you want to work, and then select the Address Leases folder.

2. Right-click the lease you want to work with. On the shortcut menu, click Add To Reservation.

Otherwise, to manually reserve an IPv4 address for a client, follow these steps:

1. In the DHCP console, expand the scope with which you want to work, and then right-click the Reservations folder. On the shortcut menu, click New Reservation.
2. In the Reservation Name text box, enter a short but descriptive name for the reservation. This name is used only for identification purposes.
3. In the IP Address box, enter the IPv4 address you want to reserve for the client.

**NOTE** This IP address must be within the valid range of addresses for the currently selected scope.



The screenshot shows a 'New Reservation' dialog box with the following fields and options:

- Reservation name:** SStaff142
- IP address:** 10 . 0 . 0 . 99
- MAC address:** AB-DE-88-45-23-67
- Description:** (empty text box)
- Supported types:**
  - ☒ Both
  - ☐ DHCP
  - ☐ BOOTP

At the bottom are 'Add' and 'Close' buttons.

4. The MAC Address box specifies the MAC address for the client computer's NIC. You can obtain the MAC address by entering the command **ipconfig /all** at the command prompt on the client computer. The Physical Address entry shows the client's MAC address. You must enter this value exactly for the address reservation to work.
5. Enter an optional comment in the Description text box.
6. By default, both DHCP and BOOTP clients are supported. This option is fine, and you need to change it only if you want to exclude a particular type of client.
7. Click Add to create the address reservation. Repeat this step to add other address reservations.

8. Click Close when you have finished.

To manually reserve an IPv6 address for a client, follow these steps:

1. In the DHCP console, expand the scope with which you want to work, and then right-click the Reservations folder. On the shortcut menu, click New Reservation.
2. In the Reservation text box, enter a short but descriptive name for the reservation. This information is used only for identification purposes.
3. In the IPv6 Address box, enter the IPv6 address you want to reserve for the client.

**NOTE** This IP address must be within the valid range of addresses for the currently selected scope.

4. The device unique identifier (DUID) box specifies the MAC address for the client computer's NIC. You can obtain the MAC address by entering the command **ipconfig /all** at the command prompt on the client computer. The Physical Address entry shows the client's MAC address. You must enter this value exactly for the address reservation to work.
5. The identity association identifier (IAID) sets a unique identifier prefix for the client. Typically, this is a nine-digit value.
6. Enter an optional comment in the Description text box.
7. Click Add to create the address reservation. Repeat this step to add other address reservations.
8. Click Close when you have finished.

When you work with reserved addresses, you should take note of a couple caveats:

- Reserved addresses aren't automatically reassigned. If the address is already in use, you need to release the address to ensure that the appropriate client can obtain it. You can force a client to release an address by terminating the client's lease or by logging on to the client and entering the command **ipconfig /release** at an administrator command prompt.
- Clients don't automatically switch to the reserved address. If the client is using a different IP address, you need to force the client to release the current lease and request a new one. You can do this by terminating the client's lease or by logging on to the client and entering the command **ipconfig /renew** at an administrator command prompt.





## Modifying Reservation Properties

You can modify the properties of reservations by following these steps:

1. In the DHCP console, expand the scope with which you want to work, and then click the Reservations folder.
2. Right-click a reservation, and then click Properties. You can now modify the reservation properties. You can't modify options that are shaded, but you can modify other options. These options are the same options described in the previous section.





## Deleting Leases and Reservations

You can delete active leases and reservations by following these steps:

1. In the DHCP console, expand the scope with which you want to work, and then click the Address Leases folder or Reservations folder, as appropriate.
2. Right-click the lease or reservation you want to delete, and then click Delete.
3. Confirm the deletion by clicking Yes.
4. The lease or reservation is now removed from DHCP; however, the client isn't forced to release the IP address. To force the client to release the IP address, log on to the client that holds the lease or reservation and enter the command **ipconfig /release** at an administrator command prompt.





# Backing Up and Restoring the DHCP Database

DHCP servers store DHCP lease and reservation information in database files. By default, these files are stored in the %SystemRoot%\System32\DHCP directory. The key files in this directory are used as follows:

- **Dhcp.mdb** The primary database file for the DHCP server
- **J50.log** A transaction log file used to recover incomplete transactions in case of a server malfunction
- **J50.chk** A checkpoint file used in truncating the transaction log for the DHCP server
- **J50000NN.log** A reserved log file for the DHCP server
- **Tmp.edb** A temporary working file for the DHCP server



## Backing Up the DHCP Database

The %SystemRoot%\System32\DHCP\Backup folder contains the backup information for the DHCP configuration and the DHCP database. By default, the DHCP database is backed up every 60 minutes automatically. To manually back up the DHCP database at any time, follow these steps:

1. In the DHCP console, right-click the server you want to back up, and then click Backup.
2. In the Browse For Folder dialog box, select the folder that will contain the backup DHCP database, and then click OK.

Registry keys that control the location and timing of DHCP backups, in addition to other DHCP settings, are located under HKEY\_LOCAL\_MACHINE in SYSTEM\CurrentControlSet\Services\DHCP\Parameters.

The following keys control the DHCP database and backup configuration:

- **BackupDatabasePath** Sets the location of the DHCP database. You should set this option through the DHCP Properties dialog box. Click the Advanced tab, and then set the Database Path as appropriate.
- **DatabaseName** Sets the name of the primary DHCP database file. The default value is DHCP.mdb.
- **BackupInterval** Determines how often the DHCP client information database is backed up. The default is 60 minutes.
- **DatabaseCleanupInterval** Determines how often the DHCP service deletes expired records from the DHCP client information database. The default is four hours.

## Restoring the DHCP Database From Backup

In the case of a server crash and recovery, you might need to restore and then reconcile the DHCP database. To force DHCP to restore the database from backup, follow these steps:

1. If necessary, restore a good copy of the %SystemRoot%\System32\DHCP\Backup directory from the archive. Afterward, start the DHCP console, right-click the server you want to restore, and then click Restore.
2. In the Browse For Folder dialog box, select the folder that contains the backup you want to restore, and then click OK.



3. During the restoration of the database, the DHCP Server service is stopped. As a result, DHCP clients are temporarily unable to contact the DHCP server to obtain IP addresses.





## Using Backup and Restore to Move the DHCP Database to a New Server

If you need to rebuild a server providing DHCP services, you might want to move the DHCP services to another server prior to rebuilding the server. To do this, you need to perform several tasks on the source and destination servers. On the destination server, do the following:

1. Install the DHCP Server service on the destination server, and then restart the server.
2. Stop the DHCP Server service in the Services console.
3. Delete the contents of the %SystemRoot%\System32\DHCP folder.

On the source server, do the following:

1. Stop the DHCP Server service in the Services console.
2. After the DHCP Server service is stopped, disable the service so that it can no longer be started.
3. Copy the entire contents of the %SystemRoot%\System32\DHCP folder to the %SystemRoot%\System32\DHCP folder on the destination server.

Now all the necessary files are on the destination server. Start the DHCP Server service on the destination server to complete the migration.

## Forcing the DHCP Server Service to Regenerate the DHCP Database

If the DHCP database becomes corrupt and Windows is unable to repair the database when you stop and restart the DHCP Server service, you can attempt to restore the database as described in “Restoring the DHCP Database From Backup” earlier in this chapter. If this fails or you prefer to start with a fresh copy of the DHCP database, follow these steps:

1. Stop the DHCP Server service in the Services console.
2. Delete the contents of the %SystemRoot%\System32\DHCP folder. If you want to force a complete regeneration of the database and not allow the server to restore from a previous backup, you should also delete the contents of the Backup folder.

**CAUTION** Don't delete DHCP files if the DHCP Server registry keys aren't intact. These keys must be available to restore the DHCP database.



- 
3. Restart the DHCP Server service.
  4. No active leases or other information for scopes are displayed in the DHCP console. To regain the active leases for each scope, you must reconcile the server scopes as discussed in the next section.
  5. To prevent conflicts with previously assigned leases, you should enable address conflict detection for the next few days, as discussed in “Avoiding IP Address Conflicts” earlier in Chapter 32.

## Reconciling Leases and Reservations

Reconciling checks the client leases and reservations against the DHCP database on the server. If inconsistencies are found between what is registered in the Windows registry and what is recorded in the DHCP server database, you can select and reconcile any inconsistent entries. After the entries you select are reconciled, DHCP either restores the IP address to the original owner or creates a temporary reservation for the IP address. When the lease time expires, the address is recovered for future use.

You can reconcile scopes individually, or you can reconcile all scopes on a server. To reconcile a scope individually, follow these steps:

1. In the DHCP console, right-click the scope with which you want to work, and then click Reconcile.
2. In the Reconcile dialog box, click Verify.
3. Inconsistencies are reported in the status window. Select the displayed addresses, and then click Reconcile to repair inconsistencies.
4. If no inconsistencies are found, click OK.

To reconcile all scopes on a server, follow these steps:

1. In the DHCP console, expand the server entry, right-click the IPv4 node, and then click Reconcile All Scopes.
2. In the Reconcile All Scopes dialog box, click Verify.
3. Inconsistencies are reported in the status window. Select the displayed addresses, and then click Reconcile to repair inconsistencies.
4. If no inconsistencies are found, click OK.



## Chapter 34. Implementing DNS

DNS is a name-resolution service that resolves computer names to IP addresses. When you use DNS, a fully qualified host name—pc15.tvpress.com, for example—can be resolved to an IP address, which enables computers to find one another. DNS operates over the TCP/IP protocol stack and can be integrated with Windows Internet Name Service (WINS), Dynamic Host Configuration Protocol (DHCP), and Active Directory Domain Services (AD DS). Fully integrating DNS with these Windows networking features enables you to optimize DNS for Windows Server domains.







# Understanding DNS

DNS organizes groups of computers into domains. These domains are organized into a hierarchical structure that can be defined on an Internet-wide basis for public networks or on an enterprise-wide basis for private networks (also known as extranets and intranets, respectively). The various levels within the hierarchy identify individual computers, organizational domains, and top-level domains. For the fully qualified host name `server4.reagentpress.com`, `server4` represents the host name for an individual computer, `reagentpress` is the organizational domain, and `com` is the top-level domain.

Top-level domains are at the root of the DNS hierarchy and are also called root domains. These domains are organized geographically, by organization type, and by function. Typical corporate domains, such as `microsoft.com`, are also referred to as parent domains because they're the parents of an organizational structure. You can divide parent domains into subdomains you can use for groups or departments within your organization.

Subdomains are often referred to as child domains. For example, the fully qualified domain name (FQDN) for a computer within a human resources group could be designated as `beowulf.hr.reagentpress.com`. Here, `beowulf` is the host name, `hr` is the child domain, and `reagentpress.com` is the parent domain.

## Integrating Active Directory and DNS

Active Directory domains use DNS to implement their naming structure and hierarchy. Active Directory and DNS are tightly integrated, so much so that you should install DNS on the network before you can install Active Directory Domain Services.

During installation of the first domain controller on an Active Directory network, you have the opportunity to automatically install DNS if a DNS server can't be found on the network. You can also specify whether DNS and Active Directory should be integrated fully. In most cases, you should respond affirmatively to both requests. With full integration, DNS information is stored directly in Active Directory, which enables you to take advantage of Active Directory's capabilities.

Understanding the difference between partial integration and full integration is very important:

- **Partial integration** With partial integration, the domain uses standard file storage. DNS information is stored in text-based files that end with the `.dns` extension. The default location of these files is `%SystemRoot%\System32\Dns`.



Updates to DNS are handled through a single authoritative DNS server. This server is designated as the primary DNS server for the particular domain or an area within a domain called a zone . Clients that use dynamic DNS updates through DHCP must be configured to use the primary DNS server in the zone. If they aren't, their DNS information won't be updated. Likewise, dynamic updates through DHCP can't be made if the primary DNS server is offline.

- **Full integration** With full integration, the domain uses directory-integrated storage. DNS information is stored directly in Active Directory and is available through the container for the dnsZone object. Because the information is part of Active Directory, any domain controller can access the data, and you can use a multimaster approach for dynamic updates through DHCP. This enables any domain controller running the DNS Server service to handle dynamic updates. Furthermore, clients that use dynamic DNS updates through DHCP can use any DNS server within the zone. An added benefit of directory integration is the ability to use directory security to control access to DNS information.

If you look at the way DNS information is replicated throughout the network, you will find more advantages to full integration with Active Directory. With partial integration, DNS information is stored and replicated separately from Active Directory. By having two separate structures, you reduce the effectiveness of both DNS and Active Directory and make administration more complex. Because DNS is less efficient than Active Directory at replicating changes, you might also increase network traffic and the amount of time required to replicate DNS changes throughout the network.

In early releases of the DNS Server service for Windows servers, restarting a DNS server could take an hour or more in large organizations with extremely large AD DS–integrated zones. The operation took this much time because the zone data was loaded in the foreground while the server was starting the DNS service. To ensure that DNS servers can be responsive after a restart, the DNS Server service loads zone data from AD DS in the background while the service restarts. This ensures that the DNS server is responsive and can handle requests for data from other zones.

At startup, DNS servers perform the following tasks:

- Enumerate all zones to be loaded
- Load root hints from files or AD DS storage
- Load all zones that are stored in files rather than in AD DS
- Begin responding to queries and Remote Procedure Calls (RPCs)
- Create one or more threads to load the zones that are stored in AD DS





Because separate threads load zone data, the DNS server is able to respond to queries while zone loading is in progress. If a DNS client performs a query for a host in a zone that has already been loaded, the DNS server responds appropriately. If the query is for a host that has not yet been loaded into memory, the DNS server reads the host's data from AD DS and updates its record list accordingly.





## Enabling DNS on the Network

To enable DNS on the network, you need to configure DNS clients and servers. When you configure DNS clients, you tell the clients the IP addresses of DNS servers on the network. By using these addresses, clients can communicate with DNS servers anywhere on the network, even if the servers are on different subnets.

The DNS client built into computers running Windows supports DNS traffic over Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6). By default, IPv6 automatically configures the site-local address of DNS servers. To add the IPv6 addresses of your DNS servers, use the properties of the Internet Protocol Version 6 (TCP/IPv6) component in Network Connections or the following command: `netsh interface IPV6 ADD DNSSERVERS`

In Windows PowerShell, you can use `Get-NetIPInterface` to list the available interfaces and then use `Set-DNSClientServerAddress` to set the IPv6 address on a specified interface.

DNS servers support IPv6 addresses as fully as they support IPv4 addresses. In the DNS Manager console, host addresses are displayed as IPv4 or IPv6 addresses. The `Dnscmd` command-line tool also accepts addresses in either format. Additionally, DNS servers can now send recursive queries to IPv6-only servers, and the server forwarder list can contain both IPv4 and IPv6 addresses. Finally, DNS servers now support the `ip6.arpa` domain namespace for reverse lookups.

When the network uses DHCP, you should configure DHCP to work with DNS. DHCP clients can register IPv6 addresses along with or instead of IPv4 addresses. To ensure proper integration of DHCP and DNS, you need to set the DHCP scope options as specified in “Setting Scope Options” in Chapter 33, “Managing and Maintaining DHCP.” For IPv4, you should set the 006 DNS Servers and 015 DNS Domain Name scope options. For IPv6, you should set the 00023 DNS Recursive Name Server IPv6 Address List and 00024 Domain Search List scope options. Additionally, if computers on the network need to be accessible from other Active Directory domains, you need to create records for them in DNS. DNS records are organized into zones, where a zone is simply an area within a domain.

DNS client computers running Windows can use Link-Local Multicast Name Resolution (LLMNR) to resolve names on a local network segment when a DNS server is not available. They also periodically search for a domain controller in the domain to which they belong. This functionality helps avoid performance problems that might occur if a network or server failure causes a DNS client to create an





association with a distant domain controller located on a slow link rather than a local domain controller.

Previously, this association continued until the client was forced to seek a new domain controller, such as when the client computer was disconnected from the network for a long period of time. By periodically renewing its association with a domain controller, a DNS client can reduce the probability that it will be associated with an inappropriate domain controller.

The DNS client service for current versions of Windows has several interoperability and security enhancements specific to LLMNR and NetBIOS. To improve security for mobile networking, the service

- Does not send outbound LLMNR queries over mobile broadband or VPN interfaces
- Does not send outbound NetBIOS queries over mobile broadband

For better compatibility with devices in power-saving mode, the LLMNR query timeout is set to 410 milliseconds (msec) for the first retry and 410 msec for the second retry, making the total timeout value 820 msec. To improve response times for all queries, the DNS client service does the following:

- Issues LLMNR and NetBIOS queries in parallel, and optimizes for IPv4 and IPv6
- Divides interfaces into networks to send parallel DNS queries
- Uses asynchronous DNS cache with an optimized response timing

**NOTE** You can configure a DNS client computer running Windows 7 or later, in addition to Windows Server 2008 R2 or later, to locate the nearest domain controller instead of searching randomly. This can improve performance in networks containing domains that exist across slow links. However, because of the network traffic this process generates, locating the nearest domain controller can have a negative impact on network performance.

Windows Server supports read-only primary zones and the GlobalNames zone. To support read-only domain controllers (RODCs), the primary read-only zone is created automatically. When a computer becomes an RODC, it replicates a full read-only copy of all the application directory partitions that DNS uses, including the domain partition, ForestDNSZones, and DomainDNSZones. This ensures that the DNS server running on the RODC has a full read-only copy of any DNS zones. As an administrator of an RODC, you can view the contents of a primary read-only zone. You cannot, however, change the contents of a zone on the RODC. You can change the contents of the zone only on a standard domain controller.





To support all DNS environments and single-label name resolution, you can create a zone named GlobalNames . For optimal performance and cross-forest support, you should integrate this zone with AD DS and configure each authoritative DNS server with a local copy. When you use Service Location (SRV) resource records to publish the location of the GlobalNames zone, this zone provides unique, single-label computer names across the forest. Unlike WINS, the GlobalNames zone is intended to provide single-label name resolution for a subset of host names—typically, the CNAME resource records for your corporate servers. The GlobalNames zone is not intended to be used for peer-to-peer name resolution, such as name resolution for workstations. This is what LLMNR is for.

When the GlobalNames zone is configured appropriately, single-label name resolution works as follows:

1. The client's primary DNS suffix is appended to the single-label name that the client is looking up, and the query is submitted to the DNS server.
2. If that computer's full name is not resolved, the client requests resolution by using its DNS suffix search lists, if any.
3. If none of those names can be resolved, the client requests resolution by using the single-label name.
4. If the single-label name appears in the GlobalNames zone, the DNS server hosting the zone resolves the name. Otherwise, the query fails over to WINS.

Dynamic updates in the GlobalNames zone are not supported.




# Configuring Name Resolution on DNS Clients

The best way to configure name resolution for DNS clients depends on the configuration of your network. If computers use DHCP, you probably want to configure DNS through settings on the DHCP server. If computers use static IP addresses or you want to configure DNS specifically for an individual system, you should configure DNS manually.

You can configure DNS settings on the DNS tab of the Advanced TCP/IP Settings dialog box. To access this dialog box, follow these steps:

1. In Settings, click Network & Internet on the home page and then click Ethernet. On the Ethernet page, click Change Adapter Options.



2. In Network Connections, right-click the connection with which you want to work, and then click Properties.
3. Double-click Internet Protocol Version 6 (TCP/IPv6) or Internet Protocol Version 4 (TCP/IPv4), depending on the type of IP address you are configuring.
4. If the computer is using DHCP and you want DHCP to specify the DNS server address, select Obtain DNS Server Address Automatically. Otherwise, select Use The Following DNS Server Addresses, and then enter primary and alternate DNS server addresses in the text boxes provided.
5. Click Advanced to display the Advanced TCP/IP Settings dialog box. In this dialog box, click the DNS tab.

You use the options of the DNS tab as follows:

- **DNS Server Addresses, In Order Of Use** Use this area to specify the IP address of each DNS server that is used for domain name resolution. Click Add if you want to add a server IP address to the list. Click Remove to remove a selected server





address from the list. Click Edit to edit the selected entry. You can specify multiple servers for DNS resolution. Their priority is determined by the order. If the first server isn't available to respond to a host name resolution request, the next DNS server in the list is accessed, and so on. To change the position of a server in the list box, select it, and then use the up or down arrow.

- **Append Primary And Connection Specific DNS Suffixes** Typically, this option is selected by default. Select this option to resolve unqualified computer names in the primary domain. For example, if the computer name Gandolf is used and the parent domain is microsoft.com, the computer name would resolve to gandolf.microsoft.com. If the fully qualified computer name doesn't exist in the parent domain, the query fails.

The parent domain used is the one set on the Computer Name tab in the System Properties dialog box. (Click System And Security\System in Control Panel, click Change Settings, and then display the Computer Name tab to check the settings.) •

- **Append Parent Suffixes Of The Primary DNS Suffix** This option is selected by default. Select this option to resolve unqualified computer names by using the parent/child domain hierarchy. If a query fails in the immediate parent domain, the suffix for the parent of the parent domain is used to try to resolve the query. This process continues until the top of the DNS domain hierarchy is reached. For example, if the computer name Gandolf is used in the dev.microsoft.com domain, DNS would attempt to resolve the computer name to gandolf.dev.microsoft.com. If this didn't work, DNS would attempt to resolve the computer name to gandolf.microsoft.com. •

- **Append These DNS Suffixes (In Order)** Select this option to set specific DNS suffixes to use rather than resolving through the parent domain. Click Add if you want to add a domain suffix to the list. Click Remove to remove a selected domain suffix from the list. Click Edit to edit the selected entry. You can specify multiple domain suffixes, which are used in order. If the first suffix is not resolved properly, DNS attempts to use the next suffix in the list. If this fails, the next suffix is used, and so on. To change the order of the domain suffixes, select the suffix, and then use the up or down arrow to change its position. This option is especially useful in hybrid namespaces where there are multiple parent domain names.

- **DNS Suffix For This Connection** This option sets a specific DNS suffix for the connection that overrides DNS names already configured for use on this connection. You usually set the DNS domain name on the Computer Name tab of the System Properties dialog box.

- **Register This Connection's Addresses In DNS** Select this option if you want all IP addresses for this connection to be registered in DNS under the computer's fully qualified domain name. This option is selected by default.







**NOTE** Dynamic DNS updates are used in conjunction with DHCP to enable a client to update its A (Host Address) record if its IP address changes and to enable the DHCP server to update the PTR (Pointer) record for the client on the DNS server. You can also configure DHCP servers to update both the A and PTR records on the client's behalf. Dynamic DNS updates are supported by DNS servers with BIND 8.2.1 or later in addition to all server versions of Windows.

- **Use This Connection's DNS Suffix In DNS Registration** Select this option if you want all IP addresses for this connection to be registered in DNS under the parent domain.





# Installing DNS Servers

You can configure any Windows Server 2016 system as a DNS server. Four types of DNS servers are available:

- **Active Directory–integrated primary server** A DNS server that's fully integrated with Active Directory. All DNS data is stored directly in Active Directory.
- **Primary server** The main DNS server for a domain that is partially integrated with Active Directory. This server stores a master copy of DNS records and the domain's configuration files. These files are stored as text files with the .dns extension.
- **Secondary server** A DNS server that provides backup services for the domain. This server stores a copy of DNS records obtained from a primary server and relies on zone transfers for updates. Secondary servers obtain their DNS information from a primary server when they are started, and they maintain this information until the information is refreshed or expired.
- **Forwarding-only server** A server that caches DNS information after lookups and always passes requests to other servers. These servers maintain DNS information until it's refreshed or expired or until the server is restarted. Unlike secondary servers, forwarding-only servers don't request full copies of a zone's database files. This means that when you start a forwarding-only server, its database contains no information.

Before you configure a DNS server, you must install the DNS Server service. Then you can configure the server to provide integrated, primary, secondary, or forwarding-only DNS services.

## Installing and Configuring the DNS Server Service


All domain controllers can act as DNS servers, and you might be prompted to install and configure DNS during installation of the domain controller. If you respond affirmatively to the prompts, DNS is already installed, and the default configuration is set automatically. You don't need to reinstall.

If you're working with a member server instead of a domain controller, or if you haven't installed DNS, follow these steps to install DNS:

1. In Server Manager, click Manage, and then click Add Roles And Features, or select Add Roles And Features in the Quick Start pane to start the Add Roles And Features Wizard. If the wizard displays the Before You Begin page, read the Welcome text, and then click Next.



2. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Click Next.
3. On the Server Selection page, you can choose to install roles and features on running servers or virtual hard disks. Either select a server from the server pool or select a server from the server pool on which to mount a virtual hard disk (VHD). If you are adding roles and features to a VHD, click Browse and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, click Next.
4. On the Server Roles page, select DNS Server. If additional features are required to install a role, you'll get an additional dialog box. Click Add Features to close the dialog box, and add the required features to the server installation. When you are ready to continue, click Next three times.




5. If the server on which you want to install the DNS Server role doesn't have all the required binary source files, the server gets the files via Windows Update by default or from a location specified in Group Policy.


**NOTE** You also can specify an alternate path for the required source files. To do this, click the Specify An Alternate Source Path link, enter that alternate path in the box provided, and then click OK. For network shares, enter the UNC path to the share, such as \\CorpServer82\\WinServer2012\\. For mounted Windows images, enter the WIM path prefixed with WIM: and including the index of the image to use, such as WIM:\\CorpServer82\\WinServer2012\\install.wim:4.

6. Click Install to begin the installation process. The Installation Progress page tracks the progress of the installation. If you close the wizard, click the Notifications icon in Server Manager, and then click the link provided to reopen the wizard.


- When Setup finishes installing the DNS Server role, the Installation Progress page will be updated to reflect this. Review the installation details to ensure that the installation was successful.




- From now on, the DNS Server service should start automatically each time you restart the server. If it doesn't start, you need to start it manually. (See "Starting and Stopping a DNS Server" later in Chapter 35.)
- After you install a DNS server, you use the DNS console to configure and manage DNS. In Server Manager, click Tools, and then click DNS to open the DNS Manager console.



- If the server you want to configure isn't listed in the tree view, you need to connect to the server. Right-click DNS in the tree view, and then click Connect To DNS Server. Now do one of the following:
  - If you're trying to connect to a local server, select This Computer, and then click OK.
  - If you're trying to connect to a remote server, select The Following Computer, enter the server's name or IP address, and then click OK.



11. An entry for the DNS server should be listed in the tree view pane of the DNS Manager console. Right-click the server entry, and then click Configure A DNS Server to start the Configure A DNS Server Wizard. Click Next.
12. On the Select Configuration Action page, select Configure Root Hints Only to specify that only the base DNS structures should be created at this time.




13. Click Next. The wizard searches for existing DNS structures and modifies them as necessary.
14. Click Finish to complete the process.

**REAL WORLD** If the wizard wasn't able to configure the root hints, you might need to configure them manually or copy them from another server. However, a default set of root hints is included with DNS Server, and these root hints should be added automatically. To confirm, right-click the server entry in the DNS console, and then select Properties. In the Properties dialog box, the currently configured root hints are shown on the Root Hints tab.

# Configuring a Primary DNS Server

Every domain should have a primary DNS server. You can integrate this server with Active Directory, or it can act as a standard primary server. Primary servers should have forward lookup zones and reverse lookup zones. You use forward lookups to resolve domain names to IP addresses. You need reverse lookups to authenticate DNS requests by resolving IP addresses to domain names or hosts.



After you install the DNS Server service on the server, you can configure a primary server by following these steps:

1. Start the DNS Manager console. If the server you want to configure isn't listed, connect to it as described previously in "Installing and configuring the DNS Server service."
2. An entry for the DNS server should be listed in the tree view pane of the DNS Manager console. Right-click the server entry, and then click New Zone to start the New Zone Wizard. Click Next.
3. You can now select the zone type. If you're configuring a primary server integrated with Active Directory (on a domain controller), select Primary Zone and be sure that the Store The Zone In Active Directory check box is selected. If you don't want to integrate DNS with Active Directory, select Primary Zone, and then clear the Store The Zone In Active Directory check box. Click Next.



- 
4. If you're integrating the zone with Active Directory, choose one of the following replication strategies; otherwise, proceed to step 6.

- **To All DNS Servers Running On Domain Controllers In This Forest** Choose this strategy if you want the widest replication strategy. Remember, the Active Directory forest includes all domain trees that share the directory data with the current domain.

- **To All DNS Servers Running On Domain Controllers In This Domain** Choose this strategy if you want to replicate DNS information within the current domain.
- **To All Domain Controllers In This Domain** Choose this strategy if you want to replicate DNS information to all domain controllers within the current domain, as needed for legacy compatibility. Although this strategy gives wider replication for DNS information within the domain and supports compatibility with legacy Windows, not every domain controller is a DNS server as well (nor do you need to configure every domain controller as a DNS server).

5. Click Next. Select Forward Lookup Zone, and then click Next.
6. Enter the full DNS name for the zone. The zone name should help determine how the server or zone fits into the DNS domain hierarchy. For example, if you're creating the primary server for the imaginedlands.local domain, you would enter **imaginedlands.local** as the zone name. Click Next.
7. If you're configuring a primary zone that isn't integrated with Active Directory, you need to set the zone file name. A default name for the zone's DNS database file should be filled in for you. You can use this name or enter a new file name. Click Next.
8. Specify whether dynamic updates are allowed. You have three options:

- **Allow Only Secure Dynamic Updates** When the zone is integrated with Active Directory, you can use access control lists (ACLs) to restrict which clients can perform dynamic updates. With this option selected, only clients with authorized computer accounts and approved ACLs can dynamically update their resource records in DNS when changes occur.
- **Allow Both Nonsecure And Secure Dynamic Updates** Choose this option to allow any client to update its resource records in DNS when changes occur. Clients can be secure or nonsecure.
- **Do Not Allow Dynamic Updates** Choose this option to disable dynamic updates in DNS. You should use this option only when the zone isn't integrated with Active Directory.

9. Click Next, and then click Finish to complete the process. The new zone is added to the server, and basic DNS records are created automatically.
10. A single DNS server can provide services for multiple domains. If you have multiple parent domains, such as microsoft.com and msn.com, you can repeat this process to configure other forward lookup zones. You also need to configure reverse lookup zones. Follow the steps listed in “Configuring Reverse Lookups” later in this chapter.
11. You need to create additional records for any computers you want to make accessible to other DNS domains. To do this, follow the steps listed in “Managing DNS Records” in Chapter 35.

**REAL WORLD** Most organizations have private and public areas of their network. The public network areas might be where web and external email servers reside. Your organization’s public network areas shouldn’t allow unrestricted access. Instead, public network areas should be configured as part of perimeter networks. (Perimeter networks are also known as DMZs , demilitarized zones, and screened subnets . These are areas protected by your organization’s firewall that have restricted external access and no access to the internal network.) Otherwise, public network areas should be in a completely separate and firewall-protected area.

The private network areas are where the organization’s internal servers and work stations reside. On the public network areas, your DNS settings are in the public Internet space. Here, you might use a .com, .org, or .net DNS name that you’ve registered with an Internet registrar and public IP addresses that you’ve purchased or leased. On the private network areas, your DNS settings are in the private network space. Here, you might use a local DNS name and private IP addresses, as discussed in Chapter 18.

## Configuring a Secondary DNS Server

Secondary servers provide backup DNS services on the network. If you’re using full Active Directory integration, you don’t really need to configure secondaries. Instead, you should configure multiple domain controllers to handle DNS services. Active Directory replication will then handle replicating DNS information to your domain controllers. On the other hand, if you’re using partial integration, you might want to configure secondaries to lessen the load on the primary server. On a small or mediumsize network, you might be able to use the name servers of your Internet service provider (ISP) as secondaries. In this case, you should contact your ISP to configure secondary DNS services for you. Alternatively, you can put your public DNS records on a dedicated, external DNS service while hosting your private DNS records entirely on your internal DNS servers.






Because secondary servers use forward lookup zones for most types of queries, you might not need reverse lookup zones. But reverse lookup zone files are essential for primary servers, and you must configure them for proper domain name resolution.

If you want to set up your own secondaries for backup services and load balancing, follow these steps:

1. Start the DNS Manager console. If the server you want to configure isn't listed, connect to it as described previously.
2. Right-click the server entry, and then click New Zone to start the New Zone Wizard. Click Next.
3. For Zone Type, select Secondary Zone. Click Next.
4. Secondary servers can use both forward and reverse lookup zone files. You create the forward lookup zone first, so select Forward Lookup Zone, and then click Next.
5. As you are creating a copy of an existing zone, click Browse when prompted to enter the zone name and then use the dialog box provided to select the zone you are copying on the source server. Note the name of the source server for the next step. When you are ready to continue, click Next.



- Click in the Master Servers list, type name or IP address of the primary server for the zone, and then press Enter. The wizard then attempts to validate the server. If an error occurs, be sure the server is connected to the network and that you've entered the correct IP address. Also ensure you've enabled zone transfers on the primary. If you want to copy zone data from other servers in case the first server isn't available, repeat this step.





7. Click Next, and then click Finish. On a busy or large network, you might need to configure reverse lookup zones on secondaries. If so, follow the steps listed in the next section.





## Configuring Reverse Lookups

Forward lookups are used to resolve domain names to IP addresses. Reverse lookups are used to resolve IP addresses to domain names. Each segment on your network should have a reverse lookup zone. For example, if you have the subnets 192.168.10.0, 192.168.11.0, and 192.168.12.0, you should have three reverse lookup zones.

The standard naming convention for reverse lookup zones is to enter the network ID in reverse order and then use the suffix `in-addr.arpa`. With the previous example, you'd have reverse lookup zones named `10.168.192.in-addr.arpa`, `11.168.192.in-addr.arpa`, and `12.168.192.in-addr.arpa`. Records in the reverse lookup zone must be in sync with the forward lookup zone. If the zones get out of sync, authentication might fail for the domain.

You create reverse lookup zones by following these steps:


1. Start the DNS Manager console. If the server you want to configure isn't listed, connect to it as described previously.
2. Right-click the server entry, and then click New Zone to start the New Zone Wizard. Click Next.
3. If you're configuring a primary server integrated with Active Directory (a domain controller), select Primary Zone and be sure that Store The Zone In Active Directory is selected. If you don't want to integrate DNS with Active Directory, select Primary Zone, and then clear the Store The Zone In Active Directory check box. Click Next.





4. If you're configuring a reverse lookup zone for a secondary server, select Secondary Zone, and then click Next.
5. If you're integrating the zone with Active Directory, choose one of the following replication strategies:
  - **To All DNS Servers Running On Domain Controllers In This Forest** Choose this strategy if you want the widest replication strategy. Remember, the Active Directory forest includes all domain trees that share the directory data with the current domain.
  - **To All DNS Servers Running On Domain Controllers In This Domain** Choose this strategy if you want to replicate DNS information within the current domain.
  - **To All Domain Controllers In This Domain (For Windows 2000 Compatibility)** Choose this strategy if you want to replicate DNS information to all domain controllers within the current domain, as needed for Windows 2000 compatibility. Although this strategy gives wider replication for DNS information within the domain, not every domain controller is a DNS server as well (and you don't need to configure every domain controller as a DNS server either).
6. Select Reverse Lookup Zone, and then click Next.
7. Choose whether you want to create a reverse lookup zone for IPv4 or IPv6 addresses, and then click Next. Do one of the following:

- If you are configuring a reverse lookup zone for IPv4, enter the network ID for the reverse lookup zone. The values you enter set the default name for the reverse lookup zone. Click Next.



The screenshot shows the 'New Zone Wizard' window. The title bar says 'New Zone Wizard' with a close button. The main heading is 'Reverse Lookup Zone Name' with a subtext: 'A reverse lookup zone translates IP addresses into DNS names.' Below this, it says 'To identify the reverse lookup zone, type the network ID or the name of the zone.' There are two radio buttons: 'Network ID:' (selected) and 'Reverse lookup zone name:'. The 'Network ID:' field contains '192 .168 .10 .'. Below this field, there is explanatory text: 'The network ID is the portion of the IP addresses that belongs to this zone. Enter the network ID in its normal (not reversed) order.' and 'If you use a zero in the network ID, it will appear in the zone name. For example, network ID 10 would create zone 10.in-addr.arpa, and network ID 10.0 would create zone 0.10.in-addr.arpa.' The 'Reverse lookup zone name:' field contains '10.168.192.in-addr.arpa'.

**TIP** If you have multiple subnets on the same network, such as 192.168.10 and 192.168.11, you can enter only the network portion for the zone name. For example, you could enter 192.168. In this case, you'd have 168.192.in-addr.arpa as the zone name and allow the DNS Manager console to create the necessary subnet zones when needed.

- If you are configuring a reverse lookup zone for IPv6, Enter the IPv6 prefix for the addresses and networks that are part of the reverse lookup zone. Be sure to use full prefix notation, such as FE80::/64. The values you enter are used to automatically generate the related zone names. Depending on the prefix you enter, the New Zone wizard will create up to eight zones. Click Next.

**NOTE** Although DHCP for Windows Server will only work with address ranges and networks that use the /64 network prefix, DNS for Windows Server doesn't have these limitations. You can use subnetting as appropriate.

The screenshot shows the 'New Zone Wizard' dialog box with the title bar 'New Zone Wizard' and a close button. The main heading is 'Reverse Lookup Zone Name' with a subtext: 'A reverse lookup zone translates IP addresses into DNS names.' To the right is a small server icon. The instructions state: 'To name the reverse lookup zone, enter an IPv6 address prefix to auto generate the zone name(s). Depending on the prefix you enter, up to 8 zones may be created.' There is a text box labeled 'IPv6 Address Prefix:' containing 'FE80::/64'. Below it is a section 'Reverse Lookup Zones' with a list box containing one entry: a checked checkbox followed by '0.0.0.0.0.0.0.0.0.0.0.0.8.e.f.ip6.arpa'. At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

8. If you're configuring a primary or secondary server that isn't integrated with Active Directory, you need to set the zone file name. A default name for the zone's DNS database file should be filled in for you. You can use this name or enter a new file name. Click Next.
9. Specify whether dynamic updates are allowed. You have three options:
  - **Allow Only Secure Dynamic Updates** When the zone is integrated with Active Directory, you can use ACLs to restrict which clients can perform dynamic updates. With this option selected, only clients with authorized computer accounts and approved ACLs can dynamically update their resource records in DNS when changes occur.
  - **Allow Both Nonsecure And Secure Dynamic Updates** Choose this option to allow any client to update its resource records in DNS when changes occur. Clients can be secure or nonsecure.
  - **Do Not Allow Dynamic Updates** Choosing this option disables dynamic updates in DNS. You should use this option only when the zone isn't integrated with Active Directory.
10. Click Next, and then click Finish. The new zone is added to the server, and basic DNS records are created automatically.




After you set up the reverse lookup zones, you need to ensure that delegation for the zones is handled properly. Contact your networking team or your ISP to ensure that the zones are registered with the parent domain.

## Configuring Global Names

The GlobalNames zone is a specially named forward lookup zone that should be integrated with AD DS. Deploying a GlobalNames zone creates static, global records with single-label names, without relying on WINS. This enables users to access hosts by using single-label names rather than fully qualified domain names. You should use the GlobalNames zone when name resolution depends on DNS, such as when your organization is no longer using WINS and you are planning to deploy only IPv6. Because dynamic updates cannot be used to register updates in the GlobalNames zone, you should configure single-label name resolution only for your primary servers.

You can deploy a GlobalNames zone by completing the following steps:

1. In the DNS Manager console, select a DNS server that is also a domain controller. If the server you want to configure isn't listed, connect to it as described previously in "Installing and Configuring the DNS Server Service."
2. Right-click the Forward Lookup Zones node, and then click New Zone. In the New Zone Wizard, click Next to accept the defaults to create a primary zone integrated with AD DS. On the Active Directory Zone Replication Scope page, choose to replicate the zone throughout the forest, and then click Next. On the Zone Name page, enter **GlobalNames** as the zone name. Click Next twice, and then click Finish.



3. On every authoritative DNS server in the forest now and in the future, you need to enter the following at an administrator shell prompt:

```
Set-DnsServerGlobalNameZone -ComputerName ServerName -Enable $True
```





ServerName is the name of the DNS server that hosts the GlobalNames zone. To specify the local computer, simply omit the –ComputerName parameter, such as `Set-DnsServerGlobalNameZone –Enable $True`

4. For each server that you want users to be able to access by using a single-label name, add an alias (CNAME) record to the GlobalNames zone. In the DNS Manager console, right-click the GlobalNames node, select New Alias (CNAME), and then use the dialog box provided to create the new resource record.

**NOTE** An authoritative DNS server tries to resolve queries in the following order: by using local zone data, by using the GlobalNames zone, by using DNS suffixes, by using WINS. For dynamic updates, an authoritative DNS server checks the GlobalNames zone before checking the local zone data.

**TIP** If you want DNS clients in another forest to use the GlobalNames zone for resolving names, you need to add an SRV resource record with the service name `_globalnames._msdcs` to that forest's forestwide DNS partition. The record must specify the FQDN of the DNS server that hosts the GlobalNames zone.







## Chapter 35. Managing and Maintaining DNS

After installing a DNS server, you'll need to create records, manage zone properties, the Start Of Authority (SOA) Record, and other options to ensure smooth operations. You may also need to review and modify the server's configuration and security settings. Related tasks are discussed in this chapter.






# Managing DNS Servers

The DNS Manager console is the tool you use to manage local and remote DNS servers. The DNS Manager console's main window is divided into two panes. The left pane makes it possible for you to access DNS servers and their zones. The right pane shows the details for the currently selected item. You can work with the DNS Manager console in three ways:

- Double-click an entry in the left pane to expand the list of files for the entry.
- Select an entry in the left pane to display details such as zone status and domain records in the right pane.
- Right-click an entry to display a context menu.



The Forward Lookup Zones and Reverse Lookup Zones folders provide access to the domains and subnets configured for use on this server. When you select domain or subnet folders in the left pane, you can manage DNS records for the domain or subnet.

## Adding and Removing Servers to Manage

You can use the DNS Manager console to manage servers running DNS by following these steps:

1. Right-click DNS in the console tree, and then click **Connect To DNS Server**.
2. If you're trying to connect to the local computer, select **This Computer**. Otherwise, select **The Following Computer**, and then enter the IP address or fully qualified host name of the remote computer with which you want to connect.
3. Click **OK**. Windows Server 2016 attempts to contact the server. If it makes contact, it adds the server to the console.

**NOTE** If a server is offline or otherwise inaccessible because of security restrictions or problems with the Remote Procedure Call (RPC) service, the






connection fails. You can still add the server to the console by clicking Yes when prompted.

In the DNS Manager console, you can delete a server by selecting its entry and then pressing Delete. When prompted, click Yes to confirm the deletion. Deleting a server only removes it from the server list in the console tree. It doesn't actually delete the server.



## Starting and Stopping a DNS Server

To manage DNS servers, you use the DNS Server service. You can start, stop, pause, resume, and restart the DNS Server service in the Services node of Server Manager or from the command line. You can also manage the DNS Server service in the DNS Manager console. Right-click the server you want to manage in the DNS Manager console, point to All Tasks, and then click Start, Stop, Pause, Resume, or Restart as appropriate.



**NOTE** In Server Manager, under the DNS Server node, expand the DNS node and then right-click the server with which you want to work. On the shortcut menu, select Start Service, Stop Service, Pause Service, Resume Service, or Restart Service as appropriate.

## Using DNSSEC and Signing Zones

All current versions of Windows support DNS Security Extensions (DNSSEC). DNSSEC is defined in several Request For Comments (RFCs), including RFCs 4033, 4034, and 4035. These RFCs add origin authority, data integrity, and authenticated denial of existence to DNS. With DNSSEC, there are the following additional resource records to learn about:

- [DNSKEY \(Domain Name System Key\)](#)
- [RRSIG \(Resource Record Signature\)](#)
- [NSEC \(NextSECure\)](#)
- [DS \(Domain Services\)](#)

The DNS client running on these operating systems can send queries that indicate support for DNSSEC, process related records, and determine whether a DNS server has validated records on its behalf. On Windows servers, DNSSEC allows your DNS servers to securely sign zones, to host DNSSEC-signed zones, to process related records, and to perform both validation and authentication. The way a DNS client



works with DNSSEC is configured through the Name Resolution Policy Table (NRPT), which stores settings that define the DNS client's behavior. Typically, you manage the NRPT through Group Policy.

When a DNS server hosting a signed zone receives a query, the server returns the digital signatures in addition to the requested records. A resolver or another server configured with a trust anchor for a signed zone or for a parent of a signed zone can obtain the public key of the public/private key pair and validate that the responses are authentic and have not been tampered with.

As part of your predeployment planning, you need to identify the DNS zones to secure with digital signatures. DNS Server for Windows Server 2016 has the following significant enhancements for DNSSEC:

- Support for dynamic updates in Active Directory–integrated zones. Previously, if an Active Directory domain zone was signed, you needed to manually update all SRV records and other resource records. This is no longer required because DNS Server now does this automatically.
- Support for online signing, automated key management, and automated trust anchor distribution. Previously, you needed to configure and manage signings, keys, and trust anchors. This is no longer required because DNS Server now does this automatically.
- Support for validations of records signed with updated DNSSEC standards including NSEC3 and RSA/SHA-2.

With Windows Server 2016, an authoritative DNS server also can act as the Key Master for DNSSEC. The Key Master generates and manages signing keys for both Active Directory-integrated zones protected by DNSSEC and standard (file-backed) zones protected by DNSSEC. When a zone has a designated Key Master, the Key Master is responsible for the entire key signing process from key generation to storage, rollover, retirement, and deletion.

Although key signing and management tasks can only be initiated from the Key Master, other primary DNS servers can continue to use zone signing—they just do so via the Key Master. You must choose a key master when you sign a zone with DNSSEC. You can transfer the key master role to another DNS server that hosts the zone at any time.

Additionally, keep the following in mind:

- For file-backed zones, the primary server and all secondary servers hosting the zone must be a Windows Server or a DNSSEC-aware server that is running an operating system other than Windows.

- For Active Directory–integrated zones, every domain controller that is a DNS server in the domain must be running Windows Server if the signed zone is set to replicate to all DNS servers in the domain. Every domain controller that is a DNS server in the forest must be running Windows Server if the signed zone is set to replicate to all DNS servers in the forest.
- For mixed environments, all servers that are authoritative for a DNSSEC-signed zone must be DNSSEC-aware servers. DNSSEC-aware Windows clients that request DNSSEC data and validation must be configured to issue DNS queries to a DNSSEC-aware server. Non-DNSSEC-aware Windows clients can be configured to issue DNS queries to DNSSEC-aware servers. DNSSEC-aware servers can be configured to recursively send queries to a non-DNSSEC-aware DNS server.

Securing DNS zones with digital signatures is a multistep process. As part of that process, you need to designate a key master . Any authoritative server that hosts a primary copy of a zone can act as the key master. Next, you need to generate a Key Signing Key and a Zone Signing Key. A Key Signing Key (KSK) that is an authentication key has a private key and a public key associated with it. The private key is used for signing all of the DNSKEY records at the root of the zone. The public key is used as a trust anchor for validating DNS responses. A Zone Signing Key (ZSK) is used for signing zone records.

After you generate keys, you create resource records for authenticated denial of existence by using either the more secure NSEC3 standard or the less secure NSEC standard. Because trust anchors are used to validate DNS responses, you also need to specify how trust anchors are updated and distributed. Typically, you'll want to automatically update and distribute trust anchors. By default, records are signed with SHA-1 and SHA-256 encryption. You can select other encryption algorithms as well.

You don't need to go through the configuration process each time you sign a zone. The signing keys and other signing parameters are available for reuse.


To sign a zone while customizing the signing parameters, follow these steps:

1. In the DNS Manager console, right-click the zone you want to secure. On the shortcut menu, select DNSSEC, and then select Sign The Zone. This starts the Zone Signing Wizard. If the wizard displays a welcome page, read the Welcome text, and then click Next.
2. On the Signing Options page, select Customize Zone Signing Parameters, and then click Next.





3. Select a key master for the zone. Any authoritative server that hosts a primary copy of a zone can act as the key master. When you are ready to continue, click Next twice.



4. On the Key Signing Key page, define up to three Key Signing Keys. Configure a KSK by clicking Add, accepting or changing the default values for key properties and rollover, and then clicking OK. When you are ready to continue, click Next twice.

**Zone Signing Wizard**

**Key Signing Key (KSK)**  
Configure one or more KSKs

Configure parameters for at least one KSK. A maximum of three KSKs can be specified for each of the available cryptographic algorithms.

| Algorithm   | Key length | KSP                | Replication | Rollover state | Initial rollover o... |
|-------------|------------|--------------------|-------------|----------------|-----------------------|
| RSA/SHA-256 | 2048       | Microsoft Softw... | Enabled     | Enabled        | 0                     |
| RSA/SHA-512 | 4096       | Microsoft Softw... | Enabled     | Enabled        | 0                     |

< >

Add Edit Remove

**New Key Signing Key (KSK)**

Guid  
Guid: {00000000-0000-0000-0000-000000000000}

Key Generation

☒ Generate new signing keys.  
☐ Use pre-generated keys

Use this key as active key:   
Use this key as standby key:

Key Properties


Cryptographic algorithm: RSA/SHA-256  
Key length (Bits): 2048  
Select a key storage provider to generate and store keys: Microsoft Software Key Storage Prov  
DNSKEY RRSET signature validity period (hours): 168  
☒ Replicate this private key to all DNS servers authoritative for this zone.  
(Applicable only to AD integrated zones)

Key Rollover

☒ Enable automatic rollover  
Rollover frequency (days): 755  
Delay the first rollover by (days): 0

OK Cancel

- On the Zone Signing Key page, define up to three Zone Signing Keys. Configure a ZSK by clicking Add, accepting or changing the default values for key properties and rollover, and then clicking OK. When you are ready to continue, click Next five times.
- After the wizard signs the zone, click Finish.




To sign a zone and use existing signing parameters, follow these steps:

1. In the DNS Manager console, right-click the zone you want to secure. On the shortcut menu, select DNSSEC and then select Sign The Zone. This starts the Zone Signing Wizard. If the wizard displays a welcome page, read the Welcome text, and then click Next.
2. On the Signing Options page, select Sign The Zone With Parameters Of An Existing Zone. Enter the name of an existing signed zone, such as **imaginedlands.local** , and then click Next.



3. On the Key Master page, select a key master for the zone. Any authoritative server that hosts a primary copy of a zone can act as the key master. Click Next twice.




4. After the wizard signs the zone, click Finish.

## Creating Child Domains Within Zones

By using the DNS Manager console, you can create child domains within a zone. For example, if you create the primary zone microsoft.com, you could create the subdomains hr.microsoft.com and mis.microsoft.com for the zone. You create child domains by following these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server with which you want to work.
2. Right-click the parent domain entry, and then click New Domain.

3. Enter the name of the new domain, and then click OK. For tech.imaginedlands.local, you would enter **tech** . For mis.imaginedlands.local, you would enter **mis** .



A screenshot of a Windows-style dialog box titled "New DNS Domain". The dialog box has a close button (X) in the top right corner. Inside the dialog, there is a label "Type the new DNS domain name:" followed by a text input field. The input field contains the text "tech". At the bottom of the dialog, there are two buttons: "OK" and "Cancel". The "OK" button is highlighted with a blue dashed border.

# Creating Child Domains in Separate Zones


As your organization grows, you might want to organize the DNS namespace into separate zones. At your corporate headquarters, you could have a zone for the parent domain microsoft.com. At branch offices, you could have zones for each office, such as memphis.microsoft.com, newyork.microsoft.com, and la.microsoft.com.

You create child domains in separate zones by following these steps:

1. Install a DNS server in each child domain, and then create the necessary forward and reverse lookup zones for the child domain as described earlier in “Installing DNS Servers.”
2. On the authoritative DNS server for the parent domain, you delegate authority to each child domain. Delegating authority enables the child domain to resolve and respond to DNS queries from computers inside and outside the local subnet.

You delegate authority to a child domain by following these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server with which you want to work.
2. Right-click the parent domain entry, and then click New Delegation to start the New Delegation Wizard. Click Next.
3. Enter the name of the delegated domain, such as **resources** , and then click Next. The name you enter updates the value in the Fully Qualified Domain Name text box.



New Delegation Wizard

**Delegated Domain Name**  
Authority for the DNS domain you supply will be delegated to a different zone.

Specify the name of the DNS domain you want to delegate.

Delegated domain:  
resources

Fully qualified domain name (FQDN):  
resources.imaginedlands.local

4. Click Add. This displays the New Name Server Record dialog box.
5. In the Server Fully Qualified Domain Name text box, type the fully qualified host name of a DNS server for the child domain, such as



**corpserver01.tn.tvpress.local** , and then click Resolve. The server then performs a lookup query and adds the resolved IP address to the IP Address list.

6. Repeat step 5 to specify additional name servers. The order of the entries determines which IP address is used first. Change the order as necessary by using the Up and Down buttons. When you are ready to continue, click OK to close the New Name Server Record dialog box.
7. Click Next, and then click Finish.





## Deleting a Domain or Subnet

Deleting a domain or subnet permanently removes it from the DNS server. To delete a domain or subnet, follow these steps:

1. In the DNS Manager console, right-click the domain or subnet entry.
2. On the shortcut menu, click Delete, and then confirm the action by clicking Yes.
3. If the domain or subnet is integrated with Active Directory, you'll receive a warning prompt. Confirm that you want to delete the domain or subnet from Active Directory by clicking Yes.

**NOTE** Deleting a domain or subnet deletes all DNS records in a zone file but doesn't actually delete the zone file on a primary or secondary server that isn't integrated with Active Directory. The actual zone file remains in the %SystemRoot%\System32\Dns directory. You can delete this file after you have deleted the zones from the DNS Manager console.






# Managing DNS Records

After you create the necessary zone files, you can add records to the zones. Computers that need to be accessed from Active Directory and DNS domains must have DNS records. Although there are many types of DNS records, most of these record types aren't commonly used. So rather than focus on record types you probably won't use, let's focus on the ones you will use:

- **A (IPv4 address)** Maps a host name to an IPv4 address. When a computer has multiple adapter cards, IPv4 addresses, or both, it should have multiple address records.
- **AAAA (IPv6 address)** Maps a host name to an IPv6 address. When a computer has multiple adapter cards, IPv6 addresses, or both, it should have multiple address records.
- **CNAME (canonical name)** Sets an alias for a host name. For example, by using this record, zeta.microsoft.com can have an alias of www.microsoft.com.
- **MX (mail exchanger)** Specifies a mail exchange server for the domain, which enables email messages to be delivered to the correct mail servers in the domain.
- **NS (name server)** Specifies a name server for the domain, which enables DNS lookups within various zones. Each primary and secondary name server should be declared through this record.
- **PTR (pointer)** Creates a pointer that maps an IP address to a host name for reverse lookups.
- **SOA (start of authority)** Declares the host that's the most authoritative for the zone and, as such, is the best source of DNS information for the zone. Each zone file must have an SOA record (which is created automatically when you add a zone). Also declares other information about the zone, such as the responsible person, refresh interval, retry interval, and so on.
- **SRV (service location)** Locates a server providing a specific service. Active Directory uses SRV records to locate domain controllers, global catalog servers, LDAP servers, and Kerberos servers. Most SRV records are created automatically. For example, Active Directory creates an SRV record when you promote a domain controller. LDAP servers can add an SRV record to indicate they are available to handle LDAP requests in a particular zone.

## Adding Address and Pointer Records

You use the A and AAAA records to map a host name to an IP address, and the PTR record creates a pointer to the host for reverse lookups. You can create address and pointer records at the same time or separately.



The screenshot shows the 'New Host' dialog box with the following fields and options:

- Name (uses parent domain name if blank):** pencil
- Fully qualified domain name (FQDN):** pencil.tech.imaginedlands.local.
- IP address:** 192.168.10.58
- ☒ Create associated pointer (PTR) record
- ☐ Allow any authenticated user to update DNS records with the same owner name
- Buttons:** Add Host, Cancel

You create a new host entry with address and pointer records by following these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server with which you want to work.
2. Right-click the domain you want to update, and then click New Host (A Or AAAA). This opens the New Host dialog box.
3. Enter the single-part computer name, such as **pencil** , and then the IP address, such as **192.168.10.58** .
4. Select the Create Associated Pointer (PTR) Record check box.


**NOTE** You can create PTR records only if the corresponding reverse lookup zone is available. You can create this file by following the steps listed in “Configuring Reverse Lookups” in Chapter 34. The Allow Any Authenticated User option is available only when a DNS server is configured on a domain controller.

5. Click Add Host, and then click OK. Repeat these steps as necessary to add other hosts.
6. Click Done when you have finished.

## Adding a PTR Record Later

If you need to add a PTR record later, you can do so by following these steps:

1. In the DNS Manager console, expand the Reverse Lookup Zones folder for the server with which you want to work.
2. Right-click the subnet you want to update, and then click New Pointer (PTR).



New Resource Record

Pointer (PTR)

Host IP Address:  
192.168.10.58

Fully qualified domain name (FQDN):  
58.10.168.192.in-addr.arpa

Host name:  
lnpc14 Browse...

☐ Allow any authenticated user to update all DNS records with the same name. This setting applies only to DNS records for a new name.

3. Complete the host IP address. For example, if you are configuring the host IP address 192.168.10.42 for the 192.168.10 network, type **42**. Then enter the host name, such as **lnpc14**. Click OK.

## Adding DNS Aliases with CNAME

You specify host aliases by using CNAME records. Aliases enable a single host computer to appear to be multiple host computers. For example, the host matrix.imaginedlands.com can be made to appear as www.imaginedlands.com and ftp.imaginedlands.com.

New Resource Record

Alias (CNAME)

Alias name (uses parent domain if left blank):  
intranet

Fully qualified domain name (FQDN):  
intranet.tech.imaginedlands.local.

Fully qualified domain name (FQDN) for target host:  
fileserver82b4.imaginedlands.local Browse...

☐ Allow any authenticated user to update all DNS records with the same name. This setting applies only to DNS records for a new name.

To create a CNAME record, follow these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server with which you want to work.
2. Right-click the domain you want to update, and then click New Alias (CNAME).
3. In the Alias Name text box, enter the alias. The alias is a single-part host name, such as www or ftp .
4. In the Fully Qualified Domain Name (FQDN) For Target Host text box, enter the full host name of the computer for which the alias is to be used.
5. Click OK.

## Adding Mail Exchange Servers

MX records identify mail exchange servers for the domain. These servers are responsible for processing or forwarding email within the domain. When you create an MX record, you must specify a preference number for the mail server. A preference number is a value from 0 to 65,535 that denotes the mail server's priority within the domain. The mail server with the lowest preference number has the highest priority and is the first to receive mail. If mail delivery fails, the mail server with the next lowest preference number is tried.

New Resource Record

Mail Exchanger (MX)

Host or child domain:

By default, DNS uses the parent domain name when creating a Mail Exchange record. You can specify a host or child name, but in most deployments, the above field is left blank.

Fully qualified domain name (FQDN):

imaginedlands.local.

Fully qualified domain name (FQDN) of mail server:

mailer1.imagedlands.local Browse...

Mail server priority:

10

You create an MX record by following these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server with which you want to work.
  2. Right-click the domain you want to update, and then click New Mail Exchanger (MX).
  3. You can now create a record for the mail server by filling in these text boxes:
- **Host Or Child Domain** Using a single-part name, enter the name of the subdomain for which the server specified in this record is responsible. In most cases, you will leave this box blank, which specifies that there is no subdomain and the server is responsible for the domain in which this record is created.
  - **Fully Qualified Domain Name (FQDN)** Enter the FQDN of the domain to which this mail exchange record should apply, such as **imaginedlands.com**.
  - **Fully Qualified Domain Name (FQDN) Of Mail Server** Enter the FQDN of the mail server that should handle mail receipt and delivery, such as **corpmail.imagedlands.com**. Email for the previously specified domain is routed to this mail server for delivery.
  - **Mail Server Priority** Enter a preference number for the host from 0 to 65,535.

**NOTE** Assign preference numbers that leave room for growth. For example, use 10 for your highest priority mail server, 20 for the next, and 30 for the one after that.

**REAL WORLD** You can't enter a multipart name in the Host Or Child Domain text box. If you need to enter a multipart name, you are creating the MX record at the



wrong level of the DNS hierarchy. Create or access the additional domain level, and then add an MX record at this level for the subdomain.

4. Click OK.




# Adding Name Servers

NS records specify the name servers for the domain. Each primary and secondary name server should be declared through this record. If you obtain secondary name services from an ISP, be sure to insert the appropriate NS records.

You create an NS record by following these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server with which you want to work.
2. Display the DNS records for the domain by selecting the domain folder in the tree view.
3. Right-click an existing NS record in the view pane, and then click Properties. This opens the Properties dialog box for the domain with the Name Servers tab selected.



4. Click Add. This displays the New Name Server Record dialog box.
5. In the Server Fully Qualified Domain Name text box, enter the name of a DNS server for the child domain, such as **corpserver01.imaginedlands.local** and then click Resolve. The server then performs a lookup query and adds the resolved IP address to the IP Address list.
6. Repeat step 5 to specify additional name servers. The order of the entries determines which IP address is used first. Change the order as necessary by



using the Up and Down buttons. When you are ready to continue, click OK to close the New Name Server Record dialog box.

7. Click OK to save your changes.







# Viewing and Updating DNS Records

To view or update DNS records, follow these steps:

1. Double-click the zone with which you want to work. Records for the zone should be displayed in the right pane.
2. Double-click the DNS record you want to view or update. This opens the record's Properties dialog box. Make the necessary changes, and then click OK.





# Updating Zone Properties and the SOA Record

Each zone has separate properties you can configure. These properties set general zone parameters by using the SOA record, change notification, and WINS integration. In the DNS Manager console, you set zone properties by doing one of the following:

- [Right-click the zone you want to update, and then click Properties.](#)
- [Select the zone, and then click Properties on the Action menu.](#)

The Properties dialog boxes for forward and reverse lookup zones are identical except for the WINS and WINS-R tabs. In forward lookup zones, you use the WINS tab to configure lookups for NetBIOS computer names. In reverse lookup zones, you use the WINS-R tab to configure reverse lookups for NetBIOS computer names.



# Modifying the SOA Record

An SOA record designates the authoritative name server for a zone and sets general zone properties, such as retry and refresh intervals. You can modify this information by following these steps:

1. In the DNS Manager console, right-click the zone you want to update, and then click Properties.
2. Click the Start Of Authority (SOA) tab, and then update the settings as appropriate.

The screenshot shows the 'imaginedlands.local Properties' dialog box with the 'Start of Authority (SOA)' tab selected. The dialog has three main tabs: 'WINS', 'Zone Transfers', and 'Security'. Under 'Zone Transfers', there are sub-tabs: 'General', 'Start of Authority (SOA)', and 'Name Servers'. The 'Start of Authority (SOA)' tab is active, showing the following fields and controls:

- Serial number:** A text box containing '56494994' and an 'Increment' button.
- Primary server:** A text box containing 'corpserver12.imaginedlands.local.' and a 'Browse...' button.
- Responsible person:** A text box containing 'hostmaster.imaginedlands.local.' and a 'Browse...' button.
- Refresh interval:** A numeric box with '15' and a dropdown menu set to 'minutes'.
- Retry interval:** A numeric box with '10' and a dropdown menu set to 'minutes'.
- Expires after:** A numeric box with '1' and a dropdown menu set to 'days'.
- Minimum (default) TTL:** A numeric box with '1' and a dropdown menu set to 'hours'.
- TTL for this record:** A text box showing '0 :1 :0 :0' with the format '(DDDD:HH.MM.SS)' to its right.

At the bottom of the dialog are four buttons: 'OK', 'Cancel', 'Apply', and 'Help'.

You use the text boxes on the Start Of Authority (SOA) tab as follows:

- **Serial Number** A serial number that indicates the version of the DNS database files. The number is updated automatically whenever you make changes to zone files. You can also update the number manually. Secondary servers use this number to determine whether the zone's DNS records have changed. If the primary server's serial number is larger than the secondary server's serial number, the records have changed, and the secondary server can request the DNS records for the zone. You



can also configure DNS to notify secondary servers of changes (which might speed up the update process).

- **Primary Server** The FQDN for the name server followed by a period. The period is used to terminate the name and ensure that the domain information isn't appended to the entry.
- **Responsible Person** The email address of the person in charge of the domain. The default entry is hostmaster followed by a period followed by your domain name, meaning hostmaster@your\_domain.com. If you change this entry, substitute a period in place of the @ symbol in the email address and terminate the address with a period.
- **Refresh Interval** The interval at which a secondary server checks for zone updates. The default value is 15 minutes. You reduce network traffic by increasing this value. However, keep in mind that if the interval is set to 60 minutes, NS record changes might not be propagated to a secondary server for up to an hour.
- **Retry Interval** The time the secondary server waits after a failure to download the zone database. If the interval is set to 10 minutes and a zone database transfer fails, the secondary server waits 10 minutes before requesting the zone database once more.
- **Expires After** The period of time for which zone information is valid on the secondary server. If the secondary server can't download data from a primary server within this period, the secondary server lets the data in its cache expire and stops responding to DNS queries. Setting Expires After to seven days enables the data on a secondary server to be valid for seven days.
- **Minimum (Default) TTL** The minimum time-to-live (TTL) value for cached records on a secondary server. The value can be set in days, hours, minutes, or seconds. When this value is reached, the secondary server causes the associated record to expire and discards it. The next request for the record needs to be sent to the primary server for resolution. Set the minimum TTL to a relatively high value, such as 24 hours, to reduce traffic on the network and increase efficiency. Keep in mind that a higher value slows down the propagation of updates through the Internet.
- **TTL For This Record** The TTL value for this particular SOA record. The value is set in the format Days : Hours : Minutes : Seconds and generally should be the same as the minimum TTL for all records.

## Allowing and Restricting Zone Transfers

Zone transfers send a copy of zone information to other DNS servers. These servers can be in the same domain or in other domains. For security reasons, Windows Server 2016 disables zone transfers. To enable zone transfers for secondaries you've






configured internally or with ISPs, you need to permit zone transfers and then specify the types of servers to which zone transfers can be made.

Although you can allow zone transfers with any server, this opens the server to possible security problems. Instead of opening the floodgates, you should restrict access to zone information so that only servers you've identified can request updates from the zone's primary server. This enables you to funnel requests through a select group of secondary servers, such as your ISP's secondary name servers, and to hide the details of your internal network from the outside world.

To allow zone transfers and restrict access to the primary zone database, follow these steps:

1. In the DNS Manager console, right-click the domain or subnet you want to update, and then click Properties.
2. Click the Zone Transfers tab.






3. To restrict transfers to name servers listed on the Name Servers tab, select the Allow Zone Transfers check box, and then choose Only To Servers Listed On The Name Servers Tab.
4. To restrict transfers to designated servers, select the Allow Zone Transfers check box and then choose Only To The Following Servers. Then click Edit as appropriate to display the Allow Zone Transfers dialog box. Click in the IP Address list, enter the IP address of the secondary server for the zone, and then press Enter. Windows then attempts to validate the server. If an error occurs, make sure the server is connected to the network and that you've entered the correct IP address. If you want to copy zone data from other servers in case the first server isn't available, you can add IP addresses for other servers as well. Click OK.
5. Click OK to save your changes.

## Notifying Secondaries of Changes

You set properties for a zone with its SOA record. These properties control how DNS information is propagated on the network. You can also specify that the primary server

should notify secondary name servers when changes are made to the zone database. To do this, follow these steps:

1. In the DNS Manager console, right-click the domain or subnet you want to update, and then click Properties.
2. On the Zone Transfers tab, click Notify. This displays the Notify dialog box.




3. To notify secondary servers listed on the Name Servers tab, select the Automatically Notify check box, and then choose Servers Listed On The Name Servers Tab.
4. If you want to designate specific servers to notify, select the Automatically Notify check box, and then choose The Following Servers. Click in the IP Address list, enter the IP address of the secondary server for the zone, and then press Enter. Windows then attempts to validate the server. If an error occurs, make sure the server is connected to the network and that you entered the correct IP address. If you want to notify other servers, add IP addresses for those servers as well.
5. Click OK twice.

## Setting the Zone Type

When you create zones, they're designated as having a specific zone type and an Active Directory integration mode. You can change the type and integration mode at any time by following these steps:

1. In the DNS Manager console, right-click the domain or subnet you want to update, and then click Properties.
2. Under Type on the General tab, click Change. In the Change Zone Type dialog box, select the new type for the zone.



3. To integrate the zone with Active Directory, select the Store The Zone In Active Directory check box.
4. To remove the zone from Active Directory, clear the Store The Zone In Active Directory check box.
5. Click OK twice.

## Enabling and Disabling Dynamic Updates

Dynamic updates enable DNS clients to register and maintain their own address and pointer records. This is useful for computers dynamically configured through DHCP. By enabling dynamic updates, you make it easier for dynamically configured computers to locate one another on the network. When a zone is integrated with Active Directory, you have the option of requiring secure updates. With secure updates, you use ACLs to control which computers and users can dynamically update DNS.

You can enable and disable dynamic updates by following these steps:

1. In the DNS Manager console, right-click the domain or subnet you want to update, and then click Properties.
2. Use the following options in the Dynamic Updates list on the General tab to enable or disable dynamic updates:





- **None** Disable dynamic updates.
- **Nonsecure And Secure** Enable nonsecure and secure dynamic updates.
- **Secure Only** Enable dynamic updates with Active Directory security. This is available only with Active Directory integration.

3. Click OK.

NOTE DNS integration settings must also be configured for DHCP. See “Integrating DHCP and DNS” in Chapter 32.

## Managing DNS Server Configuration and Security

You use the Server Properties dialog box to manage the general configuration of DNS servers. Through it, you can enable and disable IP addresses for the server and control access to DNS servers outside the organization. You can also configure monitoring, logging, and advanced options.





# Enabling and Disabling IP Addresses for a DNS Server

By default, multihomed DNS servers respond to DNS requests on all available network interfaces and the IP addresses they're configured to use.

Through the DNS Manager console, you can specify that the server can answer requests only on specific IP addresses. Generally, you'll want to ensure that a DNS server has at least one IPv4 interface and one IPv6 interface.

To specify which IP addresses are used for answering requests, follow these steps:

1. In the DNS Manager console, right-click the server you want to configure, and then click Properties.
2. On the Interfaces tab, select Only The Following IP Addresses. Select an IP address that should respond to DNS requests, or clear an IP address that should not respond to DNS requests. Only the selected IP addresses will be used for DNS. All other IP addresses on the server will be disabled for DNS.
3. Click OK.





# Controlling Access to DNS Servers Outside the Organization

Restricting access to zone information enables you to specify which internal and external servers can access the primary server. For external servers, this controls which servers can get in from the outside world. You can also control which DNS servers within your organization can access servers outside of your organization. To do this, you need to set up DNS forwarding within the domain.

With DNS forwarding, you configure DNS servers within the domain as one of the following:

- **Nonforwarders** Servers that must pass DNS queries they can't resolve to designated forwarding servers. These servers essentially act like DNS clients to their forwarding servers.
- **Forwarding-only** Servers that can only cache responses and pass requests to forwarders. These are also known as caching-only DNS servers.
- **Forwarders** Servers that receive requests from nonforwarders and forwarding-only servers. Forwarders use normal DNS communication methods to resolve queries and to send responses back to other DNS servers.
- **Conditional forwarders** Servers that forward requests based on the DNS domain. Conditional forwarding is useful if your organization has multiple internal domains.


**NOTE** You can't configure the root server for a domain for forwarding (except for conditional forwarding used with internal name resolution). You can configure all other servers for forwarding.

## Creating Nonforwarding and Forwarding-Only Servers


To create a nonforwarding or forwarding-only DNS server, follow these steps:

1. In the DNS Manager console, right-click the server you want to configure, and then click Properties.
2. Click the Advanced tab. To configure the server as a nonforwarder, ensure that the Disable Recursion check box is cleared, click OK, and then skip the remaining steps. To configure the server as a forwarding-only server, be sure that the Disable Recursion check box is selected.





3. On the Forwarders tab, click Edit. This displays the Edit Forwarders dialog box.
4. Click in the IP Address list, type the IP address of a forwarder for the network, and then press Enter. Windows then attempts to validate the server. If an error occurs, make sure the server is connected to the network and that you've entered the correct IP address. Repeat this process to specify the IP addresses of other forwarders.



5. Set the Forward Queries Time Out interval. This value controls how long the nonforwarder tries to query the current forwarder if it gets no response. When



the Forward Time Out interval passes, the nonforwarder tries the next forwarder on the list. The default is three seconds. Click OK.

## Creating Forwarding Servers

Any DNS server that isn't designated as a nonforwarder or a forwarding-only server will act as a forwarder. Thus, on the network's designated forwarders you should be sure that the Disable Recursion option is not selected and that you haven't configured the server to forward requests to other DNS servers in the domain.

## Configuring Conditional Forwarding

If you have multiple internal domains, you might want to consider configuring conditional forwarding, which enables you to direct requests for specific domains to specific DNS servers for resolution. Conditional forwarding is useful if your organization has multiple internal domains and you need to resolve requests between these domains.

To configure conditional forwarding, follow these steps:

1. In the DNS Manager console, select and then right-click the Conditional Forwarders folder for the server with which you want to work. Click New Conditional Forwarder on the shortcut menu.
2. In the New Conditional Forwarder dialog box, enter the name of a domain to which queries should be forwarded, such as **tvpress.local** .

New Conditional Forwarder

DNS Domain:

tvpress.local

IP addresses of the master servers:

| IP Address              | Server FQDN | Validated |
|-------------------------|-------------|-----------|
| <Click here to add a... |             |           |

Delete

Up

Down

☒ Store this conditional forwarder in Active Directory, and replicate it as follows:

All DNS servers in this forest

This will not replicate to DNS servers that are pre-Windows Server 2003 domain controllers

Number of seconds before forward queries time out: 5

The server FQDN will not be available if the appropriate reverse lookup zones and entries are not configured.

OK Cancel

3. Click in the IP Address list, type the IP address of an authoritative DNS server in the specified domain, and then press Enter. Repeat this process to specify additional IP addresses.
4. If you're integrating DNS with Active Directory, select the Store This Conditional Forwarder In Active Directory check box, and then choose one of the following replication strategies:
  - **All DNS Servers In This Forest** Choose this strategy if you want the widest replication strategy. Remember, the Active Directory forest includes all domain trees that share the directory data with the current domain.
  - **All DNS Servers In This Domain** Choose this strategy if you want to replicate forwarder information within the current domain and child domains of the current domain.
  - **All Domain Controllers In This Domain** Choose this strategy if you want to replicate forwarder information to all domain controllers within the current domain and child domains of the current domain. Although this strategy gives wider replication for forwarder information within the domain, not every domain controller is a DNS server as well (nor do you need to configure every domain controller as a DNS server).
1. Set the Forward Queries Time Out interval. This value controls how long the server tries to query the forwarder if it gets no response. When the Forward



Time Out interval passes, the server tries the next authoritative server on the list. The default is five seconds. Click OK.

2. Repeat this procedure to configure conditional forwarding for other domains.

## **Enabling and Disabling Event Logging**

By default, the DNS service tracks all events for DNS in the DNS Server event log. This log records all applicable DNS events and is accessible through the Event Viewer node in Computer Management. This means that all informational, warning, and error events are recorded. You can change the logging options by following these steps:


1. In the DNS Manager console, right-click the server you want to configure, and then click Properties.
2. Use the options on the Event Logging tab to configure DNS logging. To disable logging altogether, choose No Events.
3. Click OK.

## Using Debug Logging to Track DNS Activity

You typically use the DNS Server event log to track DNS activity on a server. This log records all applicable DNS events and is accessible through the Event Viewer node in Computer Management. If you're trying to troubleshoot DNS problems, it's sometimes useful to configure a temporary debug log to track certain types of DNS events. However, don't forget to clear these events after you finish debugging.

To configure debugging, follow these steps:

1. In the DNS Manager console, right-click the server you want to configure, and then click Properties.
2. On the Debug Logging tab, select the Log Packets For Debugging check box, and then select the check boxes for the events you want to track temporarily.



3. In the File Path And Name text box, enter the name of the log file, such as **dns.logs** . Logs are stored in the %SystemRoot%\System32\Dns directory by default.
4. Click OK. When finished debugging, turn off logging by clearing the Log Packets For Debugging check box.




# Monitoring a DNS Server

Windows Server 2016 has built-in functionality for monitoring a DNS server. Monitoring is useful to ensure that DNS resolution is configured properly.

You can configure monitoring to occur manually or automatically by following these steps:

1. In the DNS Manager console, right-click the server you want to configure, and then click Properties.
2. Click the Monitoring tab. You can perform two types of tests. To test DNS resolution on the current server, select the A Simple Query Against This DNS Server check box. To test DNS resolution in the domain, select the A Recursive Query To Other DNS Servers check box.



3. You can perform a manual test by clicking Test Now. You can schedule the server for automatic monitoring by selecting the Perform Automatic Testing At The Following Interval check box and then setting a time interval in seconds, minutes, or hours.

- 
4. The Test Results panel shows the results of testing. You'll receive a date and time stamp indicating when the test was performed and a result, such as Pass or Fail. Although a single failure might be the result of a temporary outage, multiple failures typically indicate a DNS resolution problem.

**NOTE** If all recursive query tests fail, the advanced server option Disable Recursion might be selected. Click the Advanced tab and check the server options.

**REAL WORLD** If you're actively troubleshooting a DNS problem, you might want to configure testing to occur every 10–15 seconds. This interval will provide a rapid succession of test results. If you're monitoring DNS for problems as part of your daily administrative duties, you'll want a longer time interval, such as two or three hours.



# About the Author



William R. Stanek (<http://www.williamrstanek.com>) has more than 20 years of hands-on experience with advanced programming and development. He is a leading technology expert, an award-winning author, and a pretty-darn-good instructional trainer. Over the years, his practical advice has helped millions of programmers, developers, and network engineers all over the world. In 2013, William celebrated the publication of his 150<sup>th</sup> book.


William has been involved in the commercial Internet community since 1991. His core business and technology experience comes from more than 11 years of military service.

He has substantial experience in developing server technology, encryption, and Internet solutions. He has written many technical white papers and training courses on a wide variety of topics. He frequently serves as a subject matter expert and consultant.

William has an MS with distinction in information systems and a BS in computer science, magna cum laude. He is proud to have served in the Persian Gulf War as a combat crewmember on an electronic warfare aircraft. He flew on numerous combat missions into Iraq and was awarded nine medals for his wartime service, including one of the United States of America's highest-flying honors, the Air Force Distinguished Flying Cross. Currently, he resides in the Pacific Northwest with his wife and children.

William recently rediscovered his love of the great outdoors. When he's not writing, he can be found hiking, biking, backpacking, traveling, or trekking in search of adventure with his family! In his spare time, William writes books for children, including *The Bugville Critters Explore the Solar System* and *The Bugville Critters Go on Vacation*.

Find William on Twitter at <http://www.twitter.com/WilliamStanek> and on Facebook at <http://www.facebook.com/William.Stanek.Author>.



Thank you for purchasing this book. If you found this book to be useful, helpful or informative, raise your voice and support William's work by sharing about this book online.

Unsure how to share? Here are some tips:

- [Blog about the book](#)
- [Write a review at your favorite online store](#)
- [Post about the book on Facebook or elsewhere](#)
- [Tweet about the book](#)

Stay in touch with William on Facebook and Twitter!

