


Microsoft®

# Windows Server® 2012

William R. Stanek  
*Author and Series Editor*



# Pocket Consultant

PUBLISHED BY  
Microsoft Press  
A Division of Microsoft Corporation  
One Microsoft Way  
Redmond, Washington 98052-6399

Copyright © 2012 by William R. Stanek

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2012944749  
ISBN: 978-0-7356-6633-7

Printed and bound in the United States of America.

First Printing

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Book Support at [mspinput@microsoft.com](mailto:mspinput@microsoft.com). Please tell us what you think of this book at <http://www.microsoft.com/learning/booksurvey>.

Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

The example companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the author, Microsoft Corporation, nor its resellers or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

**Acquisitions Editor:** Anne Hamilton

**Developmental Editor:** Karen Szall

**Project Editor:** Karen Szall

**Editorial Production:** Curtis Philips

**Technical Reviewer:** Bob Hogan; Technical Review services provided by Content Master, a member of CM Group, Ltd.

**Copyeditor:** Roger LeBlanc

**Indexer:** William P. Meyers

**Cover:** Twist Creative • Seattle

*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 Karen, Martin, Lucinda, Juliana, and many others who've helped out in ways both large and small.*

—WILLIAM R. STANEK



# Contents at a Glance

*Introduction*

xxv

## **PART I      WINDOWS SERVER 2012 ADMINISTRATION FUNDAMENTALS**

---

CHAPTER 1	Windows Server 2012 Administration Overview	3
CHAPTER 2	Managing Servers Running Windows Server 2012	31
CHAPTER 3	Monitoring Processes, Services, and Events	87
CHAPTER 4	Automating Administrative Tasks, Policies, and Procedures	135
CHAPTER 5	Enhancing Computer Security	191

## **PART II      WINDOWS SERVER 2012 DIRECTORY SERVICES ADMINISTRATION**

---

CHAPTER 6	Using Active Directory	217
CHAPTER 7	Core Active Directory Administration	249
CHAPTER 8	Creating User and Group Accounts	295
CHAPTER 9	Managing User and Group Accounts	347

## **PART III      WINDOWS SERVER 2012 DATA ADMINISTRATION**

---

CHAPTER 10	Managing File Systems and Drives	385
CHAPTER 11	Configuring Volumes and RAID Arrays	419
CHAPTER 12	Data Sharing, Security, and Auditing	457
CHAPTER 13	Data Backup and Recovery	519

## **PART IV      WINDOWS SERVER 2012 NETWORK ADMINISTRATION**

---

CHAPTER 14	Managing TCP/IP Networking	559
CHAPTER 15	Running DHCP Clients and Servers	573
CHAPTER 16	Optimizing DNS	615

*Index*

651



# Contents

*Introduction*

xxv

## **PART I    WINDOWS SERVER 2012 ADMINISTRATION FUNDAMENTALS**

---

<b>Chapter 1</b>	<b>Windows Server 2012 Administration Overview</b>	<b>3</b>
	Windows Server 2012 and Windows 8 . . . . .	4
	Getting to Know Windows Server 2012 . . . . .	6
	Power Management Options . . . . .	8
	Networking Tools and Protocols . . . . .	11
	Understanding Networking Options	11
	Working with Networking Protocols	12
	Domain Controllers, Member Servers, and Domain Services . . .	14
	Working with Active Directory	14
	Using Read-Only Domain Controllers	16
	Using Restartable Active Directory Domain Services	16
	Name-Resolution Services . . . . .	17
	Using Domain Name System	18
	Using Windows Internet Name Service	20
	Using Link-Local Multicast Name Resolution	22
	Frequently Used Tools . . . . .	23
	Windows PowerShell 3.0	24
	Windows Remote Management	25

---

### **What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

[microsoft.com/learning/booksurvey](http://microsoft.com/learning/booksurvey)

<b>Chapter 2</b>	<b>Managing Servers Running Windows Server 2012</b>	<b>31</b>
	Server Roles, Role Services, and Features for Windows Server 2012 . . . . .	32
	Full-Server, Minimal-Interface, and Server Core Installations . . .	40
	Navigating Server Core	40
	Installing Windows Server 2012. . . . .	43
	Performing a Clean Installation	44
	Performing an Upgrade Installation	47
	Performing Additional Administration Tasks During Installation	48
	Changing the Installation Type	55
	Managing Roles, Role Services, and Features. . . . .	57
	Performing Initial Configuration Tasks	58
	Server Manager Essentials and Binaries	62
	Managing Your Servers Remotely	65
	Connecting to and Working with Remote Servers	67
	Adding and Removing Roles, Role Services, and Features	70
	Managing System Properties. . . . .	73
	The Computer Name Tab	75
	The Hardware Tab	75
	The Advanced Tab	76
	The Remote Tab	85
<b>Chapter 3</b>	<b>Monitoring Processes, Services, and Events</b>	<b>87</b>
	Managing Applications, Processes, and Performance. . . . .	87
	Task Manager	88
	Viewing and Working with Processes	88
	Administering Processes	91
	Viewing System Services	94
	Viewing and Managing System Performance	95
	Viewing and Managing Remote User Sessions	99
	Managing System Services. . . . .	100
	Navigating Services in Server Manager	100
	Navigating Services in Computer Management	102



Starting, Stopping, and Pausing Services	103
Configuring Service Startup	103
Configuring Service Logon	104
Configuring Service Recovery	106
Disabling Unnecessary Services	107
Event Logging and Viewing . . . . .	108
Accessing Events in Server Manager	109
Accessing Events in Event Viewer	110
Filtering Event Logs	112
Setting Event Log Options	115
Clearing Event Logs	116
Archiving Event Logs	116
Monitoring Server Performance and Activity . . . . .	118
Why Monitor Your Server?	118
Getting Ready to Monitor	119
Using the Monitoring Consoles	119
Choosing Counters to Monitor	122
Performance Logging	124
Viewing Data Collector Reports	128
Configuring Performance Counter Alerts	129
Tuning System Performance . . . . .	130
Monitoring and Tuning Memory Usage	130
Monitoring and Tuning Processor Usage	132
Monitoring and Tuning Disk I/O	133
Monitoring and Tuning Network Bandwidth and Connectivity	134

**Chapter 4 Automating Administrative Tasks, Policies, and Procedures 135**

Understanding Group Policies . . . . .	138
Group Policy Essentials	138
In What Order Are Multiple Policies Applied?	139
When Are Group Policies Applied?	139
Group Policy Requirements and Version Compatibility	140
Navigating Group Policy Changes . . . . .	141

Managing Local Group Policies . . . . .	143
Local Group Policy Objects	143
Accessing the Top-Level Local Policy Settings	144
Local Group Policy Object Settings	145
Accessing Administrator, Non-Administrator, and User-Specific Local Group Policy	146
Managing Site, Domain, and Organizational Unit Policies . . . . .	147
Understanding Domain and Default Policies	147
Using the Group Policy Management Console	148
Getting to Know the Policy Editor	149
Using Administrative Templates to Set Policies	151
Creating and Linking GPOs	152
Creating and Using Starter GPOs	153
Delegating Privileges for Group Policy Management	154
Blocking, Overriding, and Disabling Policies	155
Maintaining and Troubleshooting Group Policy . . . . .	158
Refreshing Group Policy	158
Configuring the Refresh Interval	159
Modeling Group Policy for Planning Purposes	161
Copying, Pasting, and Importing Policy Objects	164
Backing Up and Restoring Policy Objects	165
Determining Current Group Policy Settings and Refresh Status	166
Disabling an Unused Part of Group Policy	166
Changing Policy Processing Preferences	167
Configuring Slow-Link Detection	167
Removing Links and Deleting GPOs	170
Troubleshooting Group Policy	171
Fixing Default Group Policy Objects	172
Managing Users and Computers with Group Policy . . . . .	173
Centrally Managing Special Folders	173
User and Computer Script Management	178
Deploying Software Through Group Policy	181
Automatically Enrolling Computer and User Certificates	186
Managing Automatic Updates in Group Policy	187

<b>Chapter 5</b>	<b>Enhancing Computer Security</b>	<b>191</b>
	Using Security Templates . . . . .	191
	Using the Security Templates and Security Configuration And Analysis Snap-ins	193
	Reviewing and Changing Template Settings	193
	Analyzing, Reviewing, and Applying Security Templates	201
	Deploying Security Templates to Multiple Computers	204
	Using the Security Configuration Wizard . . . . .	206
	Creating Security Policies	206
	Editing Security Policies	211
	Applying Security Policies	211
	Rolling Back the Last-Applied Security Policy	211
	Deploying a Security Policy to Multiple Computers	212

---

**PART II    WINDOWS SERVER 2012 DIRECTORY SERVICES ADMINISTRATION**

---

<b>Chapter 6</b>	<b>Using Active Directory</b>	<b>217</b>
	Introducing Active Directory . . . . .	217
	Active Directory and DNS	217
	Read-Only Domain Controller Deployment	219
	Active Directory Features for Windows Server 2008 R2	220
	Active Directory Features for Windows Server 2012	221
	Working with Domain Structures . . . . .	223
	Understanding Domains	224
	Understanding Domain Forests and Domain Trees	225
	Understanding Organizational Units	227
	Understanding Sites and Subnets	229
	Working with Active Directory Domains . . . . .	230
	Using Computers with Active Directory	230
	Working with Domain Functional Levels	231
	Raising or Lowering Domain and Forest Functionality	235

Understanding the Directory Structure . . . . .	237
Exploring the Data Store . . . . .	238
Exploring Global Catalogs . . . . .	239
Universal Group Membership Caching . . . . .	240
Replication and Active Directory . . . . .	241
Active Directory and LDAP . . . . .	242
Understanding Operations Master Roles . . . . .	242
Using the Active Directory Recycle Bin . . . . .	244
Preparing Schema for the Recycle Bin . . . . .	244
Recovering Deleted Objects . . . . .	245
<b>Chapter 7 Core Active Directory Administration</b> . . . . .	<b>249</b>
Tools for Managing Active Directory . . . . .	249
Active Directory Administration Tools . . . . .	249
Active Directory Command-Line Tools . . . . .	250
Active Directory Support Tools . . . . .	251
Using Active Directory Users And Computers . . . . .	252
Active Directory Administrative Center and Windows PowerShell . . . . .	256
Managing Computer Accounts . . . . .	259
Creating Computer Accounts on a Workstation or Server . . . . .	259
Creating Computer Accounts in Active Directory Administrative Center . . . . .	260
Creating Computer Accounts in Active Directory Users And Computers . . . . .	261
Viewing and Editing Computer Account Properties . . . . .	263
Deleting, Disabling, and Enabling Computer Accounts . . . . .	264
Resetting Locked Computer Accounts . . . . .	264
Moving Computer Accounts . . . . .	266
Managing Computers . . . . .	267
Joining a Computer to a Domain or Workgroup . . . . .	267
Using Offline Domain Join . . . . .	268
Managing Domain Controllers, Roles, and Catalogs. . . . .	270
Installing and Demoting Domain Controllers . . . . .	270
Viewing and Transferring Domainwide Roles . . . . .	273

Viewing and Transferring the Domain Naming Master Role	275
Viewing and Transferring Schema Master Roles	275
Transferring Roles Using the Command Line	276
Seizing Roles Using the Command Line	276
Configuring Global Catalogs	280
Configuring Universal Group Membership Caching	281
Managing Organizational Units . . . . .	281
Creating Organizational Units	281
Viewing and Editing Organizational Unit Properties	282
Renaming and Deleting Organizational Units	282
Moving Organizational Units	282
Managing Sites . . . . .	282
Creating Sites	283
Creating Subnets	284
Associating Domain Controllers with Sites	285
Configuring Site Links	285
Configuring Site Link Bridges	288
Maintaining Active Directory . . . . .	289
Using ADSI Edit	289
Examining Intersite Topology	291
Troubleshooting Active Directory . . . . .	292
<b>Chapter 8 Creating User and Group Accounts</b>	<b>295</b>
The Windows Server Security Model . . . . .	296
Authentication Protocols	296
Access Controls	297
Claims-Based Access Controls	297
Central Access Policies	299
Differences Between User and Group Accounts . . . . .	300
User Accounts	301
Group Accounts	302
Default User Accounts and Groups . . . . .	306
Built-in User Accounts	307
Predefined User Accounts	307

Built-in and Predefined Groups	308
Implicit Groups and Special Identities	309
Account Capabilities . . . . .	309
Privileges	310
Logon Rights	312
Built-in Capabilities for Groups in Active Directory	313
Using Default Group Accounts . . . . .	315
Groups Used by Administrators	316
Implicit Groups and Identities	317
User Account Setup and Organization . . . . .	318
Account Naming Policies	319
Password and Account Policies	320
Configuring Account Policies . . . . .	322
Configuring Password Policies	323
Configuring Account Lockout Policies	325
Configuring Kerberos Policies	326
Configuring User Rights Policies . . . . .	327
Configuring Global User Rights	328
Configuring Local User Rights	330
Adding a User Account . . . . .	330
Creating Domain User Accounts	330
Creating Local User Accounts	334
Adding a Group Account . . . . .	335
Creating a Global Group	336
Creating a Local Group and Assigning Members	337
Handling Global Group Membership . . . . .	338
Managing Individual Membership	339
Managing Multiple Memberships in a Group	340
Setting the Primary Group for Users and Computers	340
Implementing Managed Accounts . . . . .	341
Creating and Using Managed Service Accounts	342
Configuring Services to Use Managed Service Accounts	344
Removing Managed Service Accounts	344

Moving Managed Service Accounts	345
Using Virtual Accounts	346
<b>Chapter 9 Managing User and Group Accounts</b>	<b>347</b>
Managing User Contact Information . . . . .	347
Setting Contact Information	347
Searching for Users and Groups in Active Directory	350
Configuring the User's Environment Settings. . . . .	351
System Environment Variables	352
Logon Scripts	353
Assigning Home Directories	354
Setting Account Options and Restrictions. . . . .	355
Managing Logon Hours	355
Setting Permitted Logon Workstations	357
Setting Dial-in and VPN Privileges	358
Setting Account Security Options	360
Managing User Profiles . . . . .	361
Local, Roaming, and Mandatory Profiles	362
Using the System Utility to Manage Local Profiles	365
Updating User and Group Accounts. . . . .	368
Renaming User and Group Accounts	369
Copying Domain User Accounts	371
Importing and Exporting Accounts	372
Deleting User and Group Accounts	373
Changing and Resetting Passwords	373
Enabling User Accounts	374
Managing Multiple User Accounts . . . . .	375
Setting Profiles for Multiple Accounts	376
Setting Logon Hours for Multiple Accounts	377
Setting Permitted Logon Workstations for Multiple Accounts	378
Setting Logon, Password, and Expiration Properties for Multiple Accounts	378
Troubleshooting Logon Problems. . . . .	378
Viewing and Setting Active Directory Permissions . . . . .	380

<b>Chapter 10 Managing File Systems and Drives</b>	<b>385</b>
Managing the File Services Role . . . . .	385
Adding Hard Disk Drives . . . . .	389
Physical Drives	389
Preparing a Physical Drive for Use	392
Using Disk Management	394
Removable Storage Devices	396
Installing and Checking for a New Drive	398
Understanding Drive Status	399
Working with Basic, Dynamic, and Virtual Disks . . . . .	401
Using Basic and Dynamic Disks	401
Special Considerations for Basic and Dynamic Disks	402
Changing Drive Types	402
Reactivating Dynamic Disks	404
Rescanning Disks	404
Moving a Dynamic Disk to a New System	404
Managing Virtual Hard Disks	405
Using Basic Disks and Partitions . . . . .	406
Partitioning Basics	406
Creating Partitions and Simple Volumes	407
Formatting Partitions	410
Compressing Drives and Data . . . . .	411
Compressing Drives	412
Compressing Directories and Files	412
Expanding Compressed Drives	412
Expanding Compressed Directories and Files	413
Encrypting Drives and Data . . . . .	413
Understanding Encryption and the Encrypting File System	414
Encrypting Directories and Files	415
Working with Encrypted Files and Folders	416
Configuring Recovery Policy	417
Decrypting Files and Directories	418



<b>Chapter 11</b>	<b>Configuring Volumes and RAID Arrays</b>	<b>419</b>
Using Volumes and Volume Sets . . . . .		420
Understanding Volume Basics		420
Understanding Volume Sets		422
Creating Volumes and Volume Sets		424
Deleting Volumes and Volume Sets		426
Managing Volumes		426
Improving Performance and Fault Tolerance with RAID . . . . .		426
Implementing RAID on Windows Server 2012 . . . . .		427
Implementing RAID-0: Disk Striping		427
Implementing RAID-1: Disk Mirroring		428
Implementing RAID-5: Disk Striping with Parity		431
Managing RAID and Recovering from Failures . . . . .		432
Breaking a Mirrored Set		432
Resynchronizing and Repairing a Mirrored Set		432
Repairing a Mirrored System Volume to Enable Boot		433
Removing a Mirrored Set		434
Repairing a Striped Set Without Parity		434
Regenerating a Striped Set with Parity		434
Standards-Based Storage Management . . . . .		435
Getting Started with Standards-Based Storage		435
Working with Standards-Based Storage		436
Creating Storage Pools and Allocating Space		438
Creating a Storage Space		439
Creating a Virtual Disk in a Storage Space		440
Creating a Standard Volume		441
Managing Existing Partitions and Drives . . . . .		443
Assigning Drive Letters and Paths		443
Changing or Deleting the Volume Label		444
Deleting Partitions and Drives		445
Converting a Volume to NTFS		445
Resizing Partitions and Volumes		447
Repairing Disk Errors and Inconsistencies Automatically		449
Analyzing and Optimizing Disks		453

<b>Chapter 12 Data Sharing, Security, and Auditing</b>	<b>457</b>
Using and Enabling File Sharing . . . . .	458
Configuring Standard File Sharing . . . . .	461
Viewing Existing Shares	461
Creating Shared Folders in Computer Management	463
Creating Shared Folders in Server Manager	466
Changing Shared Folder Settings	468
Managing Share Permissions . . . . .	469
Various Share Permissions	469
Viewing and Configuring Share Permissions	470
Managing Existing Shares . . . . .	474
Understanding Special Shares	474
Connecting to Special Shares	475
Viewing User and Computer Sessions	476
Stopping File and Folder Sharing	478
Configuring NFS Sharing . . . . .	479
Using Shadow Copies . . . . .	481
Understanding Shadow Copies	481
Creating Shadow Copies	482
Restoring a Shadow Copy	482
Reverting an Entire Volume to a Previous Shadow Copy	483
Deleting Shadow Copies	483
Disabling Shadow Copies	483
Connecting to Network Drives . . . . .	484
Mapping a Network Drive	484
Disconnecting a Network Drive	485
Object Management, Ownership, and Inheritance . . . . .	485
Objects and Object Managers	485
Object Ownership and Transfer	486
Object Inheritance	487

File and Folder Permissions .....	488
Understanding File and Folder Permissions	489
Setting Basic File and Folder Permissions	491
Setting Special Permissions on Files and Folders	493
Setting Claims-Based Permissions	496
Auditing System Resources .....	498
Setting Auditing Policies	498
Auditing Files and Folders	500
Auditing the Registry	502
Auditing Active Directory Objects	502
Using, Configuring, and Managing NTFS Disk Quotas .....	503
Understanding NTFS Disk Quotas and How NTFS Quotas Are Used	504
Setting NTFS Disk Quota Policies	506
Enabling NTFS Disk Quotas on NTFS Volumes	508
Viewing Disk Quota Entries	510
Creating Disk Quota Entries	510
Deleting Disk Quota Entries	512
Exporting and Importing NTFS Disk Quota Settings	512
Disabling NTFS Disk Quotas	513
Using, Configuring, and Managing Resource Manager Disk Quotas .....	513
Understanding Resource Manager Disk Quotas	514
Managing Disk Quota Templates	515
Creating Resource Manager Disk Quotas	518
<b>Chapter 13 Data Backup and Recovery</b>	<b>519</b>
Creating a Backup and Recovery Plan .....	519
Figuring Out a Backup Plan	519
The Basic Types of Backup	521
Differential and Incremental Backups	522
Selecting Backup Devices and Media	522
Common Backup Solutions	523
Buying and Using Backup Media	524
Selecting a Backup Utility	525

Backing Up Your Data: The Essentials. . . . .	526
Installing the Windows Backup and Recovery Utilities	526
Getting Started with Windows Server Backup	527
Getting Started with the Backup Command-Line Utility	529
Working with Wbadmin Commands	531
Using General-Purpose Commands	531
Using Backup Management Commands	532
Using Recovery Management Commands	533
Performing Server Backups . . . . .	534
Configuring Scheduled Backups	535
Modifying or Stopping Scheduled Backups	538
Creating and Scheduling Backups with Wbadmin	538
Running Manual Backups	540
Recovering Your Server from Hardware or Startup Failure	541
Recovering from a Failed Start	544
Starting a Server in Safe Mode	544
Backing Up and Restoring the System State	546
Restoring Active Directory	547
Restoring the Operating System and the Full System	547
Restoring Applications, Nonsystem Volumes, and Files and Folders	550
Managing Encryption Recovery Policy. . . . .	551
Understanding Encryption Certificates and Recovery Policy	551
Configuring the EFS Recovery Policy	553
Backing Up and Restoring Encrypted Data and Certificates . . .	554
Backing Up Encryption Certificates	554
Restoring Encryption Certificates	555

<b>Chapter 14    Managing TCP/IP Networking</b>	<b>559</b>
Navigating Networking in Windows Server 2012 .....	559
Managing Networking in Windows 8 and Windows Server 2012 .....	562
Installing TCP/IP Networking .....	565
Configuring TCP/IP Networking .....	566
Configuring Static IP Addresses	566
Configuring Dynamic IP Addresses and Alternate IP Addressing	568
Configuring Multiple Gateways	569
Configuring Networking for Hyper-V	570
Managing Network Connections .....	571
Checking the Status, Speed, and Activity for Network Connections	571
Enabling and Disabling Network Connections	572
Renaming Network Connections	572
<b>Chapter 15    Running DHCP Clients and Servers</b>	<b>573</b>
Understanding DHCP .....	573
Using Dynamic IPv4 Addressing and Configuration	573
Using Dynamic IPv6 Addressing and Configuration	575
Checking IP Address Assignment	578
Understanding Scopes	578
Installing a DHCP Server .....	579
Installing DHCP Components	579
Starting and Using the DHCP Console	581
Connecting to Remote DHCP Servers	582
Starting and Stopping a DHCP Server	583
Authorizing a DHCP Server in Active Directory	583

Configuring DHCP Servers . . . . .	584
Configuring Server Bindings	584
Updating DHCP Statistics	584
DHCP Auditing and Troubleshooting	585
Integrating DHCP and DNS	586
Integrating DHCP and NAP	588
Avoiding IP Address Conflicts	591
Saving and Restoring the DHCP Configuration	591
Managing DHCP Scopes . . . . .	592
Creating and Managing Superscopes	592
Creating and Managing Scopes	593
Creating and Managing Failover Scopes	602
Managing the Address Pool, Leases, and Reservations. . . . .	605
Viewing Scope Statistics	605
Enabling and Configuring MAC Address Filtering	606
Setting a New Exclusion Range	607
Reserving DHCP Addresses	608
Modifying Reservation Properties	609
Deleting Leases and Reservations	610
Backing Up and Restoring the DHCP Database . . . . .	610
Backing Up the DHCP Database	610
Restoring the DHCP Database from Backup	611
Using Backup and Restore to Move the DHCP Database to a New Server	611
Forcing the DHCP Server Service to Regenerate the DHCP Database	612
Reconciling Leases and Reservations	612

<b>Chapter 16 Optimizing DNS</b>	<b>615</b>
Understanding DNS	615
Integrating Active Directory and DNS	616
Enabling DNS on the Network	617
Configuring Name Resolution on DNS Clients	620
Installing DNS Servers	621
Installing and Configuring the DNS Server Service	622
Configuring a Primary DNS Server	624
Configuring a Secondary DNS Server	627
Configuring Reverse Lookups	628
Configuring Global Names	629
Managing DNS Servers	631
Adding and Removing Servers to Manage	631
Starting and Stopping a DNS Server	632
Using DNSSEC and Signing Zones	632
Creating Child Domains Within Zones	634
Creating Child Domains in Separate Zones	635
Deleting a Domain or Subnet	636
Managing DNS Records	636
Adding Address and Pointer Records	637
Adding DNS Aliases with CNAME	638
Adding Mail Exchange Servers	638
Adding Name Servers	639
Viewing and Updating DNS Records	640
Updating Zone Properties and the SOA Record	641
Modifying the SOA Record	641
Allowing and Restricting Zone Transfers	643
Notifying Secondaries of Changes	644
Setting the Zone Type	645
Enabling and Disabling Dynamic Updates	645

Managing DNS Server Configuration and Security . . . . .	645
Enabling and Disabling IP Addresses for a DNS Server	646
Controlling Access to DNS Servers Outside the Organization	646
Enabling and Disabling Event Logging	648
Using Debug Logging to Track DNS Activity	648
Monitoring a DNS Server	649
<i>Index</i>	651

---

**What do you think of this book? We want to hear from you!**

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

[microsoft.com/learning/booksurvey](https://microsoft.com/learning/booksurvey)



# Introduction

---

Welcome to *Windows Server 2012 Pocket Consultant*. Over the years, I've written about many different server technologies and products, but the one product I like writing about the most is Microsoft Windows Server. For anyone transitioning to Windows Server 2012 from an earlier release of Windows Server, I'll let you know right up front that I believe this is the most significant update to Windows Server since the introduction of Windows 2000 Server. While the extensive UI changes are a key part of the revisions to the operating system, the deeper changes are below the surface, in the underlying architecture.

The good news is Windows Server 2012 builds off the same code base as Microsoft Windows 8. This means that you can apply much of what you know about Windows 8 to Windows Server 2012, including how Windows works with touch-based user interfaces. Although you might not install Windows Server 2012 on touch UI-capable computers, you can manage Windows Server 2012 from your touch UI-capable computers. If you do end up managing it this way, understanding the touch UI as well as the revised interface options will be crucial to your success. For this reason, I discuss both the touch UI and the traditional mouse and keyboard techniques throughout this book.

When you are working with touch UI-enabled computers, you can manipulate onscreen elements in ways that weren't possible previously. You can enter text using the onscreen keyboard and interact with screen elements in the following ways:

- **Tap** Tap an item by touching it with your finger. A tap or double-tap of elements on the screen generally is the equivalent of a mouse click or double-click.
- **Press and hold** Press your finger down and leave it there for a few seconds. Pressing and holding elements on the screen generally is the equivalent of a right-click.
- **Swipe to select** Slide an item a short distance in the opposite direction compared to how the page scrolls. This selects the items and also might bring up related commands. If pressing and holding doesn't display commands and options for an item, try using swipe to select instead.
- **Swipe from edge (slide in from edge)** Starting from the edge of the screen, swipe or slide in. Sliding in from the right edge opens the Charms panel. Sliding in from the left edge shows open apps and allows you to easily switch between them. Sliding in from the top or bottom edge shows commands for the active element.
- **Pinch** Touch an item with two or more fingers and then move the fingers toward each other. Pinching zooms in or shows less information.
- **Stretch** Touch an item with two or more fingers and then move the fingers away from each other. Stretching zooms out or shows more information.

Because I've written many top-selling Windows Server books, I was able to bring a unique perspective to this book—the kind of perspective you gain only after working with technologies for many years. Long before there was a product called Windows Server 2012, I was working with the beta product. From these early beginnings, the final version of Windows Server 2012 evolved until it became the finished product that is available today.

As you've probably noticed, a great deal of information about Windows Server 2012 is available on the web and in other printed books. You can find tutorials, reference sites, discussion groups, and more to make using Windows Server 2012 easier. However, the advantage of reading this book is that much of the information you need to learn about Windows Server 2012 is organized in one place and presented in a straightforward and orderly fashion. This book has everything you need to customize Windows Server 2012 installations, master Windows Server 2012 configurations, and maintain Windows Server 2012 servers.

In this book, I teach you how features work, why they work the way they do, and how to customize them to meet your needs. I also offer specific examples of how certain features can meet your needs, and how you can use other features to troubleshoot and resolve issues you might have. In addition, this book provides tips, best practices, and examples of how to optimize Windows Server 2012. This book won't just teach you how to configure Windows Server 2012, it will teach you how to squeeze every last bit of power out of it and make the most of the features and options it includes.

Unlike many other books about administering Windows Server 2012, this book doesn't focus on a specific user level. This isn't a lightweight beginner book. Regardless of whether you are a beginning administrator or a seasoned professional, many of the concepts in this book will be valuable to you, and you can apply them to your Windows Server 2012 installations.

---

## Who Is This Book For?

*Windows Server 2012 Pocket Consultant* covers all editions of Windows Server 2012. The book is designed for the following readers:

- Current Windows system administrators
- Accomplished users who have some administrator responsibilities
- Administrators upgrading to Windows Server 2012 from previous versions
- Administrators transferring from other platforms

To pack in as much information as possible, I had to assume that you have basic networking skills and a basic understanding of Windows Server. With this in mind, I don't devote entire chapters to explaining Windows Server architecture, Windows Server startup and shutdown, or why you want to use Windows Server. I do, however, cover Windows server configuration, Group Policy, security, auditing, data backup, system recovery, and much more.

I also assume that you are fairly familiar with Windows commands and procedures as well as the Windows user interface. If you need help learning Windows basics, you should read other resources (many of which are available from Microsoft Press).

## How This Book Is Organized

---

Rome wasn't built in a day, and this book wasn't intended to be read in a day, in a week, or even in a month. Ideally, you'll read this book at your own pace, a little each day as you work your way through all the features Windows Server 2012 has to offer. This book is organized into 16 chapters. The chapters are arranged in a logical order, taking you from planning and deployment tasks to configuration and maintenance tasks.

Ease of reference is an essential part of this hands-on guide. This book has an expanded table of contents and an extensive index for finding answers to problems quickly. Many other quick-reference features have been added to the book as well, including quick step-by-step procedures, lists, tables with fast facts, and extensive cross references.

As with all Pocket Consultants, *Windows Server 2012 Pocket Consultant* is designed to be a concise and easy-to-use resource for managing Windows servers. This is the readable resource guide that you'll want on your desktop at all times. The book covers everything you need to perform the core administrative tasks for Windows servers. Because the focus is on giving you maximum value in a pocket-size guide, you don't have to wade through hundreds of pages of extraneous information to find what you're looking for. Instead, you'll find exactly what you need to get the job done, and you'll find it quickly.

In short, the book is designed to be the one resource you turn to whenever you have questions regarding Windows Server administration. To this end, the book zeroes in on daily administration procedures, frequently performed tasks, documented examples, and options that are representative while not necessarily inclusive. One of my goals is to keep the content so concise that the book remains compact and easy to navigate while at the same time ensuring that it is packed with as much information as possible.

## Conventions Used in This Book

---

I've used a variety of elements to help keep the text clear and easy to follow. You'll find code listings in monospace type. When I tell you to actually type a command, the command appears in **bold** type. When I introduce and define a new term or use a code term in a paragraph of text, I put it in *italics*.

**NOTE** Group Policy now includes both policies and preferences. Under the Computer Configuration and User Configuration nodes, you find two nodes: Policies and Preferences. Settings for general policies are listed under the Policies node. Settings for general preferences are listed under the Preferences node. When referencing settings under the Policies node, I sometimes use shortcut references, such as User Configuration\Administrative Templates\Windows Components, or specify that the policies are found in the Administrative Templates for User Configuration under Windows Components. Both references tell you that the policy setting being discussed is under User Configuration rather than Computer Configuration and can be found under Administrative Templates\Windows Components.

Other conventions include the following:

- **Best Practices** To examine the best technique to use when working with advanced configuration and maintenance concepts
- **Caution** To warn you about potential problems you should look out for
- **More Info** To provide more information on a subject
- **Note** To provide additional details on a particular point that needs emphasis
- **Real World** To provide real-world advice when discussing advanced topics
- **Security Alert** To point out important security issues
- **Tip** To offer helpful hints or additional information

I truly hope you find that *Windows Server 2012 Pocket Consultant* provides everything you need to perform the essential administrative tasks on Windows servers as quickly and efficiently as possible. You are welcome to send your thoughts to me at [williamstaneke@aol.com](mailto:williamstaneke@aol.com). Follow me on Twitter at WilliamStanek and on Facebook at [www.facebook.com/William.Stanek.Author](http://www.facebook.com/William.Stanek.Author).

## Other Resources

---

No single magic bullet for learning everything you'll ever need to know about Windows Server 2012 exists. While some books are offered as all-in-one guides, there's simply no way one book can do it all. With this in mind, I hope you use this book as it is intended to be used—as a concise and easy-to-use resource. It covers everything you need to perform core administration tasks for Windows servers, but it is by no means exhaustive.

Your current knowledge will largely determine your success with this or any other Windows resource or book. As you encounter new topics, take the time to practice what you've learned and read about. Seek out further information as necessary to get the practical hands-on know-how and knowledge you need.

I recommend that you regularly visit the Microsoft website for Windows Server ([microsoft.com/windowsserver/](http://microsoft.com/windowsserver/)) and [support.microsoft.com](http://support.microsoft.com) to stay current with the latest changes. To help you get the most out of this book, you can visit my corresponding website at [williamstaneke.com/windows](http://williamstaneke.com/windows). This site contains information about Windows Server 2012 and updates to the book.

## Errata & Book Support

---

We've made every effort to ensure the accuracy of this book and its companion content. Any errors that have been reported since this book was published are listed on our Microsoft Press site at *oreilly.com*:

*<http://go.microsoft.com/fwlink/?Linkid=258651>*

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at:

*[mspinput@microsoft.com](mailto:mspinput@microsoft.com)*.

Please note that product support for Microsoft software is not offered through the addresses above.

## We Want to Hear from You

---

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

*<http://www.microsoft.com/learning/booksurvey>*

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

## Stay in Touch

---

Let's keep the conversation going! We're on Twitter: *<http://twitter.com/MicrosoftPress>*.



PART I

# Windows Server 2012 Administration Fundamentals

- CHAPTER 1** Windows Server 2012 Administration  
Overview **3**
- CHAPTER 2** Managing Servers Running Windows Server  
2012 **31**
- CHAPTER 3** Monitoring Processes, Services, and  
Events **87**
- CHAPTER 4** Automating Administrative Tasks, Policies, and  
Procedures **135**
- CHAPTER 5** Enhancing Computer Security **191**





# Windows Server 2012

## Administration Overview

- Windows Server 2012 and Windows 8 **4**
- Getting to Know Windows Server 2012 **6**
- Power Management Options **8**
- Networking Tools and Protocols **11**
- Domain Controllers, Member Servers, and Domain Services **14**
- Name-Resolution Services **17**
- Frequently Used Tools **23**

Microsoft Windows Server 2012 is a powerful, versatile, full-featured server operating system that builds on the enhancements that Microsoft provided in Windows Server 2008 Release 2. Windows Server 2012 and Windows 8 share a number of common features because they were part of a single development project. These features share 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 8 to Windows Server 2012.

This chapter covers getting started with Windows Server 2012 and explores the extent to which the architectural changes affect how you work with and manage Windows Server 2012. Throughout this chapter and the other chapters of this book, you'll also find discussions of the many security features and enhancements. These discussions explore all aspects of computer security, including physical security, information security, and network security. Although this book focuses on Windows Server 2012 administration, the tips and techniques it presents can help anyone who supports, develops for, or works with the Windows Server 2012 operating system.

# Windows Server 2012 and Windows 8

---

Before you deploy Windows Server 2012, you should carefully plan the server architecture. As part of your implementation planning, you need to look closely at the software configuration that will be used and modify the hardware configuration on a per-server basis to meet related requirements. For additional flexibility in server deployments, you can deploy servers using one of three installation types:

- **Server With A GUI installation** 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. This installation option provides the most dynamic solution and is recommended for deployments of Windows Server 2012 in which the server role might change over time.
- **Server Core installation** A minimal installation option that provides a fixed subset of roles but does not include the Server Graphical Shell, Microsoft Management Console, or 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 or remotely 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.
- **Server With Minimal Interface installation** An intermediate installation option where you perform a full-server installation and then remove the Server Graphical Shell. This leaves a minimal user interface, Microsoft Management Console, Server Manager, and a subset of Control Panel for local management. This installation option is ideally suited to situations in which you want to carefully control the tasks that can be performed on a server, as well as the roles and features installed, but still want the convenience of the graphical interface.

You choose the installation type during installation of the operating system. In a significant change from earlier releases of Windows Server, you can change the installation type once you've installed a server. A key difference between the installation types relates to the presence of the graphical management tools and the graphical shell. A Server Core installation has neither; a full-server installation has both; and a minimal-interface installation has only the graphical management tools.

**MORE INFO** Several server features and roles require the graphical shell. They include Fax Server, Remote Desktop Session Host, Windows Deployment Services, and the Internet Printing user interface. Additionally, in Event Viewer, the Details view requires the graphical shell, as does the graphical interface for Windows Firewall.

Like Windows 8, Windows Server 2012 has the following features:

- **Modularization for language independence and disk imaging for hardware independence** Each component of the operating system is designed as an independent module you can easily add or remove. This functionality provides the basis for the configuration architecture in Windows Server 2012. Microsoft distributes Windows Server 2012 on media with Windows Imaging Format (WIM) disk images that use compression and single-instance storage to dramatically reduce the size of image files.
- **Preinstallation and preboot environments** The Windows Preinstallation Environment 4.0 (Windows PE 4.0) replaces MS-DOS as the preinstallation environment and provides a bootable startup environment for installation, deployment, recovery, and troubleshooting. The Windows Preboot Environment provides a startup environment with a boot manager that lets you choose which boot application to run to load the operating system. On systems with multiple operating systems, you access pre-Windows 7 operating systems in the boot environment by using the legacy operating system entry.
- **User account controls and elevation of privileges** User Account Control (UAC) enhances computer security by ensuring true separation of standard user and administrator user accounts. Through UAC, all applications run using either standard user or administrator user privileges, and you see a security prompt by default whenever you run an application that requires administrator privileges. The way the security prompt works depends on Group Policy settings. Additionally, if you log on using the built-in Administrator account, you typically do not see elevation prompts.

In Windows 8 and Windows Server 2012, features with common code bases have identical management interfaces. In fact, just about every Control Panel utility that is available in Windows Server 2012 is identical to or nearly identical to its Windows 8 counterpart. Of course, exceptions exist in some cases for standard default settings. Because Windows Server 2012 does not use performance ratings, Windows servers do not have Windows Experience Index scores. Because Windows Server 2012 does not use Sleep or related states, Windows servers do not have sleep, hibernate, or resume functionality. Because you typically do not want to use extended power management options on Windows servers, Windows Server 2012 has a limited set of power options.

Windows Server 2012 does not include the Windows Aero enhancements, Windows Sidebar, Windows Gadgets, or other user-interface enhancements because Windows Server 2012 is designed to provide optimal performance for server-related tasks and is not designed for extensive personalization of the desktop appearance. That said, when you are working with a full-server installation, you can add the Desktop Experience feature and then enable some Windows 8 features on your server.

The Desktop Experience provides Windows desktop functionality on the server. Windows features added include Windows Media Player, desktop themes, Video for Windows (AVI support), Windows Defender, 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.

Because the common features of Windows 8 and Windows Server 2012 have so many similarities, I will not cover changes in the interface from previous operating system releases, discuss how UAC works, and so on. You can find extensive coverage of these features in *Windows 8 Administration Pocket Consultant* (Microsoft Press, 2012), which I encourage you to use in conjunction with this book. In addition to its coverage of broad administration tasks, *Windows 8 Administration Pocket Consultant* examines how to customize the operating system and Windows environment, configure hardware and network devices, manage user access and global settings, configure laptops and mobile networking, use remote management and remote assistance capabilities, troubleshoot system problems, and much more. This book, on the other hand, zeroes in on directory services administration, data administration, and network administration.

## Getting to Know Windows Server 2012

---

The Windows Server 2012 operating system includes several different editions. All Windows Server 2012 editions support multiple processor cores. It is important to point out that although an edition might support only one discrete-socketed processor (also referred to as a *physical processor*), that one processor could have eight processor cores (also referred to as *logical processors*).

Windows Server 2012 is a 64-bit-only operating system. In this book, I refer to 64-bit systems designed for the x64 architecture as *64-bit* systems. 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 2012 edition you're using.

When you install a Windows Server 2012 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 2012 systems that manage access to the network, to the directory database, and to shared resources.

**NOTE** In this book, *Windows Server 2012* and *Windows Server 2012 family* refer to all editions of Windows Server 2012. The various server editions support the same core features and administration tools.

Unlike Windows Server 2008, Windows Server 2012 uses a Start screen. Start is a window, not a menu. Programs can have tiles on the Start screen. Tapping or clicking a tile runs the program. When you press and hold or right-click on a program, an options panel normally is displayed. The charms bar is an options panel for Start, Desktop, and PC Settings. With a touch UI, you can display the charms by sliding in

from the right side of the screen. With a mouse and keyboard, you can display the charms by moving the mouse pointer over the hidden button in the upper-right or lower-right corner of the Start, Desktop, or PC Settings screen; or by pressing Windows key+C.

Tap or click the Search charm to display the Search panel. Any text typed while on the Start screen is entered into the Search box in the Search panel. The Search box can be focused on Apps, Settings, or Files. When focused on Apps, you can use Search to quickly find installed programs. When focused on Settings, you can use Search to quickly find settings and options in Control Panel. When focused on Files, you can use Search to quickly find files.

One way to quickly open a program is by pressing the Windows key, typing the file name of the program, and then pressing Enter. This shortcut works as long as the Apps Search box is in focus (which it typically is by default).

Pressing the Windows key toggles between the Start screen and the desktop (or, if you are working with PC Settings, between Start and PC Settings). On Start, there's a Desktop tile that you can tap or click to display the desktop. You also can display the desktop by pressing Windows key+D or, to peek at the desktop, press and hold Windows key+Comma. From Start, you access Control Panel by tapping or clicking the Control Panel tile. From the desktop, you can display Control Panel by accessing the charms, tapping or clicking Settings, and then tapping or clicking Control Panel. Additionally, because File Explorer is pinned to the desktop taskbar by default you typically can access Control Panel on the desktop by following these steps:

1. Open File Explorer by tapping or clicking the taskbar icon.
2. Tap or click the leftmost option button (down arrow) in the address list.
3. Tap or click Control Panel.

Start and Desktop have a handy menu that you can display by pressing and holding or right-clicking the lower-left corner of the Start screen or the desktop. Options on the menu include Command Prompt, Command Prompt (Admin), Device Manager, Event Viewer, System, and Task Manager. On Start, the hidden button in the lower-left corner shows a thumbnail view of the desktop when activated, and tapping or clicking the thumbnail opens the desktop. On the desktop, the hidden button in the lower-left corner shows a thumbnail view of Start when activated and tapping or clicking the thumbnail opens Start. Pressing and holding or right-clicking the thumbnail is what displays the shortcut menu.

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

1. Display Start options by sliding in from the right side of the screen or moving the mouse pointer to the bottom right or upper right corner of the screen.
2. Tap or click Settings and then tap or click Power.
3. Tap or click Shut Down or Restart as appropriate.

Alternatively, 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, servers are configured to turn off the display after 10 minutes of inactivity.

Windows 8 and Windows Server 2012 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 tapping or clicking System And Security and then tapping or clicking Power Options. Windows Server 2012 includes the Power Configuration (Powercfg.exe) utility for managing power options from the command line. At a command prompt, you can view the configured power plans by typing **powercfg /l**. The active power plan is marked with an asterisk.

The default, active power plan in Windows Server 2012 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 2012 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.

## Power Management Options

---

When working with power management, important characteristics to focus on include the following:

- Cooling modes
- Device states
- Processor states

ACPI defines active and passive cooling modes. These cooling modes are inversely related to each other:

- Passive cooling reduces system performance but is quieter because there’s less fan noise. With passive cooling, Windows lessens power consumption to

reduce the operating temperature of the computer but at the cost of system performance. Here, Windows reduces the processor speed in an attempt to cool the computer before increasing fan speed, which would increase power consumption.

- Active cooling allows maximum system performance. With active cooling, Windows increases power consumption to reduce the temperature of the machine. Here, Windows increases fan speed to cool the computer before attempting to reduce processor speed.

Power policy includes an upper and lower limit for the processor state, referred to as the *maximum processor state* and the *minimum processor state*, respectively. These states are implemented by making use of a feature of ACPI 3.0 and later versions called processor throttling, and they determine the range of currently available processor performance states that Windows can use. By setting the maximum and minimum values, you define the bounds for the allowed performance states, or you can use the same value for each to force the system to remain in a specific performance state. Windows reduces power consumption by throttling the processor speed. For example, if the upper bound is 100 percent and the lower bound is 5 percent, Windows can throttle the processor within this range as workloads permit to reduce power consumption. In a computer with a 3-GHz processor, Windows would adjust the operating frequency of the processor between .15 GHz and 3.0 GHz.

Processor throttling and related performance states were introduced with Windows XP and are not new, but these early implementations were designed for computers with discrete-socketed processors and not for computers with processor cores. As a result, they are not effective in reducing the power consumption of computers with logical processors. Windows 7 and later releases of Windows reduce power consumption in computers with multicore processors by leveraging a feature of ACPI 4.0 called *logical processor idling* and by updating processor throttling features to work with processor cores.

Logical processor idling is designed to ensure that Windows uses the fewest number of processor cores for a given workload. Windows accomplishes this by consolidating workloads onto the fewest cores possible and suspending inactive processor cores. As additional processing power is required, Windows activates inactive processor cores. This idling functionality works in conjunction with management of process performance states at the core level.


ACPI defines processor performance states, referred to as *p-states*, and processor idle sleep states, referred to as *c-states*. Processor performance states include P0 (the processor/core uses its maximum performance capability and can consume maximum power), P1 (the processor/core is limited below its maximum and consumes less than maximum power), and P<sub>*n*</sub> (where state *n* is a maximum number that is processor dependent, and the processor/core is at its minimal level and consumes minimal power while remaining in an active state).

Processor idle sleep states include C0 (the processor/core can execute instructions), C1 (the processor/core has the lowest latency and is in a nonexecuting power state), C2 (the processor/core has longer latency to improve power savings over the C1 state), and C3 (the processor/core has the longest latency to improve power savings over the C1 and C2 states).

**MORE INFO** ACPI 4.0 was finalized in June 2009 and ACPI 5.0 was finalized in December 2011. Computers manufactured prior to this time will likely not have firmware that is fully compliant, and you will probably need to update the firmware when a compatible revision becomes available. In some cases, and especially with older hardware, you might not be able to update a computer's firmware to make it fully compliant with ACPI 4.0 or ACPI 5.0. For example, if you are configuring the power options and you don't have minimum and maximum processor state options, the computer's firmware isn't fully compatible with ACPI 3.0 and likely will not fully support ACPI 4.0 or ACPI 5.0 either. Still, you should check the hardware manufacturer's website for firmware updates.

Windows switches processors/cores between any p-state and from the C1 state to the C0 state nearly instantaneously (fractions of milliseconds) and tends not to use the deep sleep states, so you don't need to worry about performance impact to throttle or wake up processors/cores. The processors/cores are available when they are needed. That said, the easiest way to limit processor power management is to modify the active power plan and set the minimum and maximum processor states to 100 percent. Logical processor idling is used to reduce power consumption by removing a logical processor from the operating system's list of nonprocessor-affinitized work. However, because processor-affinitized work reduces the effectiveness of this feature, you'll want to plan carefully prior to configuring processing affinity settings for applications. Windows System Resource Manager allows you to manage processor resources through percent processor usage targets and processor affinity rules. Both techniques reduce the effectiveness of logical processor idling.

Windows saves power by putting processor cores in and out of appropriate p-states and c-states. On a computer with four logical processors, Windows might use p-states 0 to 5, where P0 allows 100 percent usage, P1 allows 90 percent usage, P2 allows 80 percent usage, P3 allows 70 percent usage, P4 allows 60 percent usage, and P5 allows 50 percent usage. When the computer is active, logical processor 0 would likely be active with a p-state of 0 to 5, and the other processors would likely be at an appropriate p-state or in a sleep state. Figure 1-1 shows an example. Here, logical processor 1 is running at 90 percent, logical processor 2 is running at 80 percent, logical processor 3 is running at 50 percent, and logical processor 4 is in the sleep state.



**FIGURE 1-1** Understanding processor states



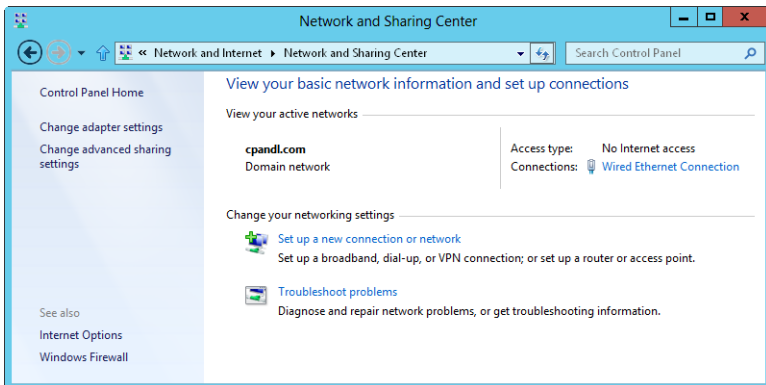
**REAL WORLD** ACPI 4.0 and ACPI 5.0 define four global power states. In G0, the working state in which software runs, power consumption is at its highest and latency is at its lowest. In G1, the sleeping state, in which software doesn't run, latency varies with sleep state and power consumption is less than the G0 state. In G2 (also referred to as S5 sleep state), the soft off state where the operating system doesn't run, latency is long and power consumption is very near zero. In G3, the mechanical off state, where the operating system doesn't run, latency is long, and power consumption is zero. There's also a special global state, known as S4 nonvolatile sleep, in which the operating system writes all system context to a file on nonvolatile storage media, allowing system context to be saved and restored.

Within the global sleeping state, G1, are sleep-state variations. S1 is a sleeping state where all system context is maintained. S2 is a sleeping state similar to S1 except that the CPU and system-cache contexts are lost and control starts from a reset. S3 is a sleeping state where all CPU, cache, and chip-set context are lost and hardware maintains memory context and restores some CPU and L2 cache configuration context. S4 is a sleeping state in which it is assumed that the hardware has powered off all devices to reduce power usage to a minimum and only the platform context is maintained. S5 is a sleeping state in which it is assumed that the hardware is in a soft off state, where no context is maintained and a complete boot is required when the system wakes.

Devices have power states as well. D0, the fully on state, consumes the highest level of power. D1 and D2 are intermediate states that many devices do not use. D3hot is a power-saving state, where the device is software enumerable and can optionally preserve device context. D3 is the off state, where the device context is lost and the operating system must reinitialize the device to turn it back on.

## Networking Tools and Protocols

Windows Server 2012 has a suite of networking tools that includes Network Explorer, Network And Sharing Center, and Network Diagnostics. Figure 1-2 shows Network And Sharing Center.



**FIGURE 1-2** Network And Sharing Center provides quick access to sharing, discovery, and networking options.

# Understanding Networking Options

The sharing and discovery configuration in Network And Sharing Center controls basic network settings. When network discovery settings are turned on and a server is connected to a network, the server can see other network computers and devices and is visible on the network. When sharing settings are turned on or off, the various sharing options are allowed or restricted. As discussed in Chapter 12, “Data Sharing, Security, and Auditing,” sharing options include file sharing, public folder sharing, printer sharing, and password-protected sharing.

In Windows 8 and Windows Server 2012, 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.
- **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: home or work, 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, you see a dialog box that allows you to 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.

Based on the network category, Windows Server configures settings that turn discovery on or off. The On (enabled) state means that the computer can discover other computers and devices on the network and that other computers on the network can discover the computer. The Off (disabled) state means that the computer cannot discover other computers and devices on the network and that other computers on the network cannot discover the computer.

Using either the Network window or Advanced Sharing Settings in Network And Sharing Center, you can enable discovery and file sharing. However, 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

To allow a server to access a network, you must install TCP/IP networking and a network adapter. 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 7 and later releases of Windows have a dual IP-layer 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, including the Internet. IPv6, on the other hand, has 128-bit addresses and is the next-generation version of IP.

**NOTE** DirectAccess clients only send IPv6 traffic across the DirectAccess connection to the DirectAccess server. Thanks to the NAT64/DNS64 support on a Windows Server 2012 DirectAccess server, DirectAccess clients can now 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 feature was introduced with Windows Vista and Windows Server 2008. This feature 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 or 10 Gbps network adapter, you must enable TCP offloading by entering the following command at an elevated, 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 (RSS) and network direct memory access (NetDMA). You can enable or disable RSS 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 the internal machine (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:BECB: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:BECB:FE4F:961D. Three or more 0 blocks would be compressed in the same way. For example, FFE8:0:0:0:0:0:1 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. The modified IP architecture in Windows 7 and later releases of Windows is referred to as the *Next Generation TCP/IP stack*, and it includes many enhancements that improve the way IPv4 and IPv6 are used.

## Domain Controllers, Member Servers, and Domain Services

---

When you install Windows Server 2012 on a new system, you can configure the server to be a member server, a domain controller, or a standalone server. The differences between these types of servers are extremely important. 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. Standalone servers aren't part of a domain. Because standalone servers have their own user databases, they authenticate logon requests independently.

### Working with Active Directory

Windows Server 2012 supports a multimaster replication model. In this 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 as well as 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.

Microsoft changed Active Directory in several fundamental ways for the original release of Windows Server 2008. As a result, Microsoft realigned the directory functionality and created a family of related services, including the following:

- **Active Directory Certificate Services (AD CS)** AD CS provides functions necessary for issuing and revoking digital certificates for users, client

computers, and servers. AD CS uses certificate authorities (CAs), which are responsible for confirming the identity of users and computers and then issuing certificates to confirm these identities. Domains have enterprise root CAs, which are the certificate servers at the root of certificate hierarchies for domains and the most trusted certificate servers in the enterprise, and subordinate CAs, which are members of a particular enterprise certificate hierarchy. Workgroups have standalone root CAs, which are the certificate servers at the root of nonenterprise certificate hierarchies, and standalone subordinate CAs, which are members of a particular nonenterprise certificate hierarchy.

- **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. Once 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.
- **Active Directory Federation Services (AD FS)** AD FS complements the authentication and access-management features of AD DS by extending them to the World Wide Web. AD FS uses web agents to provide users with access to internally hosted web applications and proxies to manage client access. Once AD FS is configured, users can use their digital identities to authenticate themselves over the Web and access internally hosted web applications with a web browser such as Internet Explorer.
- **Active Directory Lightweight Directory Services (AD LDS)** AD LDS provides a data store for directory-enabled applications that do not require AD DS and do not need to be deployed on domain controllers. AD LDS does not run as an operating system service and can be used in both domain and workgroup environments. Each application that runs on a server can have its own data store implemented through AD LDS.
- **Active Directory Rights Management Services (AD RMS)** AD RMS provides a layer of protection for an organization's information that can extend beyond the enterprise, allowing email messages, documents, intranet webpages, and more to be protected from unauthorized access. AD RMS uses a certificate service to issue rights account certificates that identify trusted users, groups, and services; a licensing service that provides authorized users, groups, and services with access to protected information; and a logging service to monitor and maintain the rights management service. Once trust has been established, users with a rights account certificate can assign rights to information. These rights control which users can access the information and what they can do with it. Users with rights account certificates can also access protected content to which they've been granted access. Encryption ensures that access to protected information is controlled both inside and outside the enterprise.

Microsoft introduced additional changes in Windows Server 2012. These changes include a new domain functional level, called *Windows Server 2012 domain functional level*, and a new forest functional level, called *Windows Server 2012 forest functional level*. The many other changes are discussed in Chapter 6, “Using Active Directory.”

## Using Read-Only Domain Controllers

Windows Server 2008 and later releases support read-only domain controllers and Restartable Active Directory Domain Services. A read-only domain controller (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.

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 that is running Windows Server 2008 or later. 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.

**TIP** 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. Although RODCs can pull information from domain controllers running Windows Server 2003, RODCs can pull updates of the domain partition only from a writable domain controller running Windows Server 2008 or later in the same domain.

## Using Restartable Active Directory Domain Services

Restartable Active Directory Domain Services 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. While AD DS is stopped on a server, 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 using the Directory Services Restore Mode account and password.

All domain controllers running Windows Server 2008 or later support Restartable Active Directory Domain Services—even RODCs. As an administrator, you can start

or stop AD DS by using the Domain Controller entry in the Services utility. Because of Restartable Active Directory Domain Services, domain controllers running Windows Server 2008 or later 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 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 when you stop AD DS. 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 normally and then stop AD DS.

## Name-Resolution Services

---

Windows operating systems use name resolution to make it easier to communicate with other computers on a network. Name resolution associates computer names with the numerical IP addresses that are used for network communications. Thus, rather than using long strings of digits, users can access a computer on the network by using a friendly name.

Current Windows operating systems natively support three name-resolution systems:

- Domain Name System (DNS)
- Windows Internet Name Service (WINS)
- Link-Local Multicast Name Resolution (LLMNR)

The sections that follow examine these services.

## Using Domain Name System

DNS is a name-resolution service that resolves computer names to IP addresses. Using DNS, the fully qualified host name `computer84.cpanidl.com`, for example, can be resolved to an IP address, which allows it and other computers to find one another. DNS operates over the TCP/IP protocol stack and can be integrated with WINS, Dynamic Host Configuration Protocol (DHCP), and Active Directory Domain Services. As discussed in Chapter 15, “Running DHCP Clients and Servers,” DHCP is used for dynamic IP addressing and TCP/IP configuration.

DNS organizes groups of computers into domains. These domains are organized into a hierarchical structure, which can be defined on an Internet-wide basis for public networks or on an enterprise-wide basis for private networks (also known as *intranets* and *extranets*). The various levels within the hierarchy identify individual computers, organizational domains, and top-level domains. For the fully qualified host name `computer84.cpanidl.com`, *computer84* represents the host name for an individual computer, *cpanidl* is the organizational domain, and *com* is the top-level domain.

Top-level domains are at the root of the DNS hierarchy; they are also called *root domains*. These domains are organized geographically, by organization type, and by function. Normal domains, such as `cpanidl.com`, are also referred to as *parent domains*. They're called parent domains because they're the parents of an organizational structure. Parent domains can be divided into subdomains that can be used for groups or departments within an 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 `jacob.hr.cpanidl.com`. Here, *jacob* is the host name, *hr* is the child domain, and *cpanidl.com* is the parent domain.

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 domain controllers using Active Directory. During installation of the first domain controller on an Active Directory network, you're given the opportunity to install DNS automatically if a DNS server can't be found on the network. You are also able to specify whether DNS and Active Directory should be fully integrated. In most cases, you should respond affirmatively to both requests. With full integration, DNS information is stored directly in Active Directory. This allows you to take advantage of Active Directory's capabilities. 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, and 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 a multimaster approach can be used for dynamic updates through DHCP. This allows 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 can see more advantages to full integration with Active Directory. With partial integration, DNS information is stored and replicated separately from Active Directory. Having two separate structures reduces the effectiveness of both DNS and Active Directory and makes 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 it takes to replicate DNS changes throughout 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. Using these addresses, clients can communicate with DNS servers anywhere on the network, even if the servers are on different subnets.

When the network uses DHCP, you should configure DHCP to work with DNS. To do this, you need to set the DHCP scope options 006 DNS Servers and 015 DNS Domain Name as specified in "Setting Scope Options" in Chapter 15. 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; a zone is simply an area within a domain. Configuring a DNS server is explained in "Configuring a Primary DNS Server" in Chapter 16, "Optimizing DNS."

When you install the DNS Server service on an RODC, the RODC is able to pull a read-only replica of all application directory partitions that are used by DNS, including *ForestDNSZones* and *DomainDNSZones*. Clients can then query the RODC for name resolution as they would query any other DNS server. However, as with directory updates, the DNS server on an RODC does not support direct updates. This means that the RODC does not register name server (NS) resource records for any Active Directory-integrated zone that it hosts. When a client attempts to update its DNS records against an RODC, the server returns a referral to a DNS server that the client can use for the update. The DNS server on the RODC should receive the updated record from the DNS server that receives details about the update using a special replicate-single-object request that runs as a background process.

Windows 7 and later releases add support for DNS Security Extensions (DNSSEC). 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, this allows your DNS servers to securely sign zones and to host DNSSEC-signed zones. It also allows DNS servers to process related records and perform both validation and authentication.

## Using Windows Internet Name Service

WINS is a service that resolves computer names to IP addresses. Using WINS, the computer name COMPUTER84, for example, can be resolved to an IP address that enables computers on a Microsoft network to find one another and transfer information. WINS is needed to support pre-Windows 2000 systems and older applications that use NetBIOS over TCP/IP, such as the .NET command-line utilities. If you don't have pre-Windows 2000 systems or applications on the network, you don't need to use WINS.

WINS works best in client/server environments in which WINS clients send single-label (host) name queries to WINS servers for name resolution and WINS servers resolve the query and respond. When all your DNS servers are running Windows Server 2008 or later, deploying a Global Names zone creates static, global records with single-label names without relying on WINS. This allows users to access hosts using single-label names rather than FQDNs and removes the dependency on WINS. To transmit WINS queries and other information, computers use NetBIOS. NetBIOS provides an application programming interface (API) that allows computers on a network to communicate. NetBIOS applications rely on WINS or the local LMHOSTS file to resolve computer names to IP addresses. On pre-Windows 2000 networks, WINS is the primary name resolution service available. On Windows 2000 and later networks, DNS is the primary name resolution service and WINS has a different function. This function is to allow pre-Windows 2000 systems to browse lists of resources on the network and to allow Windows 2000 and later systems to locate NetBIOS resources.

To enable WINS name resolution on a network, you need to configure WINS clients and servers. When you configure WINS clients, you tell the clients the IP addresses for WINS servers on the network. Using the IP addresses, clients can communicate with WINS servers anywhere on the network, even if the servers are on different subnets. WINS clients can also communicate by using a broadcast method through which clients broadcast messages to other computers on the local network segment requesting their IP addresses. Because messages are broadcast, the WINS server isn't used. Any non-WINS clients that support this type of message broadcasting can also use this method to resolve computer names to IP addresses.

When clients communicate with WINS servers, they establish sessions that have the following three key parts:

- **Name registration** During name registration, the client gives the server its computer name and its IP address and asks to be added to the WINS database. If the specified computer name and IP address aren't already in use on the network, the WINS server accepts the request and registers the client in the WINS database.

- **Name renewal** Name registration isn't permanent. Instead, the client can use the name for a specified period known as a *lease*. The client is also given a time period within which the lease must be renewed, which is known as the renewal interval. The client must reregister with the WINS server during the renewal interval.
- **Name release** If the client can't renew the lease, the name registration is released, allowing another system on the network to use the computer name, IP address, or both. The names are also released when you shut down a WINS client.

After a client establishes a session with a WINS server, the client can request name-resolution services. The method used to resolve computer names to IP addresses depends on how the network is configured. The following four name-resolution methods are available:

- **B-node (broadcast)** Uses broadcast messages to resolve computer names to IP addresses. Computers that need to resolve a name broadcast a message to every host on the local network, requesting the IP address for a computer name. On a large network with hundreds or thousands of computers, these broadcast messages can use up valuable network bandwidth.
- **P-node (peer-to-peer)** Uses WINS servers to resolve computer names to IP addresses. As explained earlier, client sessions have three parts: name registration, name renewal, and name release. In this mode, when a client needs to resolve a computer name to an IP address, the client sends a query message to the server and the server responds with an answer.
- **M-node (mixed)** Combines b-node and p-node. With m-node, a WINS client first tries to use b-node for name resolution. If the attempt fails, the client then tries to use p-node. Because b-node is used first, this method has the same problems with network bandwidth usage as b-node.
- **H-node (hybrid)** Also combines b-node and p-node. With h-node, a WINS client first tries to use p-node for peer-to-peer name resolution. If the attempt fails, the client then tries to use broadcast messages with b-node. Because peer-to-peer is the primary method, h-node offers the best performance on most networks. H-node is also the default method for WINS name resolution.

If WINS servers are available on the network, Windows clients use the p-node method for name resolution. If no WINS servers are available on the network, Windows clients use the b-node method for name resolution. Windows computers can also use DNS and the local files LMHOSTS and HOSTS to resolve network names. Working with DNS is covered in detail in Chapter 16.

When you use DHCP to assign IP addresses dynamically, you should set the name resolution method for DHCP clients. To do this, you need to set DHCP scope options for the 046 WINS/NBT Node Type as specified in "Setting Scope Options" in Chapter 15. The best method to use is h-node. You'll get the best performance and have reduced traffic on the network.

## Using Link-Local Multicast Name Resolution

LLMNR fills a need for peer-to-peer name-resolution services for devices with an IPv4 address, an IPv6 address, or both, allowing IPv4 and IPv6 devices on a single subnet without a WINS or DNS server to resolve each other's names—a service that neither WINS nor DNS can fully provide. Although WINS can provide both client/server and peer-to-peer name-resolution services for IPv4, it does not support IPv6 addresses. DNS, on the other hand, supports IPv4 and IPv6 addresses, but it depends on designated servers to provide name-resolution services.

Windows 7 and later releases support LLMNR. LLMNR is designed for both IPv4 and IPv6 clients in configurations where other name-resolution systems are not available, such as

- Home or small office networks
- Ad hoc networks
- Corporate networks where DNS services are not available

LLMNR is designed to complement DNS by enabling name resolution in scenarios in which conventional DNS name resolution is not possible. Although LLMNR can replace the need for WINS in cases where NetBIOS is not required, LLMNR is not a substitute for DNS because it operates only on the local subnet. Because LLMNR traffic is prevented from propagating across routers, it cannot accidentally flood the network.

As with WINS, you use LLMNR to resolve a host name, such as COMPUTER84, to an IP address. By default, LLMNR is enabled on all computers running Windows 7 and later releases, and these computers use LLMNR only when all attempts to look up a host name through DNS fail. As a result, name resolution works like the following for Windows 7 and later releases:

1. A host computer sends a query to its configured primary DNS server. If the host computer does not receive a response or receives an error, it tries each configured alternate DNS server in turn. If the host has no configured DNS servers or fails to connect to a DNS server without errors, name resolution fails over to LLMNR.
2. The host computer sends a multicast query over User Datagram Protocol (UDP) requesting the IP address for the name being looked up. This query is restricted to the local subnet (also referred to as the *local link*).
3. Each computer on the local link that supports LLMNR and is configured to respond to incoming queries receives the query and compares the name to its own host name. If the host name is not a match, the computer discards the query. If the host name is a match, the computer transmits a unicast message containing its IP address to the originating host.

You can also use LLMNR for reverse mapping. With a reverse mapping, a computer sends a unicast query to a specific IP address, requesting the host name of the target computer. An LLMNR-enabled computer that receives the request sends a unicast reply containing its host name to the originating host.

LLMNR-enabled computers are required to ensure that their names are unique on the local subnet. In most cases, a computer checks for uniqueness when it

starts, when it resumes from a suspended state, and when you change its network-interface settings. If a computer has not yet determined that its name is unique, it must indicate this condition when responding to a name query.

**REAL WORLD** By default, LLMNR is automatically enabled on computers running Windows 7 and later releases. You can disable LLMNR through registry settings. To disable LLMNR for all network interfaces, create and set the following registry value to 0: HKLM/SYSTEM/CurrentControlSet/Services/Dnscache/Parameters/EnableMulticast. To disable LLMNR for a specific network interface, create and set the following registry value to 0: HKLM/SYSTEM/CurrentControlSet/Services/Tcpip/Parameters/AdapterGUID/EnableMulticast.

Here, *AdapterGUID* is the globally unique identifier (GUID) of the network-interface adapter for which you want to disable LLMNR. You can enable LLMNR again at any time by setting these registry values to 1. You also can manage LLMNR through Group Policy.

## Frequently Used Tools

---

Many utilities are available for administrating Windows Server 2012 systems. The tools you use the most include the following:

- **Control Panel** A collection of tools for managing system configuration. You can organize Control Panel in different ways according to the view you're using. A view is simply a way of organizing and presenting options. You change the view by using the View By list. Category view is the default view, and it provides access to tools by category, tool, and key tasks. The Large Icons and Small Icons views are alternative views that list each tool separately by name.
- **Graphical administrative tools** The key tools for managing network computers and their resources. You can access these tools by selecting them individually from the Administrative Tools program group.
- **Administrative wizards** Tools designed to automate key administrative tasks. You can access many administrative wizards in Server Manager—the central administration console for Windows Server 2012.
- **Command-line utilities** You can launch most administrative utilities from the command prompt. In addition to these utilities, Windows Server 2012 provides others that are useful for working with Windows Server 2012 systems.

To learn how to use any of the .NET command-line tools, type **NET HELP** at a command prompt followed by the command name, such as **NET HELP SHARE**. Windows Server 2012 then provides an overview of how the command is used.

## Windows PowerShell 3.0

For additional flexibility in your command-line scripting, you might want to use Windows PowerShell 3.0. Windows PowerShell 3.0 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 with existing PowerShell host applications, and Windows PowerShell Web Access, which lets a server act as a web gateway for managing the server remotely using PowerShell and a web client.

**REAL WORLD** You can install these additional Windows PowerShell features using the Add Roles And Features Wizard. On the desktop, tap or click the Server Manager button on the taskbar. This option is included by default. In Server Manager, tap or click Manage and then tap or click Add Roles And Features. This runs the Add Roles And Features Wizard, which you use to add these features. Note, however, that with Windows Server 2012, not only can you disable a role or feature, but you also can remove the binaries needed for that role or feature. Binaries needed to install roles and features are referred to as *payloads*.

The Windows PowerShell console (*Powershell.exe*) is a 32-bit or 64-bit environment for working with Windows PowerShell at the command line. On 32-bit versions of Windows, you'll find the 32-bit executable in the %SystemRoot%\System32\WindowsPowerShell\v1.0 directory. On 64-bit versions of Windows, you'll find the 32-bit executable in the %SystemRoot%\SysWow64\WindowsPowerShell\v1.0 directory, and the 64-bit executable in the %SystemRoot%\System32\WindowsPowerShell\v1.0 directory.

On the desktop, you can open the Windows PowerShell console by tapping or clicking the PowerShell button on the taskbar. This option is included by default. On 64-bit systems, the 64-bit version of PowerShell is started by default. If you want to use the 32-bit PowerShell console on a 64-bit system, you must select the Windows PowerShell (x86) option.

You can start Windows PowerShell from a Windows command shell (*Cmd.exe*) by entering the following:

```
powershell
```

**NOTE** The directory path for Windows PowerShell should be in your command path by default. This ensures that you can start Windows PowerShell from a command prompt without first having to change to the related directory.

After starting Windows 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 can also execute cmdlets in scripts. Cmdlets are named 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 Windows PowerShell prompt, you can get a complete list of cmdlets by typing `get-help *-*`. To get help documentation on a specific cmdlet, type `get-help` followed by the cmdlet name, such as `get-help get-variable`.

All cmdlets also have configurable aliases that act as shortcuts for executing a cmdlet. To list all aliases available, type `get-item -path 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

```
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.

**REAL WORLD** Generally speaking, anything you can type at a command prompt can be typed at the PowerShell prompt as well. This is possible because PowerShell looks for external commands and utilities as part of its normal 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, and 5) external commands, utilities, and files. Thus, if any element in 1 to 4 of the execution order has the same name as a command, that element will run instead of the expected command.

## Windows Remote Management

The Windows PowerShell remoting features are supported by the `WS-Management` protocol and the Windows Remote Management (WinRM) service that implements `WS-Management` in Windows. Computers running Windows 7 and later, as well as Windows Server 2008 R2 or later include WinRM 2.0 or later. If you want to manage

a Windows server from a workstation, you need to be sure that WinRM 2.0 and Windows PowerShell 3.0 are installed and that the server has a WinRM listener enabled. An IIS extension, installable as a Windows feature called WinRM IIS Extension, lets a server act as a web gateway for managing the server remotely using WinRM and a web client.

### Enabling and Using WinRM

You can verify the availability of WinRM 2.0 and configure Windows PowerShell for remoting by following these steps:

1. Tap or click Start, and then point to Windows PowerShell. Start Windows PowerShell as an administrator by pressing and holding or right-clicking the Windows PowerShell shortcut and selecting Run As Administrator.
2. The WinRM service is configured for manual startup 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 configure it to start automatically in the future:

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

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

```
Enable-PSRemoting -force
```

You can enable remoting only when your computer is connected to a domain or a private network. If your computer is connected to a public network, you need to disconnect from the public network and connect to a domain or private network and then repeat this step. If one or more of your computer's connections has the Public Network connection type but you are actually connected to a domain or private network, you need to change the network connection type in Network And Sharing Center and then repeat this step.

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, 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, type the following:

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



Here *RemoteComputer* is the name of the remote computer, such as

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

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 see 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 see 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.

Whenever you use Windows PowerShell remoting features, you must start Windows PowerShell as an administrator by pressing and holding or right-clicking the Windows PowerShell shortcut and selecting Run As Administrator. When starting Windows PowerShell from another program, such as the command prompt, you must start that program as an administrator.

## Configuring WinRM

When you are working with an elevated, administrator command prompt, you can use the WinRM command-line utility to view and manage the remote management configuration. Type **winrm get winrm/config** to display detailed information about the remote management configuration.

If you examine the configuration listing, you'll notice there is a hierarchy of information. The base of this hierarchy, the Config level, is referenced with the path `winrm/config`. Then there are sublevels for client, service, and WinRS, referenced as `winrm/config/client`, `winrm/config/service`, and `winrm/config/winrs`. You can change the value of most configuration parameters by using the following command:

```
winrm set ConfigPath @{ParameterName="Value"}
```

Here *ConfigPath* is the configuration path, *ParameterName* is the name of the parameter you want to work with, and *Value* sets the value for the parameter, such as

```
winrm set winrm/config/winrs @{MaxShellsPerUser="10"}
```

Here, you set the *MaxShellsPerUser* parameter under `winrm/config/winrs`. This parameter controls the maximum number of connections to a remote computer that can be active per user. (By default, each user can have only five active connections.) Keep in mind that some parameters are read-only and cannot be set in this way.

WinRM requires at least one listener to indicate the transports and IP addresses on which management requests can be accepted. The transport must be HTTP, HTTPS, or both. With HTTP, messages can be encrypted using NTLM or Kerberos encryption. With HTTPS, Secure Sockets Layer (SSL) is used for encryption. You can examine the configured listeners by typing **winrm enumerate winrm/config/listener**. As Listing 1-1 shows, this command displays the configuration details for configured listeners.

**LISTING 1-1** Sample Configuration for Listeners

---

```
Listener
  Address = *
  Transport = HTTP
  Port = 80
  Hostname
  Enabled = true
  URLPrefix = wsman
  CertificateThumbprint
  ListeningOn = 127.0.0.1, 192.168.1.225
```

By default, your computer is probably configured to listen on any IP address. If so, you won't see any output. To limit WinRM to specific IP addresses, the computer's local loopback address (127.0.0.1) and assigned IPv4 and IPv6 addresses can be explicitly configured for listening. You can configure a computer to listen for requests over HTTP on all configured IP addresses by typing the following:

```
winrm create winrm/config/listener?Address=*+Transport=HTTP
```

You can listen for requests over HTTPS on all IP addresses configured on the computer by typing this:

```
winrm create winrm/config/listener?Address=*+Transport=HTTPS
```

Here, the asterisk (\*) indicates all configured IP addresses. Note that the *CertificateThumbprint* property must be empty to share the SSL configuration with another service.

You can enable or disable a listener for a specific IP address by typing

```
winrm set winrm/config/listener?Address=IP:192.168.1.225+Transport=HTTP
@{Enabled="true"}
```

or

```
winrm set winrm/config/listener?Address=IP:192.168.1.225+Transport=HTTP
@{Enabled="false"}
```

You can enable or disable basic authentication on the client by typing

```
winrm set winrm/config/client/auth @{Basic="true"}
```

or

```
winrm set winrm/config/client/auth @{Basic="false"}
```

You can enable or disable Windows authentication using either NTLM or Kerberos (as appropriate) by typing

```
winrm set winrm/config/client @{TrustedHosts="<local>"}
```

or

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

In addition to managing WinRM at the command line, you can manage the service by using Group Policy. As a result, Group Policy settings might override any settings you enter.



# Managing Servers Running Windows Server 2012

- Server Roles, Role Services, and Features for Windows Server 2012 **32**
- Full-Server, Minimal-Interface, and Server Core Installations **40**
- Installing Windows Server 2012 **43**
- Managing Roles, Role Services, and Features **57**
- Managing System Properties **73**

Servers are the heart of any Microsoft Windows network. One of your primary responsibilities as an administrator is to manage these resources. Windows Server 2012 comes with several integrated management tools. The one you'll use for handling core system administration tasks is Server Manager. Server Manager provides setup and configuration options for the local server as well as 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

- 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 fine control over system environment settings and properties. This is where the System utility comes into the picture. You can use this utility to do the following:

- Change a computer's name
- Configure application performance, virtual memory, and registry settings
- Manage system and user environment variables
- Set system startup and recovery options

# Server Roles, Role Services, and Features for Windows Server 2012

---

Windows Server 2012 uses the same configuration architecture as Windows Server 2008 and Windows Server 2008 Release 2 (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 Service (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, a Microsoft Management Console (MMC). 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. Table 2-1 provides an overview of the primary roles and the related role services you can deploy on a server running Windows Server 2012.

**TABLE 2-1** Primary Roles and Related Role Services for Windows Server 2012

<b>ROLE</b>	<b>DESCRIPTION</b>
Active Directory Certificate Services (AD CS)	Provides functions necessary for issuing and revoking digital certificates for users, client computers, and servers. Includes these role services: Certification Authority, Certification Authority Web Enrollment, Online Responder, Network Device Enrollment Service, Certificate Enrollment Web Service, and Certificate Enrollment Policy Web Service.
Active Directory Domain Services (AD DS)	Provides functions necessary for storing information about users, groups, computers, and other objects on the network, and makes this information available to users and computers. Active Directory domain controllers give network users and computers access to permitted resources on the network.
Active Directory Federation Services (AD FS)	Complements the authentication and access management features of AD DS by extending them to the World Wide Web. Includes these role services and subservices: Federation Service, Federation Service Proxy, AD FS Web Agents, Claims-Aware Agent, and Windows Token-Based Agent.
Active Directory Lightweight Directory Services (AD LDS)	Provides a data store for directory-enabled applications that do not require AD DS and do not need to be deployed on domain controllers. Does not include additional role services.
Active Directory Rights Management Services (AD RMS)	Provides controlled access to protected email messages, documents, intranet pages, and other types of files. Includes these role services: Active Directory Rights Management Server and Identity Federation Support.
Application Server	Allows a server to host distributed applications built using ASP.NET, Enterprise Services, and Microsoft .NET Framework 4.5. Includes more than a dozen role services.
DHCP Server	DHCP provides centralized control over IP addressing. DHCP servers can assign dynamic IP addresses and essential TCP/IP settings to other computers on a network. Does not include additional role services.
DNS Server	DNS is a name-resolution system that resolves computer names to IP addresses. DNS servers are essential for name resolution in Active Directory domains. Does not include additional role services.

ROLE	DESCRIPTION
Fax Server	Provides centralized control over sending and receiving faxes in the enterprise. A fax server can act as a gateway for faxing and allows you to manage fax resources, such as jobs and reports, and fax devices on the server or on the network. Does not include additional role services.
File And Storage Services	Provides essential services for managing files and storage, and the way they are made available and replicated on the network. A number of server roles require some type of file service. Includes these role services and subservices: BranchCache for Network Files, Data Deduplication, Distributed File System, DFS Namespaces, DFS Replication, File Server, File Server Resource Manager, Services for Network File System (NFS), File Server VSS Agent Service, iSCSI Target Server, iSCSI Target Storage Provider, and Storage Services.
Hyper-V	Provides services for creating and managing virtual machines that emulate physical computers. Virtual machines have separate operating system environments from the host server.
Network Policy and Access Services (NPAS)	Provides essential services for managing network access policies. Includes these role services: Network Policy Server (NPS), Health Registration Authority (HRA), and Host Credential Authorization Protocol (HCAP).
Print And Document Services	Provides essential services for managing network printers, network scanners, and related drivers. Includes these role services: Print Server, LPD Service, Internet Printing, and Distributed Scan Server.
Remote Access	Provides services for managing routing and remote access to networks. Use this role if you need to configure Virtual Private Networks (VPN), Network Address Translation (NAT), and other routing services. Includes these role services: DirectAccess and VPN (RAS) and Routing.
Remote Desktop Services	Provides services that allow users to run Windows-based applications that are installed on a remote server. When users run an application on a terminal server, the execution and processing occur on the server and only the data from the application is transmitted over the network.
Volume Activation Services	Provides services for automating the management of volume license keys and volume key activation.



ROLE	DESCRIPTION
Web Server (IIS)	Used to host websites and web-based applications. Websites hosted on a web server can have both static content and dynamic content. You can build web applications hosted on a web server by using ASP.NET and .NET Framework 4.5. When you deploy a web server, you can manage the server configuration using IIS 8 modules and administration tools. Includes several dozen role services.
Windows Deployment Services (WDS)	Provides services for deploying Windows computers in the enterprise. Includes these role services: Deployment Server and Transport Server.
Windows Server Update Services (WSUS)	Provides services for Microsoft Update, allowing you to distribute updates from designated servers.

Table 2-2 provides an overview of the primary features you can deploy on a server running Windows Server 2012. Unlike early releases of Windows, Windows Server 2012 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.

**TABLE 2-2** Primary Features for Windows Server 2012

FEATURE	DESCRIPTION
Background Intelligent Transfer Service (BITS)	Provides intelligent background transfers. When this feature is installed, the server can act as a BITS server that can receive file uploads from clients. This feature isn't necessary for downloads to clients using BITS. Additional subfeatures include BITS IIS Server Extension and BITS Compact Server.
BitLocker Drive Encryption	Provides hardware-based security to protect data through full-volume encryption that prevents disk tampering while the operating system is offline. Computers that have Trusted Platform Module (TPM) can use BitLocker Drive Encryption in Startup Key or TPM-Only mode. Both modes provide early integrity validation.
BitLocker Network Unlock	Provides support for network-based key protectors that automatically unlock BitLocker-protected operating system drives when a domain-joined computer is restarted.

FEATURE	DESCRIPTION
BranchCache	Provides services needed for BranchCache client and server functionality. Includes HTTP protocol, Hosted Cache, and related services.
Client for NFS	Provides functionality for accessing files on UNIX-based NFS servers.
Data Center Bridging	Supports a suite of IEEE standards for enhancing LANs and enforcing bandwidth allocation.
Enhanced Storage	Provides support for Enhanced Storage Devices.
Failover Clustering	Provides clustering functionality that allows multiple servers to work together to provide high availability for services and applications. Many types of services can be clustered, including file and print services. Messaging and database servers are ideal candidates for clustering.
Group Policy Management	Installs the Group Policy Management Console (GPMC), which provides centralized administration of Group Policy.
Ink and Handwriting Services	Provides support for use of a pen or stylus and handwriting recognition.
IP Address Management Server	Provides support for central management of the enterprise's IP address space and the related infrastructure servers.
Internet Printing Client	Provides functionality that allows clients to use HTTP to connect to printers on web print servers.
Internet Storage Naming Server (iSNS) Server Service	Provides management and server functions for Internet SCSI (iSCSI) devices, allowing the server to process registration requests, deregistration requests, and queries from iSCSI devices.
LPR Port Monitor	Installs the LPR Port Monitor, which allows printing to devices attached to UNIX-based computers.
Media Foundation	Provides essential functionality for Windows Media Foundation.
Message Queuing	Provides management and server functions for distributed message queuing. A group of related subfeatures is available as well.
Multipath I/O (MPIO)	Provides functionality necessary for using multiple data paths to a storage device.

FEATURE	DESCRIPTION
.NET Framework 4.5	Provides APIs for application development. Additional subfeatures include .NET Framework 4.5, ASP.NET 4.5, and Windows Communication Foundation (WCF) Activation Components.
Network Load Balancing (NLB)	NLB provides failover support and load balancing for IP-based applications and services by distributing incoming application requests among a group of participating servers. Web servers are ideal candidates for load balancing.
Peer Name Resolution Protocol (PNRP)	Provides Link-Local Multicast Name Resolution (LLMNR) functionality that allows peer-to-peer, name-resolution services. When you install this feature, applications running on the server can use LLMNR to register and resolve names.
Quality Windows Audio Video Experience	A networking platform for audio video (AV) streaming applications on IP home networks.
RAS Connection Manager Administration Kit	Provides the framework for creating profiles for connecting to remote servers and networks.
Remote Assistance	Allows a remote user to connect to the server to provide or receive Remote Assistance.
Remote Differential Compression	Provides support for differential compression by determining which parts of a file have changed and replicating only the changes.
Remote Server Administration Tools (RSAT)	Installs role-management and feature-management tools that can be used for remote administration of other Windows Server systems. Options for individual tools are provided, or you can install tools by top-level category or subcategory.
Remote Procedure Call (RPC) over HTTP Proxy	Installs a proxy for relaying RPC messages from client applications to the server over HTTP. RPC over HTTP is an alternative to having clients access the server over a VPN connection.
Simple TCP/IP Services	Installs additional TCP/IP services, including Character Generator, Daytime, Discard, Echo, and Quote of the Day.
Simple Mail Transfer Protocol (SMTP) Server	SMTP is a network protocol for controlling the transfer and routing of email messages. When this feature is installed, the server can act as a basic SMTP server. For a full-featured solution, you need to install a messaging server, such as Microsoft Exchange Server.

FEATURE	DESCRIPTION
Simple Network Management Protocol (SNMP) Services	SNMP is a protocol used to simplify management of TCP/IP networks. You can use SNMP for centralized network management if your network has SNMP-compliant devices. You can also use SNMP for network monitoring via network management software.
Subsystem for UNIX-Based Applications (SUA)	Provides functionality for running UNIX-based programs. You can download additional management utilities from the Microsoft website. (Deprecated)
Telnet Client	Allows a computer to connect to a remote Telnet server and run applications on that server.
Telnet Server	Hosts the remote sessions for Telnet clients. When Telnet Server is running on a computer, users can connect to the server with a Telnet client from a remote computer.
User Interfaces And Infrastructure	Allows you to control the user experience and infrastructure options (Graphical Management Tools And Infrastructure, Desktop Experience, or Server Graphical Shell).
Windows Biometric Framework	Provides functionality required for using fingerprint devices.
Windows Internal Database	Allows the server to use relational databases with Windows roles and features that require an internal database, such as AD RMS, UDDI Services, WSUS, Windows SharePoint Services, and Windows System Resource Manager.
Windows PowerShell	Allows you to manage the Windows PowerShell features of the server. Windows PowerShell 3.0 and the PowerShell ISE are installed by default.
Windows PowerShell Web Access	Allows the server to act as a web gateway for remotely managing servers in a web browser.
Windows Process Activation Service	Provides support for distributed, web-based applications that use HTTP and non-HTTP protocols.
Windows Standards-Based Storage Management	Provides support for managing standards-based storage and includes management interfaces as well as extensions for WMI and Windows PowerShell.
Windows Server Backup	Allows you to back up and restore the operating system, system state, and any data stored on a server.
Windows System Resource Manager (WSRM)	Allows you to manage resource usage on a per-processor basis. (Deprecated)

FEATURE	DESCRIPTION
Windows TIFF IFilter	Focuses on text-based documents, which means that searching is more successful for documents that contain clearly identifiable text (for example, black text on a white background).
WinRM IIS Extension	Provides an Internet Information Services (IIS)-based hosting model. WinRM IIS Extension can be enabled at either the website or virtual-directory level.
WINS Server	A name-resolution service that resolves computer names to IP addresses. Installing this feature allows the computer to act as a WINS server.
Wireless LAN Service	Allows the server to use wireless networking connections and profiles.
WoW64 Support	Supports WoW64, which is required on a full-server installation. Removing this feature converts a full-server installation to a Server Core installation.
XPS Viewer	A program you can use to view, search, set permissions for, and digitally sign XPS documents.

**NOTE** Desktop Experience is now a subfeature of the top-level feature called User Interfaces And Infrastructure. Desktop Experience provides Windows desktop functionality on the server. Windows features added include Windows Media Player, desktop themes, Video for Windows (AVI support), Windows Defender, 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.

As an administrator, you might be asked to install or uninstall dynamic-link libraries (DLLs), particularly if you work with IT development teams. The utility you use to work with DLLs is Regsvr32. This utility is run at the command line.

After you open a Command Prompt window, you install or register a DLL by typing **regsvr32 name.dll**—for example:

```
regsvr32 mylibs.dll
```

If necessary, you can uninstall or unregister a DLL by typing **regsvr32 /u name.dll**—for example:

```
regsvr32 /u mylibs.dll
```

Windows File Protection prevents the replacement of protected system files. You can replace only DLLs installed by the Windows Server operating system as part of a hotfix, service pack update, Windows update, or Windows upgrade. Windows File Protection is an important part of the Windows Server security architecture.

# Full-Server, Minimal-Interface, and Server Core Installations

---

Windows Server 2012 supports full-server, minimal-interface, and Server Core installations. Full-server installations, also referred to as Server With A GUI Installations, have the Graphical Management Tools And Infrastructure and Server Graphical Shell features (which are part of the User And Infrastructure feature) and the WoW64 Support framework installed. Minimal-interface installations, also referred to as Server With Minimal Interface Installations, are full-server installations with the Server Graphical Shell removed. Server Core installations have a limited user interface and do not include any of the User Interfaces And Infrastructure features or the WoW64 Support framework.

As discussed in “Changing the Installation Type” later in the chapter, the installation type can be changed at any time. With a full-server installation, you have a complete working version of Windows Server 2012 you can deploy with any permitted combination of roles, role services, and features. With a minimal-interface installation, you also can deploy any permitted combination of roles, role services, and features. However, with a Server Core installation, you have a minimal installation of Windows Server 2012 that supports a limited set of roles and role combinations. The supported roles include AD CS, AD DS, AD LDS, DHCP Server, DNS Server, File Services, Hyper-V, Media Services, Print And Document Services, Routing And Remote Access Server, Streaming Media Services, Web Server (IIS), and Windows Server Update Server. In its current implementation, a Server Core installation is not a platform for running server applications.

While all three installation types use the same licensing rules and can be managed remotely using any available and permitted remote-administration technique, full-server, minimal-interface, and Server Core installations are completely different when it comes to local console administration. With a full-server installation, you’re provided with a user interface that includes a full desktop environment for local console management of the server. With a minimal interface, you have only Microsoft Management Consoles, Server Manager, and a subset of Control Panel available for management tasks. Missing from both a minimal-interface installation and a Server Core installation are File Explorer, taskbar, notification area, Internet Explorer, built-in help system, themes, Metro-style apps, and Windows Media Player.

## Navigating Server Core

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

- 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 (Timedate.cpl) for viewing or setting the date, time, and time zone.
- Region And Language control panel (Intl.cpl) for viewing or setting regional and language options, including formats and the keyboard layout.
- Server Configuration utility (Sconfig), 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 log on just as you do with a full-server installation. 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 this command prompt for administration of the server. If you accidentally close the command prompt, you can open a new command prompt by following these steps:

1. Press Ctrl+Shift+Esc to display Task Manager.
2. On the File menu, tap or click Run New Task.
3. In the Create New Task dialog box, type **cmd** in the Open text box, and then tap or click OK.

You can use this technique to open additional Command Prompt windows as well. Although you can work with Notepad and Regedit by typing **notepad.exe** or **regedit.exe** instead of **cmd**, 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) 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
- Log off, restart, or shut down

When you are logged on, you can display the Windows Logon screen at any time by pressing Ctrl+Alt+Delete. In a Server Core installation, the Windows Logon screen has the same options as with a full-server installation, allowing you to lock the computer, switch users, log off, change a password, or start Task Manager. 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 if all of their dependencies are available in the Server Core installation.

Although a Server Core installation supports a limited set of roles and role services, you can install most features. Windows Server 2012 also supports the .NET Framework, Windows PowerShell 3.0, and Windows Remote Management (WinRM) 2.0. This support allows you to perform local and remote administration using PowerShell. You also can use Remote Desktop Services to manage a Server Core installation remotely. Some of the common tasks you might want to perform when you are logged on locally are summarized in Table 2-3.

**TABLE 2-3** Helpful Commands and Utilities for Managing Server Core Installations

COMMAND	TASK
Cscript Scregedit.wsf	Configure the operating system. Use the <i>/cli</i> parameter to list available configuration areas.
Diskraid.exe	Configure software RAID.
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. Type <b>netsh interface ipv4</b> to configure IPv4 settings. Type <b>netsh interface ipv6</b> to configure IPv6 settings.
Ocsetup.exe	Add or remove roles, role services, and features.
Pnputil.exe	Install or update hardware device drivers.
Sc query type=driver	List installed device drivers.
Serverwerroptin.exe	Configure Windows Error Reporting.
Slmgr -ato	Windows Software Licensing Management tool used to activate the operating system. Runs <i>Cscript slmgr.vbs -ato</i> .
Slmgr -ipk	Install or replace the product key. Runs <i>Cscript slmgr.vbs -ipk</i> .



COMMAND	TASK
SystemInfo	List the system configuration details.
Wevtutil.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 <i>Cscript winrm.vbs quickconfig</i> . Enter without the <i>quickconfig</i> 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 Microsoft 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.

## Installing Windows Server 2012

You can install Windows Server 2012 on new hardware or as an upgrade. When you install Windows Server 2012 on a computer with an existing operating system, you can perform a clean installation or an upgrade. With a clean installation, the Windows Server 2012 Setup program replaces the original operating system on the computer, and all user or application settings are lost. With an upgrade, the Setup program performs a clean installation of the operating system and then migrates user settings, documents, and applications from the earlier version of Windows.

Windows Server 2012 supports only 64-bit architecture. You can install the operating system only on computers with 64-bit processors. Before you install Windows Server 2012, you should be sure that your computer meets the minimum

requirements of the edition you plan to use. Microsoft provides both minimum requirements and recommended requirements. If your computer doesn't meet the minimum requirements, you will not be able to install Windows Server 2012. If your computer doesn't meet the recommended requirements, you will experience performance issues.

Windows Server 2012 requires at least 10 GB of disk space for installation of the base operating system. Microsoft recommends that a computer running Windows Server 2012 have 32 GB or more of available disk space. Additional disk space is required for paging and dump files as well as for the features, roles, and role services you install. For optimal performance, you should have at least 10 percent of free space on a server's disks at all times.

When you install Windows Server 2012, the Setup program 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 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.

## 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 boot the computer using the distribution media. If you don't boot using the distribution media, these options won't be available, and you'll only be able to manage disk partitions at a command prompt using the DiskPart utility.

You can perform a clean installation of Windows Server 2012 by following these steps:

1. Start the Setup program by using one of the following techniques:
  - For a new installation, turn on the computer with the Windows Server 2012 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 select advanced boot options and then boot from media rather than hard disk, or you might need to change the computer's firmware settings to allow booting from media.
  - For a clean installation over an existing installation, you can boot from the distribution media, or you can start the computer and log on using an account with administrator privileges. When you insert the Windows Server 2012 distribution media into the computer's disc drive, Setup should start automatically. If Setup doesn't start automatically, use File Explorer to access the distribution media and then double-tap or double-click Setup.exe.
2. If you started the computer using the distribution media, 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 2012 you are installing are different, you might see unexpected characters as you type. Be sure that you select the correct keyboard language to avoid this. When you are ready to continue with the installation, tap or click Next.

3. Choose Install Now to start the installation. After Setup copies the temporary files to the computer, choose whether to get updates for Setup during the installation. If you started Setup after logging on to an existing installation of Windows, choose either Go Online To Install Updates Now or No, Thanks.
4. With volume and enterprise licensed editions of Windows Server 2012, you might not need to provide a product key during installation. With retail editions, however, you need to enter a product key when prompted. Tap or click Next to continue. The Activate Windows When I'm Online check box is selected by default to ensure that you are prompted to activate the operating system the next time you connect to the Internet.

**NOTE** You must activate Windows Server 2012 after installation. If you don't activate Windows Server 2012 in the allotted time, you see an error stating "Your activation period has expired" or that you have a "Non-genuine version of Windows Server 2012 installed." Windows Server 2012 will then run with reduced functionality. You need to activate and validate Windows Server 2012 as necessary to regain full functionality.

5. On the Select The Operating System You Want To Install page, options are provided for full-server and Server Core installations. Make the appropriate selection, and then tap or click Next.
6. The license terms for Windows Server 2012 have changed from previous releases of Windows. After you review the license terms, tap or click I Accept The License Terms, and then tap or click Next.
7. On the Which Type Of Installation Do You Want page, select the type of installation you want Setup to perform. Because you are performing a clean installation to replace an existing installation or configure a new computer, select Custom Install Windows Only (Advanced) as the installation type. If you started Setup from the boot prompt rather than from Windows itself, the Upgrade option is disabled. To upgrade rather than perform a clean install, you need to restart the computer and boot the currently installed operating system. After you log on, you then need to start the installation.
8. On the Where Do You Want To Install Windows page, select the disk or disk and partition on which you want to install the operating system. There are two versions of the Where Do You Want To Install Windows 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 tap or 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 “Creating, Formatting, Deleting, and Extending Disk Partitions During Installation” 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 “Creating, Formatting, Deleting, and Extending Disk Partitions During Installation” 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 “Creating, Formatting, Deleting, and Extending Disk Partitions During Installation” later in this chapter.
9. 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. Tap or click **OK**.
  10. Tap or click **Next**. Setup starts the installation of the operating system. During this procedure, Setup copies the full disk image of Windows Server 2012 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.

**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. Once networking is configured correctly, use `Simgmgr -ipk` to set the product key and `Simgmgr -ato` to activate Windows. Enter `timedate.cpl` to set the server’s date and time. If you want to enable remote management 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 /?**.

## Performing an Upgrade Installation

Although Windows Server 2012 provides an upgrade option during installation, an upgrade isn't what you think it is. 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.

During the migration portion of the upgrade, Setup moves folders and files from the previous installation to a folder named `Windows.old`. As a result, the previous installation will no longer run.

**NOTE** You cannot perform an upgrade installation of Windows Server 2012 on a computer with a 32-bit operating system, even if the computer has 64-bit processors. You need to migrate the services being provided by the computer to other servers and then perform a clean installation. The Windows Server Migration tools might be able to help you migrate your server. These tools are available on computers running Windows Server 2012.

You can perform an upgrade installation of Windows Server 2012 by following these steps:

1. Start the computer, and log on using an account with administrator privileges. When you insert the Windows Server 2012 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-tap or 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 2012 you are installing are different, you might see unexpected characters as you type.
3. Choose **Install Now** to start the installation. After Setup copies the temporary files to the computer, choose whether to get updates during the installation. Choose either **Go Online To Install Updates Now** or **No, Thanks**.
4. With volume and enterprise licensed editions of Windows Server 2012, 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. Tap or click **Next** to continue. The **Automatically Activate Windows When I'm Online** check box is selected by default to ensure that you are prompted to activate the operating system the next time you connect to the Internet.

5. On the Select The Operating System You Want To Install page, options are provided for full-server and Server Core installations. Make the appropriate selection, and then tap or click Next.
6. The license terms for Windows Server 2012 have changed from previous releases of Windows. After you review the license terms, tap or click I Accept The License Terms, and then tap or click Next.
7. On the Which Type Of Installation Do You Want page, you need to select the type of installation you want Setup to perform. Because you are performing a clean installation over an existing installation, select Upgrade. If you started Setup from the boot prompt rather than from Windows itself, the Upgrade option is disabled. To upgrade rather than perform a clean install, you need to restart the computer and boot the currently installed operating system. After you log on, you can start the installation.
8. 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 2012 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.

## Performing Additional Administration Tasks During Installation

Sometimes you might forget to perform a preinstallation task prior to starting the installation. Rather than restarting the operating system, you can access a command prompt from Setup or use advanced drive options to perform the necessary administrative tasks.

### Using the Command Line During Installation

When you access a command prompt from Setup, you access the MINWINPC (mini Windows PC) environment 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. As Table 2-4 shows, the mini Windows PC environment gives you access to many of the same command-line tools that are available in a standard installation of Windows Server 2012.

**TABLE 2-4** Command-Line Utilities in the Mini Windows PC Environment

COMMAND	DESCRIPTION
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.

COMMAND	DESCRIPTION
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 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 floppy disk or hard drive.
FTP	Transfers files.
FTYPE	Displays or modifies file types used in file-extension associations.

COMMAND	DESCRIPTION
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.



<b>COMMAND</b>	<b>DESCRIPTION</b>
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 Windows command prompt.
PUSHD	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.

COMMAND	DESCRIPTION
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 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, type **diskpart**. This starts the DiskPart utility.
3. To view a list of disks on the computer, type **list disk**.
4. Select a disk by typing **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, type **clean**.
6. When the cleaning process is finished, type **exit** to exit the DiskPart utility.
7. Type **exit** to exit the command prompt.
8. In the Install Windows dialog box, tap or click the back arrow button to return to the previous window.

9. On the Which Type Of Installation Do You Want page, tap or click Custom (Advanced) to start a custom install.
10. On the Where Do You Want To Install Windows page, tap or click the disk you previously cleaned to select it as the installation partition. As necessary, tap or click the Disk Options link to display the Delete, Format, New, and Extend partition configuration options.
11. Tap or click New. In the Size box, set the size of the partition in megabytes, and then tap or click Apply.

## Loading Disk 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 a hard-disk drive or a hard-disk controller. Typically, you use this option when a disk drive you want to use for installing the operating system isn't available for selection because the device drivers aren't available.

To load the device drivers and make the hard disk available, follow these steps:

1. During installation, on the Where Do You Want To Install Windows page, tap or click Load Driver.
2. When prompted, insert the installation media into a DVD drive or USB flash drive, and then tap or 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 tap or click Next.
  - If Setup doesn't find the device driver, tap or click Browse to use the Browse For Folder dialog box to select the device driver to load, tap or click OK, and then tap or click Next.

You can tap or click the Rescan button to have Setup rescan the computer's removable media drives for the device drivers. If you are unable to install a device driver successfully, tap or click the back arrow button in the upper-left corner of the Install Windows dialog box to go back to the previous page.

## Creating, Formatting, Deleting, and Extending 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 has additional options. You can display these options by tapping or clicking Drive Options (Advanced). These additional options are used as follows:

- **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, type **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. Once 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, tap or click *Drive Options (Advanced)* to display the advanced options for working with drives.
2. Tap or click the disk on which you want to create the partition, and then tap or click *New*.
3. In the *Size* box, set the size of the partition in megabytes and then tap or 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, tap or click *Drive Options (Advanced)* to display the advanced options for working with drives.
2. Tap or click the partition that you want to format.
3. Tap or click *Format*. When prompted to confirm that you want to format the partition, tap or 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, tap or click Drive Options (Advanced) to display the advanced options for working with drives.
2. Tap or click the partition you want to delete.
3. Tap or click Delete. When prompted to confirm that you want to delete the partition, tap or click OK. Setup then deletes the partition.

## EXTENDING DISK PARTITIONS DURING INSTALLATION

Windows Server 2012 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 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 as well, 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, tap or click Drive Options (Advanced) to display the advanced options for working with drives.
2. Tap or click the partition you want to extend.
3. Tap or click Extend. In the Size box, set the size of the partition in megabytes and then tap or click Apply to extend the selected partition.
4. When prompted to confirm that you want to extend the partition, tap or click OK. Setup then extends the partition.

## Changing the Installation Type

Unlike earlier releases of Windows Server, you can change the installation type of any server running Windows Server 2012. This is possible because a key difference between the installation types relates to whether the installation has the following User Interfaces and Infrastructure features:

- Graphical Management Tools And Infrastructure
- Desktop Experience
- Server Graphical Shell

Full-server installations have both the Graphical Management Tools And Infrastructure feature and the Server Graphical Shell feature. They also might have Desktop Experience. On the other hand, minimal-interface installations have only the Graphical Management Tools And Infrastructure feature and Server Core installations have none of these features.

Knowing that Windows also automatically installs or uninstalls dependent features, server roles, and management tools to match the installation type, you can convert from one installation type to another simply by adding or removing the appropriate User Interfaces and Infrastructure features.

## Converting Full-Server and Minimal-Interface Installations

To convert a full-server installation to a minimal-interface installation, you remove the Server Graphical Shell. Although you can use the Remove Roles And Features Wizard to do this, you also can do this at a PowerShell prompt by entering the following command:

```
uninstall-windowsfeature server-gui-shell -restart
```

This command instructs Windows Server to uninstall the Server Graphical Shell and restart the server to finalize the removal. If Desktop Experience also is installed, this feature will be removed as well.

**TIP** As a best practice before you run this or any other command that might have far-reaching effects, you should run the command with the *-Whatif* parameter. This parameter tells Windows PowerShell to confirm exactly what will happen when a command is run.

To convert a minimal-interface installation to a full-server installation, you add the Server Graphical Shell. You can use the Add Roles And Features Wizard to do this, or you can enter the following command at a PowerShell prompt:

```
install-windowsfeature server-gui-shell -restart
```

This command instructs Windows Server to install the Server Graphical Shell and restart the server to finalize the installation. If you also want to install the Desktop Experience, you can use this command instead:

```
install-windowsfeature server-gui-shell, desktop-experience -restart
```

## Converting Server Core Installations

To convert a full-server or minimal-interface installation to a Server Core installation, you remove the user interfaces for Graphical Management Tools And Infrastructure. If you remove the WoW64 Support framework, you also convert the server to a Server Core installation. Although you can use the Remove Roles And Features Wizard to remove the user interfaces, you also can do this at a PowerShell prompt by entering the following command:

```
uninstall-windowsfeature server-gui-mgmt-infra -restart
```

This command instructs Windows Server to uninstall the user interfaces for Graphical Management Tools And Infrastructure and restart the server to finalize the removal. Because many dependent roles, role services, and features might be uninstalled along with the user interfaces, run the command with the *-Whatif* parameter first to get details on what exactly will be uninstalled.

If you installed the server with the user interfaces and converted it to a Server Core installation, you can revert back to a full-server installation with the following command:

```
install-windowsfeature server-gui-mgmt-infra -restart
```

As long as the binaries for this feature and any dependent features haven't been removed, the command should succeed. If the binaries were removed, however, or Server Core was the original installation type, you need to specify a source for the required binaries.

You use the `-Source` parameter to restore required binaries from a Windows Imaging (WIM) mount point. For example, if your enterprise has a mounted Windows Image for the edition of Windows Server 2012 you are working with available at the network path `\\ImServer18\WinS12EE`, you could specify the source as follows:

```
install-windowsfeature server-gui-mgmt-infra -source \\imserver18\wins12ee
```

While many large enterprises might have standard images that can be mounted using network paths, you also can mount the Windows Server 2012 distribution media and then 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 disc 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 prompt: **`dism /get-wiminfo /wimfile:e:\sources\install.wim`**, where *e*: is the drive designator of the server's disc drive.
3. Mount the installation image by entering the following command at an elevated prompt: **`dism /mount-wim /wimfile:e:\sources\install.wim /index:2 /mountdir:c:\mountdir /readonly`**, 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. Mounting the image might take several minutes.
4. Use `Install-WindowsFeature` at a PowerShell prompt with the source specified as **`c:\mountdir\windows\winsxs`**, as shown in this example:

```
install-windowsfeature server-gui-mgmt-infra  
-source c:\mountdir\windows\winsxs
```

## Managing Roles, Role Services, and Features

---

When you want to manage server configurations, you'll primarily use Server Manager to manage roles, role services, and features. Not only can you use Server Manager to add or remove roles, role services, and features, but you can also use Server Manager to view the configuration details and status for these software components.

## Performing Initial Configuration Tasks

Server Manager is your central management console for the initial setup and configuration of roles and features. Not only can Server Manager help you quickly set up a new server, the console also can help you quickly set up your management environment.

Normally, Windows Server 2012 automatically starts Server Manager whenever you log on and you can access Server Manager on the desktop. If you don't want the console to start each time you log on, tap or click Manage and then tap or click Server Manager Properties. In the Server Manager Properties dialog box, select Do Not Start Server Manager Automatically At Logon and then tap or click OK.

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

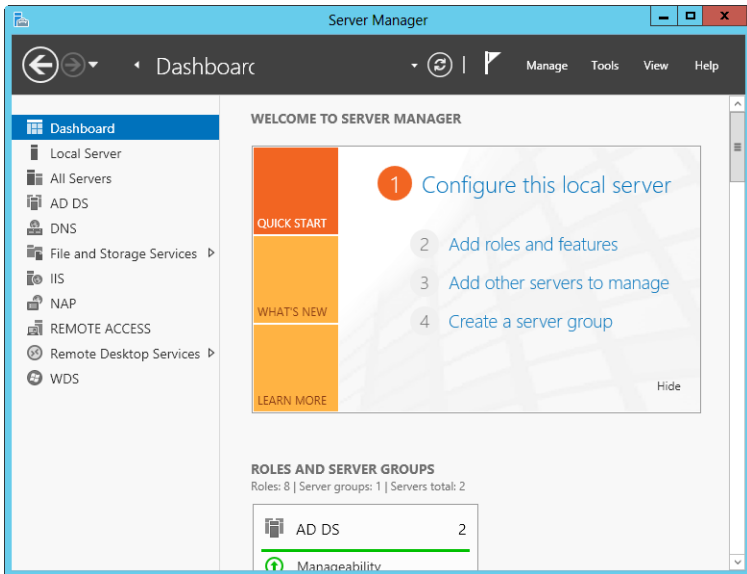
As Figure 2-1 shows, Server Manager's default view is the dashboard. The dashboard has quick links for adding roles and features to local and remote servers, adding servers to manage, and creating server groups. You'll find similar options are on the Manage menu:

- **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 Other Servers To Manage** 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. Press and hold or 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.

**TIP** When you need to connect to a server using alternate credentials, press and hold or right-click a server in the All Servers node and then select Manage As. In the Windows Security dialog box, enter your alternate credentials and then tap or click OK. 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. You need to repeat this procedure any time you change the password associated with the alternate credentials.

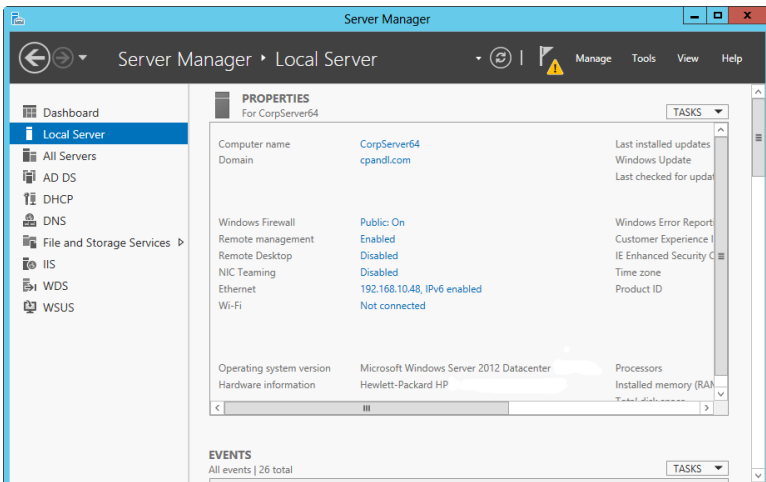
**REAL WORLD** When you are working with Server Core installations, you can use Sconfig to configure domain and workgroup membership, the computer's name, remote management, Windows Update, Remote Desktop, network settings, and the date and time. You also can use Sconfig to log off, restart, and shut down the 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.





**FIGURE 2-1** Use the dashboard for general administration.

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, as shown in Figure 2-2, you can manage the basic configuration of the server you are logged on to locally.



**FIGURE 2-2** Manage the properties of the local 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, tap or click Tasks and then tap or click Start BPA Scan.
- **Events** Provides summary information about warning and error events from the server's event logs. Tap or 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, tap or click Tasks and then tap or 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, press and hold or right-click it and then select Remove Role Or Feature.
- **Services** Lists the services running on the server by name, status and start type. Press and hold or 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. Tap or 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 tapping or clicking Change, providing the computer name and domain information, and then tapping or 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 can display the System Properties dialog box with the Computer Name tab selected by tapping or clicking System and then tapping or 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). Tap or 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, as well as 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. Tap or 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 tapping or 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 tapping or clicking Change Adapter Settings under Tasks in Network And Sharing Center.

- **IE Enhanced Security Configuration** Shows the status of Internet Explorer Enhanced Security Configuration (IE ESC). Tap or click the related link to enable or disable IE ESC. If you tap or 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 NIC teaming. Tap or click the related link to add or remove teamed interfaces and to manage related options.
- **Product ID** Shows the product identifier for Windows Server. Tap or click the related link to enter a product key and activate the operating system over the Internet.
- **Remote Desktop** Tap or 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 tapping or 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 tapping or clicking Remote Settings in the left pane.
- **Remote Management** Shows whether remote management of this server from other servers is enabled. Tap or click the related link to enable or disable remote management.

- **Time Zone** Shows the current time zone for the server. Tap or click the related link to display the Date And Time dialog box. You can then configure the server's time zone by tapping or clicking Change Time Zone, selecting the appropriate time zone, and then tapping or clicking OK twice. You can also display the Date And Time dialog box by pressing and holding or 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). Tap or 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 using Action Center. To open Action Center, tap or click the Action Center icon in the notification area of the taskbar and then select 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. Tap or 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 tapping or clicking the Windows Firewall option.
- **Windows Update** Shows the current configuration of Windows Update. Tap or 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.

## Server Manager Essentials and Binaries

The Server Manager console is designed to handle core system administration tasks. You'll spend a lot of time working with this tool, and you should get to know every detail. By default, Server Manager is started automatically. If you closed the console or disabled automatic startup, you can open the console by tapping or clicking the related option on the taskbar. Alternatively, another way to do this is by pressing the Windows key, typing **ServerManager.exe** into the Apps Search box, and then pressing Enter.

Server Manager's command-line counterpart is the ServerManager module for Windows PowerShell. When you are logged on to Windows Server 2012, this module is imported into Windows PowerShell by default. Otherwise, you need to import

the module before you can use the cmdlets it provides. You import the Server-Manager module by entering **Import-Module ServerManager** at the Windows PowerShell prompt. Once 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 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 typing **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. By using `Install-Windows-Feature` 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.

Binaries needed to install roles and features are referred to as *payloads*. With Windows Server 2012, 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 using the `-Remove` parameter of the `Uninstall-WindowsFeature` cmdlet. Subcomponents of the role or feature are removed as well. To also remove management tools, add the **-includeallmanagementtools** parameter.

When you want to install a role or feature, you can install the related components and restore any removed payloads for these components 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 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) mount point. For example, if your enterprise has a mounted Windows Image for the edition of Windows Server 2012 you are working with available at the network path `\\ImServer18\WinS12EE`, you could specify the source as follows:

```
install-windowsfeature -name ad-domain-services -includeallsubfeature  
-source \\imserver18\wins12ee
```

Keep in mind that the path you specify is only used if required binaries are not found in the Windows Side-By-Side folder on the destination server. While many large enterprises might have standard images that can be mounted using network paths, you also can mount the Windows Server 2012 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 disc drive, and then create a folder to mount the Installation image by entering the following command: **mkdir c:\mountdir**.

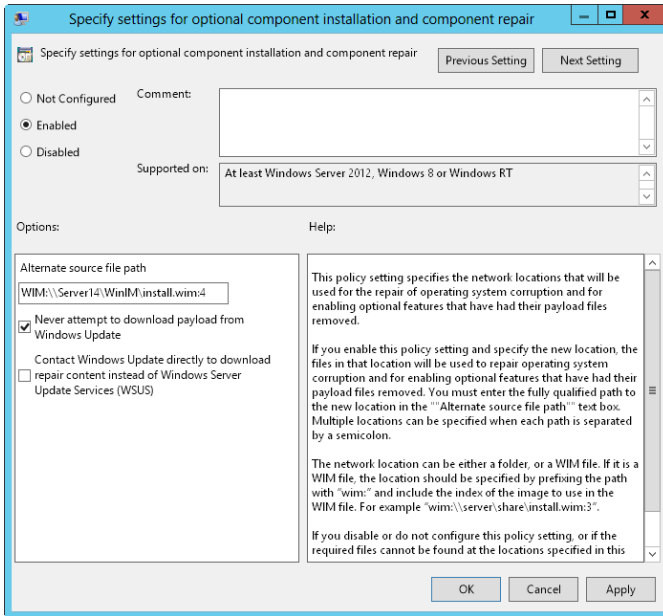
2. Locate the index number of the image you want to use by entering the following command at an elevated prompt: **dism /get-wiminfo /wimfile:e:\sources\install.wim**, where *e*: is the drive designator of the server's disc drive.
3. Mount the installation image by entering the following command at an elevated prompt: **dism /mount-wim /wimfile:e:\sources\install.wim /index:2 /mountdir:c:\mountdir /readonly**, 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. Mounting the image might take several minutes.
4. Use Install-WindowsFeature at a PowerShell 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 Computer Configuration\Administrative Templates\System. This policy also is used for obtaining payloads needed to repair components.

If you enable this policy (as shown in Figure 2-3), 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\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`.
- 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.



**FIGURE 2-3** Control component installation through Group Policy.

## Managing Your Servers Remotely

You can use Server Manager and other Microsoft Management Consoles (MMCs) to perform some management tasks on remote computers, as long as the computers are in the same domain or you are working in a workgroup and have added the remote computers in a domain as trusted hosts. You can connect to servers running full-server, minimal-interface, and Server Core installations. On the computer you want to use for managing remote computers, you should be running either Windows Server 2012 or Windows 8 and you need to install the Remote Server Administration Tools.

With Windows Server 2012, the Remote Server Administration Tools are installed as a feature 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 “Server Manager Essentials and Binaries” earlier in the chapter.

You can get the Remote Server Administration Tools for Windows 8 as a download from the Microsoft Download Center (<http://download.microsoft.com>). Different versions are available for x64 and x86 systems.

By default, remote management is enabled for servers running Windows Server 2012 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 is enabled by default for Windows Server 2012. 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 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 tap or click OK.

When you clear Enable Remote Management Of This Server From Other Computers and then tap or click OK, Server Manager performs several background tasks that disable Windows Remote Management (WinRM) 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 using Windows Remote Management. The exceptions for Windows Management Instrumentation and Windows Remote Management (Compatibility) aren't affected.

You must be a member of the Administrators group on computers you want to manage by using Server Manager. For remote connections in a workgroup-to-workgroup or workgroup-to-domain configuration, you should be logged on using the built-in Administrator account or configure the *LocalAccountTokenFilterPolicy* registry key to allow remote access from your computer. To set this key, enter the following command at an elevated, administrator command prompt:

```
reg add HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System /v LocalAccountTokenFilterPolicy /t REG_DWORD /d 1 /f
```



**NOTE** You also can enable remote management by entering **configure-SMRemoteing.exe –enable** at an elevated, administrator prompt.

**NOTE** If you want to make it possible to remotely manage a computer running Windows 8 using the WS-Management protocol, enter **winrm quickconfig** at an elevated prompt. Then each time you are prompted to make configuration changes, enter **Y**. This will start the Windows Remote Management (WinRM) 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 using Remote Desktop, you must allow related connections to the computer and configure access as discussed in Chapter 4, “Automating Administrative Tasks, Policies, and Procedures.”
- Remote Service Management must be configured as an allowed app in Windows Firewall to remotely manage a computer’s services. In the advanced firewall, there are several related rules that 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.

You can configure remote management on a Server Core installation of Windows Server 2012 using Sconfig. Start the Server Configuration utility by entering **sconfig**.

## Connecting to and Working with Remote Servers

Using Server Manager, you can connect to and manage remote servers, provided that you’ve added the server for management. To add servers one at a time to Server Manager, complete 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. The Add Servers dialog box has several panels for adding servers:
  - The Active Directory panel, selected by default, 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, tap or click Find Now.
  - The DNS panel allows you to add servers by computer name or IP address. After you enter the name or IP address, tap or click the Search button.
3. In the Name list, double-tap or double-click the server to add it to the Selected list.
4. Repeat steps 2 and 3 to add others servers. Tap or click OK.

To add many servers to Server Manager, you can use the Import process and these steps:

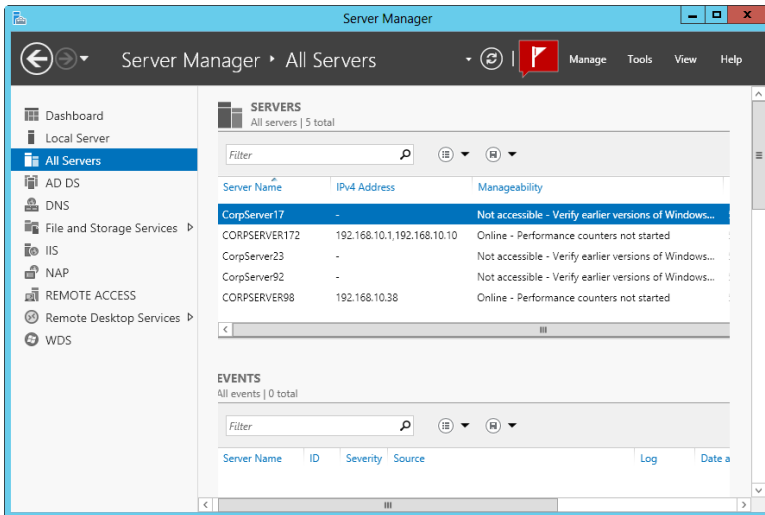
1. Create a text file that has one host name, fully qualified domain name, or IP address per line.
2. In Server Manager, select Add Servers on the Manage menu. In the Add Servers dialog box, select the Import panel.
3. Tap or click the options button to the right of the File box, and then use the Open dialog box to locate and open the server list.
4. In the Computer list, double-tap or double-click each server you want to add to the Selected list. Tap or click OK.

After you add a remote computer, the Server Manager console shows the name of the remote computer in the All Servers view. Server Manager always resolves IP addresses to host names. As shown in Figure 2-4, the All Servers view also lists the Manageability status of the server as well. If a server is listed as "Not accessible," you typically need to log on locally to resolve the problem.

In the All Servers view, the servers you add are listed in the Servers pane so that you can manage them each time you work with Server Manager. Server Manager tracks the services, events, and more for each added server, and each server is added to the appropriate server groups automatically 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 see a list of the domain controllers you added for management as well as 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 DNS, or by importing a list of host names, fully qualified domain names, or IP addresses. Any server you add to a custom group is added automatically for management as well.



**FIGURE 2-4** Note the Manageability status of each server, and take corrective actions as necessary.

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 pane, 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, tap or click Find Now. In the Name list, double-tap or 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, tap or click the Search button. In the Name list, double-tap or double-click a server to add it to the Selected list.
  - The Import panel allows you to import a list of servers. Tap or 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-tap or double-click a server to add it to the Selected list.
3. Tap or click OK to create the server group.

When you press and hold or right-click a server name in the Servers pane of a server group or in the All Servers view, you display an extended list of management options. These options perform the corresponding task or open the corresponding

management tool with the selected server in focus. For example, if you were to right-click CorpServer172 and then select Computer Management, Computer Management connects to CorpServer172 and then opens.

You can work with a remote computer using an interactive remote Windows PowerShell session. To do this, open an elevated, administrator Windows PowerShell prompt. Type **enter-*pssession* 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. When prompted to enter the authorized user’s password, type the password and then press Enter. You can now enter commands in the session as you would if you were using Windows PowerShell locally. To exit the session, enter **exit-*pssession***.

The following example enters an interactive remote session with Server85 using the credentials of Williams:

```
enter-pssession server85 –credential williams
```

## Adding and Removing Roles, Role Services, 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.

When you select a role-based group in the left pane, the Servers panel 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 tap or click an event, you can get more information about the event.
- Summary information about the status of related system services. You can press and hold or right-click a service to manage its run status.

**TIP** By default, Server Manager refreshes details every 10 minutes. You can refresh the details manually by tapping or clicking the Refresh Servers button on the toolbar. If you want to set a different default refresh interval, tap or click Manage and then tap or click Server Manager Properties. Next, set the new refresh interval in minutes and then tap or click OK.

You can manage a service by pressing and holding or right-clicking the service and then tapping or clicking Stop Service, Start Service, Pause Service, Resume Service, or Restart Service 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. See Chapter 3, “Monitoring Processes, Services, and Events,” for detailed information about working with events and system services.

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 2012, you can install roles and features on running servers (whether physical machines or virtual) as well as virtual hard disks. Servers 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, select 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 tap or click Next. You can avoid seeing the Before You Begin page the next time you start this wizard by selecting the Skip This Page By Default check box before tapping or clicking Next.
2. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Tap or 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, tap or click Browse and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, tap or click Next.

**NOTE** Only servers running Windows Server 2012 and that have been added for management in Server Manager are listed.

4. On the Server Roles page, select the role or roles to install. If additional features are required to install a role, you'll see an additional dialog box. Tap or click Add Features to close the dialog box and add the required features to the server installation. Tap or click Next to continue.

**NOTE** Some roles cannot be added at the same time as other roles. You have to install each role separately. Other roles cannot be combined with existing roles, and you'll see warning prompts about this. A server running a Server Core installation can act as a domain controller and can also hold any of the flexible single-master operations (FSMO) roles for Active Directory.

5. On the Features page, select the feature or features to install. If additional features are required to install a feature you selected, you'll see an additional dialog box. Tap or click Add Features to close the dialog box and add the required features to the server installation. When you are ready to continue, tap or click Next.
6. With some roles, you'll see an extra wizard page, which provides additional information about using and configuring the role. You may also have the opportunity to install additional role services as part of a role. For example, with Print And Document Services, Web Server Role (IIS), and WSUS, you'll see an

additional information page and a page for selecting role services to install along with the role.

7. On the Confirmation page, tap or click the Export Configuration Settings link to generate an installation report that can be displayed in Internet Explorer.
8. 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, type that alternate path in the box provided, and then tap or click OK. For example, if you mounted a Windows image and made it available on the local server as discussed in "Server Manager Essentials and Binaries" earlier, you could enter the alternate path as **c:\mountdir\windows\winsxs**. 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**.
9. After you review the installation options and save them as necessary, tap or click Install to begin the installation process. The Installation Progress page tracks the progress of the installation. If you close the wizard, tap or click the Notifications icon in Server Manager and then tap or click the link provided to reopen the wizard.
10. 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.

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.

You can remove a server role or feature by following these steps:

1. In Server Manager, select 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 tap or click Next. You can avoid seeing the Before You Begin page the next time you start this wizard by selecting the Skip This Page By Default check box before tapping or clicking Next.
2. On the Server Selection page, you can choose to remove roles and features from 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 removing roles and features from a VHD, tap or click Browse and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, tap or click Next.
3. On the 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

remove the other role as well. If you tap or click the Remove Features button, the wizard removes the dependent roles and features as well. Note that if you want to keep related management tools, you should clear the Remove Management Tools check box prior to tapping or clicking the Remove Features button and then click Continue. Tap or click Next.

4. On the 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, you'll see a warning prompt stating that you cannot remove the feature unless you also remove the other feature or role. If you tap or click the Remove Features button, the wizard removes the dependent roles and features as well. 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 tapping or clicking the Remove Features button. Tap or click Next.
5. On the Confirmation page, review the related components that the wizard will remove based on your previous selections and then tap or click Remove. The Removal Progress page tracks the progress of the removal. If you close the wizard, tap or click the Notifications icon in Server Manager and then tap or click the link provided to reopen the wizard.
6. When the wizard finishes modifying the server configuration, you'll see the Removal Progress page. 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.

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.

## Managing System Properties

---

You use the System console to view system information and perform basic configuration tasks. To access the System console, double-tap or double-click System in Control Panel. As Figure 2-5 shows, 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. The type of operating system is listed as 32-bit or 64-bit.
- **Computer Name, Domain, And Workgroup Settings** Provides the computer name, description, domain, and workgroup details. If you want to change any of this information, tap or click Change Settings and then tap or 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 2012 isn't activated yet, tap or click the link provided to start the activation process and then follow the prompts.

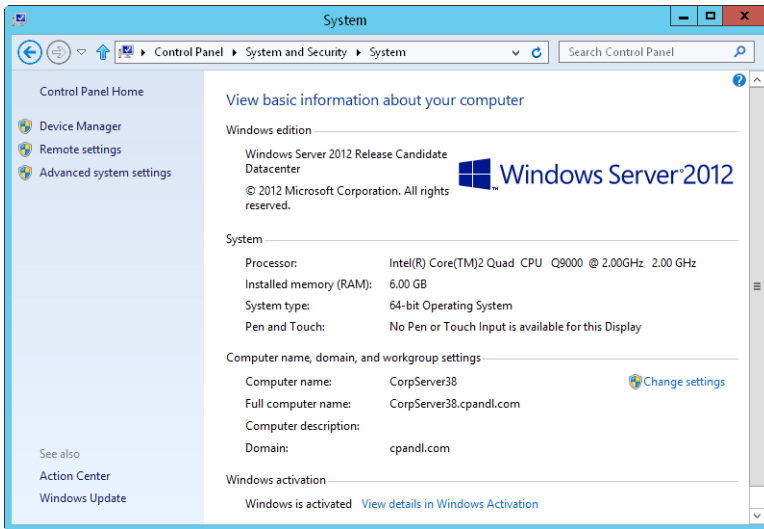


FIGURE 2-5 Use the System console to view and manage system properties.

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

Although volume-licensed versions of Windows Server 2012 might not require activation or product keys, retail versions of Windows Server 2012 require both activation and product keys. If Windows Server 2012 has not been activated, you can activate the operating system by selecting **Activate Windows Now** under **Windows Activation**. You can also activate Windows by entering **slmgr -ato** at a command prompt.

You can change the product key provided during installation of Windows Server 2012 to stay in compliance with your licensing plan. 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 using a confirmation identifier. To view this and other options, enter **slmgr** at a command prompt.

Within the System console, you can access the System Properties dialog box and use this dialog box to manage system properties. Tap or click **Change Settings** under **Computer Name, Domain, And Workgroup Settings**. The following sections examine



key areas of the operating system you can configure using the System Properties dialog box.

## The Computer Name Tab

You can display and modify the computer's network identification on the Computer Name tab of the System Properties dialog box. The Computer Name tab displays the full computer name of the system and the domain membership. The full computer name is essentially the Domain Name System (DNS) name of the computer, which also identifies the computer's place within the Active Directory hierarchy. If a computer is a domain controller or a certificate authority, you can change the computer name only after removing the related role from the computer.

You can join a computer to a domain or workgroup by following these steps:

1. On the Computer Name tab of the System Properties dialog box, tap or click Change. This displays the Computer Name/Domain Changes dialog box.
2. To put the computer in a workgroup, select the Workgroup option and then type the name of the workgroup to join.
3. To join the computer to a domain, select the Domain option, type the name of the domain to join, and then tap or click OK.
4. If you changed the computer's domain membership, you'll see a Windows Security prompt. Enter the name and password of an account with permission to add the computer to the specified domain or to remove the computer from a previously specified domain, and then tap or click OK.
5. When prompted that your computer has joined the workgroup or domain you specified, tap or click OK.
6. You'll see a prompt stating that you need to restart the computer. Tap or click OK.
7. Tap or click Close, and then tap or click Restart Now to restart the computer.

To change the name of a computer, follow these steps:

1. On the Computer Name tab of the System Properties dialog box, tap or click Change. This displays the Computer Name/Domain Changes dialog box.
2. In the Computer Name text box, type the new name for the computer.
3. You'll see a prompt stating that you need to restart the computer. Tap or click OK.
4. Tap or click Close, and then tap or click Restart Now to restart the computer.

## The Hardware Tab

The System Properties dialog box's Hardware tab provides access to Device Manager and Driver Installation Settings. To access the Hardware tab, open the System Properties dialog box and then tap or click the Hardware tab.

For installed devices, you can configure Windows Server to download driver software and realistic icons for devices. By default, Windows Server does not do this.

If you want a computer to check for drivers automatically, tap or click the Device Installation Settings button and then select either Yes, Do This Automatically or No, Let Me Choose What To Do. If you want to choose what to do, you can specify the following:

- Always install the best driver software from Windows Update
- Never install driver software from Windows Update
- Automatically get the device apps and info provided by your device manufacturer

The first two options do exactly what they say. The final option tells Windows Update that you want to get metadata and companion applications for devices. Tap or click Save Changes, and then tap or click OK to apply your changes.

## The Advanced Tab

The System utility's Advanced tab controls many of the key features of the Windows operating system, including application performance, virtual memory usage, the user profile, environment variables, and startup and recovery. To access the Advanced tab, open the System Properties dialog box and then tap or click the Advanced tab.

### Setting Windows Performance

Many graphics enhancements were added to the Windows Server 2008 interface, and these enhancements are available in later releases as well. These enhancements include many visual effects for menus, toolbars, windows, and the taskbar. You can configure Windows performance by following these steps:

1. Tap or click the Advanced tab in the System Properties dialog box, and then tap or click Settings in the Performance panel to display the Performance Options dialog box.
2. The Visual Effects tab is selected by default. You have the following options for controlling visual effects:
  - **Let Windows Choose What's Best For My Computer** Enables the operating system to choose the performance options based on the hardware configuration. For a newer computer, this option will probably have the same effect as choosing the Adjust For Best Appearance option. The key distinction, however, is that this option is chosen by Windows based on the available hardware and its performance capabilities.
  - **Adjust For Best Appearance** When you optimize Windows for best appearance, you enable all visual effects for all graphical interfaces. Menus and the taskbar use transitions and shadows. Screen fonts have smooth edges. List boxes have smooth scrolling. Folders use web views and more.
  - **Adjust For Best Performance** When you optimize Windows for best performance, you turn off the resource-intensive visual effects, such as slide transitions and smooth edges for fonts, while maintaining a basic set of visual effects.

- **Custom** You can customize the visual effects by selecting or clearing the visual effects options in the Performance Options dialog box. If you clear all options, Windows does not use visual effects.
3. Tap or click Apply when you have finished changing visual effects. Tap or click OK twice to close the open dialog boxes.

## Setting Application Performance

Application performance is related to processor-scheduling caching options you set for the Windows Server 2012 system. 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). You control application performance by following these steps:

1. Access the Advanced tab in the System Properties dialog box, and then display the Performance Options dialog box by tapping or clicking Settings in the Performance panel.
2. In the Performance Options dialog box, tap or 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 2012 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. Tap or click OK.

## Configuring Virtual Memory

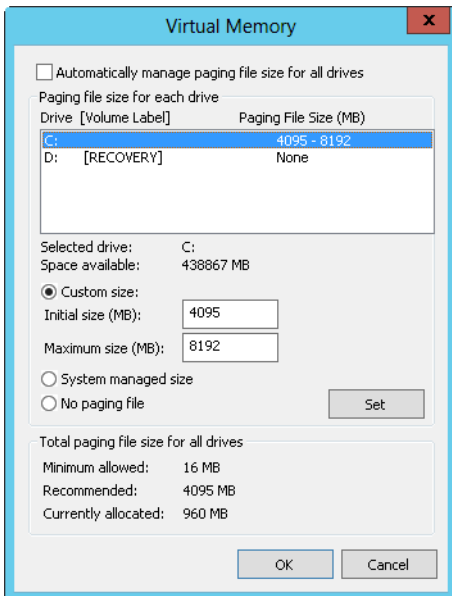
With virtual memory, you can use disk space to extend the amount of memory available on a system by using part of the hard disk as part of system memory. This feature writes RAM to disks by using a process called *paging*. With paging, a set amount of RAM, such as 8192 megabytes (MB), is written to the 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.

**REAL WORLD** Current releases of Windows Server automatically manage virtual memory much better than their predecessors. Typically, Windows Server allocates virtual memory in an amount at least as large as the total physical memory installed on the computer. This helps to ensure that paging files don't become fragmented, which can result in poor system performance. If you want to manage virtual memory manually, you can use a fixed virtual memory size in most cases. To do this, set the initial size and the maximum size 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). In most cases, for computers with 8 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.

You can configure virtual memory by following these steps:

1. Access the Advanced tab in the System Properties dialog box, and then display the Performance Options dialog box by tapping or clicking Settings in the Performance panel.
2. In the Performance Options dialog box, tap or click the Advanced tab and then tap or click Change to display the Virtual Memory dialog box, shown in Figure 2-6.



**FIGURE 2-6** Virtual memory extends the amount of RAM on a system.

The following information is provided:

- **Paging File Size For Each Drive** Provides information on 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).
3. 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.
  4. In the Drive list, select the volume you want to work with.
  5. Select Custom Size, and then enter values in the Initial Size and Maximum Size boxes.
  6. Tap or click Set to save the changes.
  7. Repeat steps 4–6 for each volume you want to configure.

**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 4 GB of RAM would need a paging file of 4 GB on the system drive.

8. Tap or click OK. If prompted to overwrite an existing Pagefile.sys file, tap or click Yes.
9. If you updated the settings for a paging file that is currently in use, you'll see a prompt indicating that you need to restart the system for the changes to take effect. Tap or click OK.
10. Tap or click OK twice to close the open dialog boxes. When you close the System utility, you'll see a prompt asking if you want to restart the system. Tap or click Restart.

You can have Windows Server 2012 automatically manage virtual memory by following these steps:

1. Access the Advanced tab in the System Properties dialog box, and then display the Performance Options dialog box by tapping or clicking Settings in the Performance panel.
2. Tap or click the Advanced tab, and then tap or click Change to display the Virtual Memory dialog box.
3. Select the Automatically Manage Paging File Size For All Drives check box.

4. Tap or click OK three times to close the open dialog boxes.

**NOTE** If you updated the settings for the paging file currently in use, you'll see a prompt indicating that you need to restart the server for the changes to take effect. Tap or click OK. When you close the System Properties dialog box, you'll see a prompt telling you that you need to restart the system for the changes to take effect. On a production server, you should schedule this reboot outside normal business hours.

## Configuring Data Execution Prevention

Data Execution Prevention (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 it from executing. This 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.

**NOTE** The 32-bit versions of Windows support DEP as implemented by Advanced Micro Devices (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. The 64-bit versions of Windows also support the NX processor feature.

### 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. Access the Advanced tab in the System Properties dialog box, and then display the Performance Options dialog box by tapping or clicking Settings in the Performance panel.
2. In the Performance Options dialog box, tap or click the Data Execution Prevention tab. The text at the bottom of this tab indicates whether the computer supports execution protection.
3. 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 tap or 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.
4. Tap or click OK.

If you turned on DEP and allowed exceptions, you can add or remove a program as an exception by following these steps:

1. Access the Advanced tab in the System Properties dialog box, and then display the Performance Options dialog box by tapping or clicking Settings in the Performance panel.
2. In the Performance Options dialog box, tap or click the Data Execution Prevention tab.
3. To add a program as an exception, tap or click Add. Use the Open dialog box to find the executable file for the program you are configuring as an exception, and then tap or click Open.
4. To temporarily disable a program as an exception (this might be necessary for troubleshooting), clear the check box next to the program name.
5. To remove a program as an exception, tap or click the program name and then tap or click Remove.
6. Tap or click OK to save your settings.

### Understanding DEP Compatibility

To be compatible with DEP, applications must be able to mark memory explicitly with 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 user-mode 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 System and User 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, shown in Figure 2-7. To access this dialog box, open the System Properties dialog box, tap or click the Advanced tab, and then tap or click Environment Variables.

### CREATING AN ENVIRONMENT VARIABLE


You can create an environment variable by following these steps:

1. Tap or 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.
2. In the Variable Name text box, type the variable name. In the Variable Value text box, type the variable value.
3. Tap or click OK.

### EDITING AN ENVIRONMENT VARIABLE

You can edit an environment variable by following these steps:

1. Select the variable in the User Variables or System Variables list.
2. Tap or 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.
3. Type a new value in the Variable Value text box, and then tap or click OK.



**FIGURE 2-7** Configure system and user environment variables in the Environment Variables dialog box.



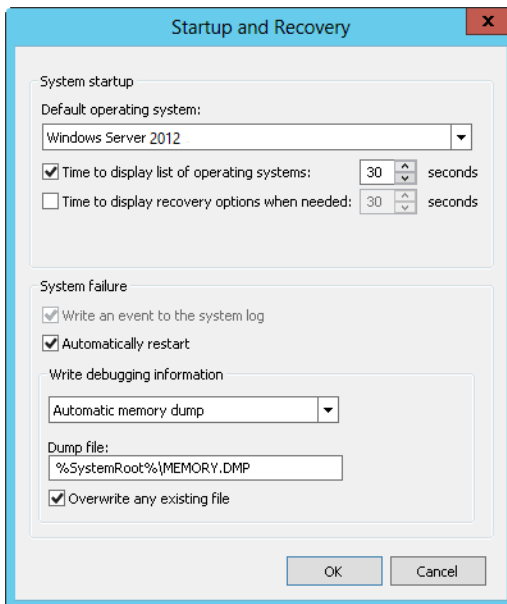
## DELETING AN ENVIRONMENT VARIABLE

To delete an environment variable, select it and tap or click Delete.

**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.

## Configuring System Startup and Recovery

You configure system startup and recovery properties in the Startup And Recovery dialog box, shown in Figure 2-8. To access this dialog box, open the System Properties dialog box, tap or click the Advanced tab, and then tap or click Settings in the Startup And Recovery panel.



**FIGURE 2-8** Configure system startup and recovery properties in the Startup And Recovery dialog box.

## 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 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 fatal 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.
- **Small Memory Dump** Use this option to dump the physical memory segment in which the error occurred. This dump is 256 KB in size.
- **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.

- **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.
- **Automatic Memory Dump** Use this option to let Windows determine which type of memory dump is best and create the dump file accordingly.

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 usually want to select Overwrite Any Existing File as well. 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.)

## The Remote Tab

The Remote tab of the System Properties dialog box controls Remote Assistance invitations and Remote Desktop connections. These options are discussed in Chapter 4.



# Monitoring Processes, Services, and Events

- Managing Applications, Processes, and Performance **87**
- Managing System Services **100**
- Event Logging and Viewing **108**
- Monitoring Server Performance and Activity **118**
- Tuning System Performance **130**

As an administrator, you need to keep an eye on network systems. The status and usage of system resources can change dramatically over time. Services might stop running. File systems might run out of space. Applications might throw exceptions that, in turn, can cause system problems. Unauthorized users might try to break into the system. The techniques discussed in this chapter can help you identify and resolve these and other system problems.

## Managing Applications, Processes, and Performance

---

Any time you start an application or type a command at the command line, Microsoft Windows Server 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, only 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.

## Task Manager

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 tap or click Task Manager.
- Press the Windows key, type **taskmgr**, and then press Enter.
- Press and hold or right-click the taskbar, and then tap or click Task Manager on the shortcut menu.

**NOTE** When you press the Windows key and type **taskmgr**, you'll see two matches. One match is the full name, Task Manager. The other match is the command name, **taskmgr**.

The following sections cover techniques you use to work with Task Manager.

## Viewing and Working with Processes

Task Manager has two general views:

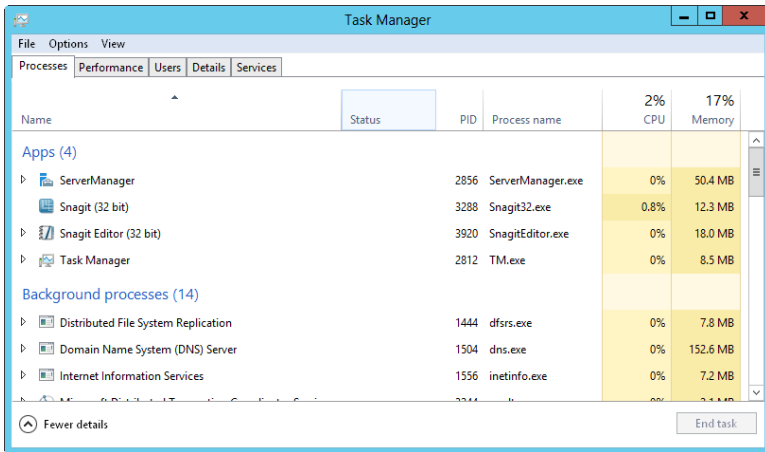
- **Summary** Shows only applications running in the foreground, which lets you quickly select and work with foreground applications
- **Expanded** Expands the view, providing additional tabs that you can use to get information about all running processes, system performance, connected users, and configured services

If you are in summary view, you can switch to expanded view by tapping or clicking More Details. If you are in the expanded view, you can switch to summary view by tapping or 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. As shown in Figure 3-1, the expanded view has multiple tabs you can select to work with running processes, system performance, connected users, and configured services. The Processes tab, also shown in Figure 3-1, 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

**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 tapping or 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.

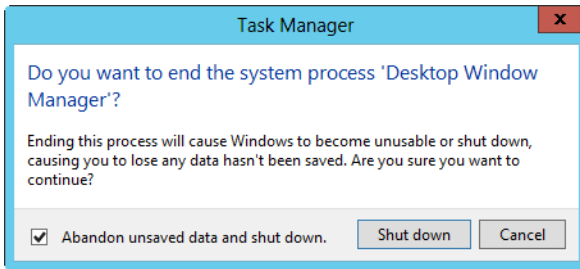


**FIGURE 3-1** View the status of processes currently running on the server.

**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 frozen before you end its related task.

You can stop a process by selecting the process and then tapping or clicking **End Task**. You shouldn't try to stop Windows processes using this technique. If you try to stop a Windows process or a group of Windows processes, Task Manager displays a warning prompt similar to the one shown in Figure 3-2. This warning states that ending this process will cause Windows to become unusable or to shut down. To proceed, you must select **Abandon Unsaved Data And Shut Down** and then tap or click **Shut Down**. Windows then displays a blue screen with an error code. After collecting error information, Windows will restart.



**FIGURE 3-2** Stopping processes for essential Windows services causes Windows to become unusable or to shut down.

Other columns on the Processes tab provide a lot of information about running processes. You can use this information to determine which processes are over consuming system resources such as CPU time and memory. Although only CPU and Memory columns are displayed by default, others columns can be added by pressing and holding or right-clicking any column header and then selecting options 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, as well as 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** Lists the publisher of the process, such as Microsoft Corporation.
- **Type** Displays the general process type as app, background process, or Windows process. This information is useful if you clear the Group By Type option on the View menu.

Pressing and holding or right-clicking an application's listing in Task Manager displays a shortcut menu 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



**NOTE** 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.

## Administering Processes

Task Manager's Details tab is shown in Figure 3-3. This 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

Name	User name	CPU	Memory (p...)	Status	Description
csrss.exe	SYSTEM	00	936 K	Running	Client Server Runtime Process
csrss.exe	SYSTEM	00	1,288 K	Running	Client Server Runtime Process
dfsr.exe	SYSTEM	00	8,108 K	Running	Distributed File System Replication
dfssvc.exe	SYSTEM	00	1,336 K	Running	Windows NT Distributed File System Service
dns.exe	SYSTEM	00	156,276 K	Running	Domain Name System (DNS) Server
dwm.exe	DWM-1	00	21,828 K	Running	Desktop Window Manager
explorer.exe	WilliamS	00	12,144 K	Running	Windows Explorer
inetinfo.exe	SYSTEM	00	7,416 K	Running	Internet Information Services
ismsserv.exe	SYSTEM	00	1,036 K	Running	Windows NT Intersite Messaging Service
lsass.exe	SYSTEM	00	13,648 K	Running	Local Security Authority Process
Microsoft.ActiveDirectory...	SYSTEM	00	9,972 K	Running	Microsoft.ActiveDirectory.WebServices
msdtc.exe	NETWORK...	00	2,152 K	Running	Microsoft Distributed Transaction Coordinator Service
ServerManager.exe	WilliamS	00	49,952 K	Running	ServerManager
services.exe	SYSTEM	00	3,164 K	Running	Services and Controller app
smss.exe	SYSTEM	00	280 K	Running	Windows Session Manager
Snagit32.exe	WilliamS	00	12,556 K	Running	Snagit
SnagitEditor.exe	WilliamS	00	20,260 K	Running	Snagit Editor
SnagPriv.exe	WilliamS	00	800 K	Running	Snagit RPC Helper
splwow64.exe	WilliamS	00	996 K	Running	Print driver host for applications
spoolsv.exe	SYSTEM	00	2,480 K	Running	Spooler SubSystem App
svchost.exe	SYSTEM	00	1,984 K	Running	Host Process for Windows Services
svchost.exe	NETWORK...	00	2,508 K	Running	Host Process for Windows Services
svchost.exe	LOCAL SE...	00	15,556 K	Running	Host Process for Windows Services

**FIGURE 3-3** The Details tab provides detailed information about running processes.

Other columns can be added by pressing and holding or right-clicking any column header and then tapping or clicking Select Columns. When you're trying to troubleshoot system problems using process information, you might want to add the following columns to the view:

- **Base Priority** Priority determines how much of the system's resources are allocated to a process. To set the priority for a process, press and hold or 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** The total amount of CPU cycle time used by a process since it was started. To quickly see the processes that are using the most CPU time, display this column and then tap or 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** 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** 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** 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** *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** The highest amount of memory used by the process. The change, or delta, between current memory usage and peak memory usage is important to note as well. 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 on startup 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 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 32-bit/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.

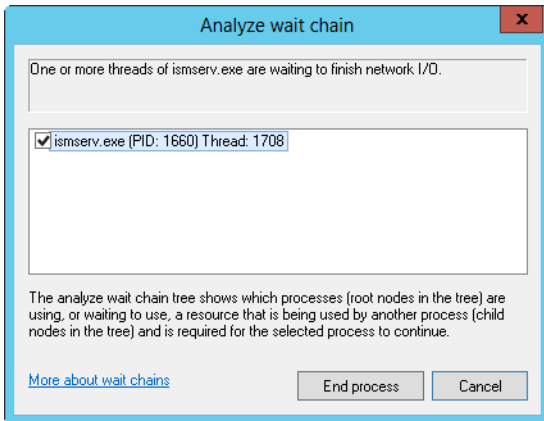
- **Process ID (PID)** The numeric identifier for the process.
- **Session ID** The identifier for the session under which the process is running.
- **Threads** The current number of threads the process is using. Most server applications are multithreaded. Multithreading allows concurrent execution of 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 Windows XP, Windows Server 2003, and earlier 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 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 amount 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 pressing and holding or right-clicking the process and then tapping or clicking Analyze Wait Chain. If the process is waiting for a resource to be released, you then see the wait chain for that process (as shown in Figure 3-4). 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 tap or 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 they are being used and free them when done. A problem occurs when a process fails to release a resource, as can happen with a poorly architected program.



**FIGURE 3-4** Analyzing wait chains.

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 pressing and holding or 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:

- Press and hold or right-click the application on the Processes tab, and then tap or click End Task.
- Press and hold or right-click the main application process on the Details tab, and then tap or click End Task.
- Press and hold or right-click the main or a dependent process on the Processes tab, and then tap or click End Process Tree.

## Viewing System Services

Task Manager's Services tab provides an overview of system services. This tab displays services by name, process ID, description, status, and group. As shown in Figure 3-5, multiple services typically run under the same process ID. You can quickly sort services by their process ID by tapping or clicking the related column heading. You can tap or click the Status column heading to sort services according to their status, Running or Stopped.

The Group column provides additional options about related identities or service host contexts under which a service runs:

- Services running under an identity with a restriction have the restriction listed in the Group column. For example, a service running under the Local

Service identity might be listed as LocalServiceNoNetwork to indicate that the service has no network access, or a service might be listed as LocalSystemNetworkRestricted to indicate that the service has restricted access to the network.

- Services that have Svchost.exe list their associated context for the *-k* parameter. For example, the RemoteRegistry service runs with the command line `svchost.exe -k regsvc`. You'll see an entry of regsvc in the Group column for this service.

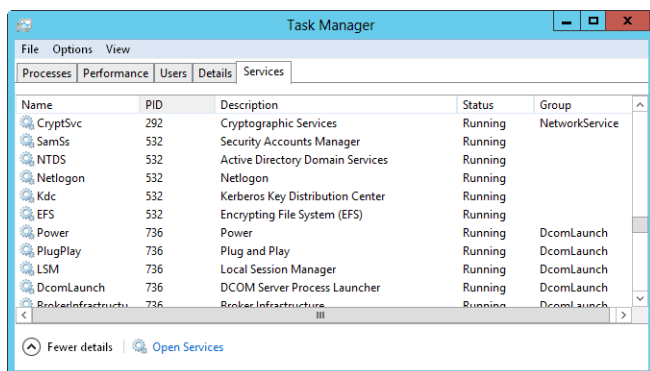


FIGURE 3-5 The Services tab provides a quick overview of the status of system services.

Pressing and holding or right-clicking a service's listing in Task Manager displays a shortcut menu that allows you to do the following:

- Start a stopped service
- Stop a started service
- Go to the related process on the Details tab

## Viewing and Managing System Performance

The Performance tab in Task Manager provides an overview of CPU and memory usage. As shown in Figure 3-6, the tab displays graphs and statistics. This information gives you a quick check of system resource usage. For more detailed information, use Performance Monitor, as explained later in this chapter.

The graphs on the Performance tab provide 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

Tap or click a summary graph in the left pane to view detailed information for that graph in the right pane. To view a close-up of any graph, double-tap or double-click the graph. Double-tapping or 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 as well as to pause the graph. Updates occur once every 4 seconds for Low, once every 2 seconds for Normal, and twice per second for High.

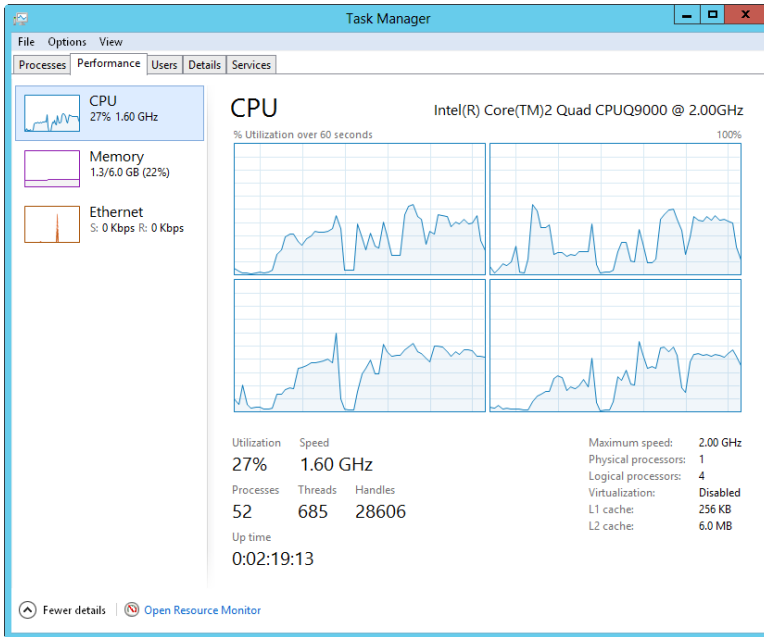


FIGURE 3-6 The Performance tab provides a quick check of system resource usage.

### CPU Usage: The Basics

When you select CPU, the % Utilization graph shows overall processor utilization for the last 60 seconds. If a system has multiple CPUs, you'll see a graph for each CPU by default. You also can view logical processors or NUMA nodes by pressing and holding or right-clicking a CPU graph, selecting Change Graph To, and then selecting Logical Processors or NUMA Nodes as appropriate.

To view kernel times, press and hold or right-click a CPU graph and then select 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 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.

You can use the CPU information provided to 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 “Tuning System Performance” later in this chapter.

## Memory Usage: The Basics

When you select Memory, 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 it 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

**NOTE** 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 take both steps.

- **Cached** Shows the amount of memory used for system caching.
- **Paged Pool** Provides information on noncritical kernel memory used by the operating system kernel.
- **Nonpaged 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.

## Network Usage: The Basics

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 usage 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. 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.

You can view detailed information on link speed, link state, bytes sent, bytes received, and more by pressing and holding or right-clicking the Network Throughput graph and selecting 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.



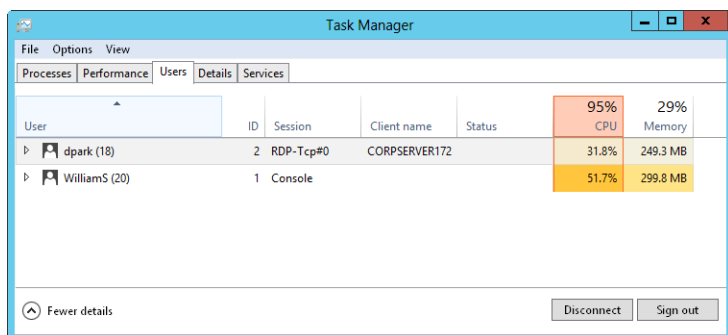
- **Bytes** Cumulative total bytes on the connection to date.

**REAL WORLD** Any time you see usage consistently approaching or exceeding 50 percent of total capacity, you should start monitoring the server more closely, and you might want to consider adding network adapters. Plan any upgrade carefully; a lot more planning is required than you might think. Consider the implications not only for that server but also for the network as a whole. You might also have connectivity problems if you exceed the allotted bandwidth of your service provider—it can often take months to obtain additional bandwidth for external connections.

## Viewing and Managing Remote User Sessions

Remote users can use Remote Desktop to connect to remote systems. Remote Desktop allows you to administer systems remotely, as if you were sitting at the console. Windows Server 2012 allows up to two active console sessions at a time.

One way to view and manage remote desktop connections is to use Task Manager. To do this, start Task Manager and then tap or click the Users tab, shown in Figure 3-7. The Users tab shows interactive user sessions for both local and remote users.



**FIGURE 3-7** The Users tab allows you to view and manage user sessions.

Each user connection is listed with user name, status, CPU utilization, and memory usage by default. Other columns can be added by pressing and holding or right-clicking any column header and then tapping or 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 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 new for Windows Server 2012 and really come in handy 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 shown in Figure 3-7, 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 press and hold or right-click a user session, you have the following options:

- **Connect** Allows you to connect a remote user session if it's inactive.
- **Disconnect** Allows you to disconnect a local or remote user session, halting all user-started applications without saving application data.
- **Sign Off** Allows you to log off a user using the normal logoff process. Application data and system state information are saved just as they are during a normal logoff.
- **Send Message** Allows you to send a console message to a logged-on user.

Also new for Windows Server 2012, the user's name is followed by the number of processes she is running. If you double-tap or double-click the user's name, you see an entry for each running process. Processes are listed by name, CPU usage, and memory usage.

## Managing System Services

---

Services provide key functions to workstations and servers. To manage system services on the local server or a remote server, you use the Services panel in Server Manager or the Services node in Computer Management. To work with services on remote servers, remote management and inbound exceptions for Remote Service Management must be enabled. For more information, see "Managing Your Servers Remotely" in Chapter 2, "Managing Servers Running Windows Server 2012."

## 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, like the one shown in Figure 3-8. 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.

**SERVICES**  
All services | 151 total

Server Name	Display Name	Service Name	Status	Start Type
CORPSEVER172	Print Spooler	Spooler	Running	Automatic
CORPSEVER172	SNMP Trap	SNMPTRAP	Stopped	Manual
CORPSEVER172	SSDP Discovery	SSDPSRV	Stopped	Disabled
CORPSEVER172	Software Protection	sppsvc	Start Pending	Automatic (Delayed Start)
CORPSEVER172	Shell Hardware Detection	ShellHWDetection	Running	Automatic
CORPSEVER172	System Event Notification Service	SENS	Running	Automatic

**FIGURE 3-8** Use the Services panels in Server Manager to manage services on local and remote servers.

The columns on the Services panel can be adjusted by pressing and holding or right-clicking any column header and then tapping or 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.
- **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.
- **Start Type** The startup setting for the service. Automatic services are started at bootup. Users or other services start manual services. Disabled services are turned off and can't be started while they remain disabled.

**TIP** 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 pressing and holding or right-clicking any column header, selecting **Group By**, and then selecting your grouping option.

## Navigating Services in Computer Management

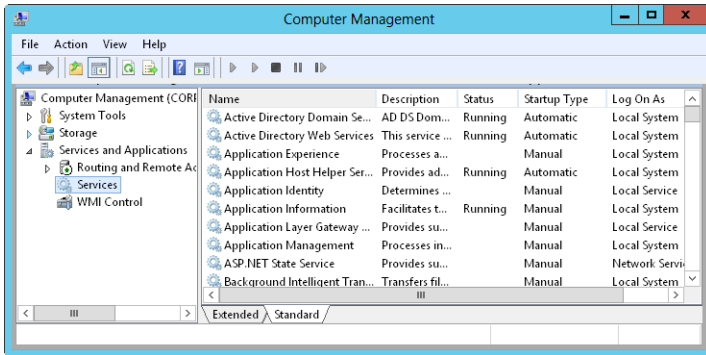
For quick and easy management of any service on a remote server, you can use the Services node in Computer Management. You can open Computer Management and automatically connect to a remote server from Server Manager. To do this, follow these steps:

1. Select All Servers or any server group node in the left pane.
2. On the Servers panel, press and hold or right-click the server to which you want to connect.
3. Tap or click Computer Management.

**TIP** When you are working with remote servers in Computer Management, many features rely on remote management and appropriate firewall exceptions being enabled as discussed in Chapter 2. 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, press and hold or 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, press and hold or right-click the server to which you want to connect and then select 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 as shown in Figure 3-9. 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-tap or 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 Chapter 2.
- **Description** A short description of the service and its purpose.
- **Status** Indicates whether the status of the service is running, paused, or stopped. (Stopped is indicated by a blank entry.)
- **Startup Type** The startup setting for the service. Automatic services are started at bootup. 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.



**FIGURE 3-9** Use the Services pane to manage services on local and remote computers.

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. Tap or click Start to start a stopped service. Tap or 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, you'll see a description that details the service's purpose.

**NOTE** Both the operating system and a user can disable services. Generally, Windows Server 2012 disables a service if a possible conflict with another service exists.

## Starting, Stopping, and Pausing Services

As an administrator, you often have to start, stop, or pause services. To start, stop, or pause a service, press and hold or 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. Additionally, if you pause a service, you can use the Resume option to resume normal operation.

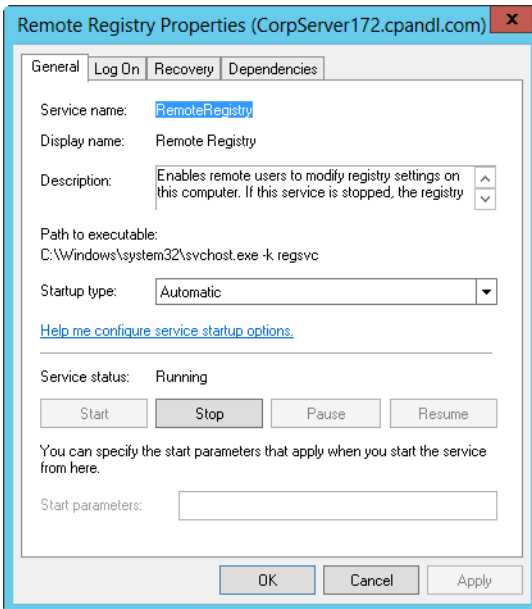
**NOTE** When services that are set to start automatically fail, the status is listed as blank, and you usually receive notification in a pop-up dialog box. Service failures can also be logged to the system's event logs. In Windows Server 2012, you can configure actions to handle service failure automatically. For example, you can have Windows Server 2012 attempt to restart the service for you. For details, see "Configuring Service Recovery" later in this chapter.

## 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. Press and hold or 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, as shown in Figure 3-10:
  - **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.



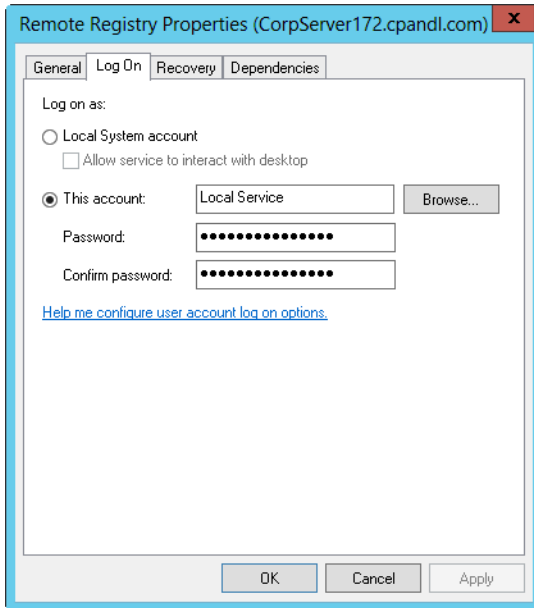
**FIGURE 3-10** Configure service startup options by using the General tab's Startup Type list.

3. Tap or click OK.

## 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, press and hold or right-click the service you want to configure, and then choose Properties.
2. Select the Log On tab, shown in Figure 3-11.



**FIGURE 3-11** Use the Log On tab to configure the service logon account.

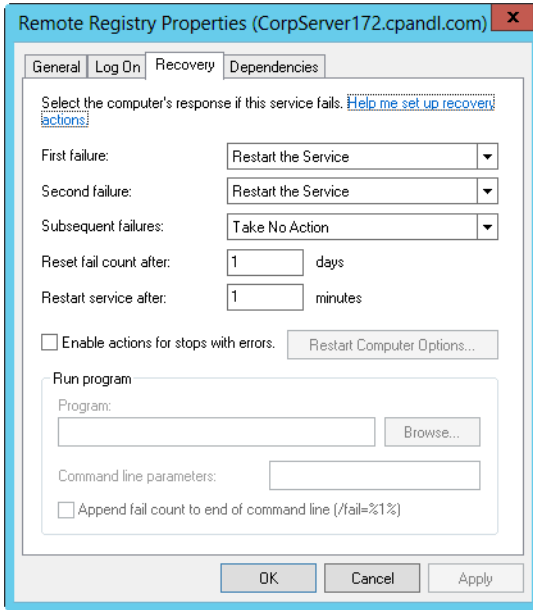
3. Select Local System Account if you want the service to log on using the system account (the default for most services). If the service provides a user interface that can be manipulated, select Allow Service To Interact With Desktop to allow users to control the service's interface.
4. Select This Account if you want the service to log on using a specific user account. Be sure to type an account name and password in the text boxes provided. Use the Browse button to search for a user account if necessary.
5. Tap or 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 normal 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, press and hold or right-click the service you want to configure, and then choose Properties.
2. Tap or click the Recovery tab, shown in Figure 3-12.



**FIGURE 3-12** Use the Recovery tab to specify actions that should be taken in case of service failure.

**NOTE** Windows Server 2012 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, such as DCOM Server Process Launcher and Group Policy Client, are configured to restart the computer if the service fails. You cannot change these settings because they are not available.

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 command-line 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 reboot 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. Tap or 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 she can't perform task X anymore, you can enable the related service again if necessary.

To disable a service, follow these steps:

1. In Computer Management, press and hold or right-click the service you want to configure, and then choose Properties. On the General tab, select Disabled in the Startup Type list.
2. Disabling a service doesn't stop a running service. It only prevents it from being started the next time the computer is booted, meaning that the security risk still exists. To address this, tap or click Stop on the General tab in the Properties dialog box and then tap or click OK.

# Event Logging and Viewing

---

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. For more information, see “Managing Your Servers Remotely” in Chapter 2.

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 you’ll see 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 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 “Configuring User Rights Policies” in Chapter 8, “Creating User and Group Accounts.”

- **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 startup. The default location is %SystemRoot%\System32\Winevt\Logs\System.evtx.

**SECURITY ALERT** As administrators, we tend to monitor the application and system logs the most—but don’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 you'll see include the following:



- **DFS Replication** This log records Distributed File System (DFS) replication activities. The default location is %SystemRoot%\System32\Winevt\Logs\DfsReplication.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, you'll see supplemental logs 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, like the one shown in Figure 3-13. When you select the server you want to work with in the Servers panel, its events are listed in the Events panel. You can use this panel as follows:

- For a server you are logged on to locally, you can use the Events panel in the Local Server node or the All Servers node to view recent warning and error events in the application and system logs.
- Automatically created server group nodes are organized by server roles, such as AD DS or DNS, and you'll be able to view recent error and warning events in logs related to the server role, if applicable. Not all roles have associated logs, but some roles, like AD DS, have multiple associated logs.
- For custom server groups created by you or other administrators, you'll be able to use the related Events panel to view recent warning and error events in the application and system logs.

**EVENTS**  
All events | 17 total

Filter   

Server Name	ID	Severity	Source	Log	Date
CORPSEVER172	5008	Error	DFSR	DFS Replication	5/17,
CORPSEVER172	5014	Warning	DFSR	DFS Replication	5/17,
CORPSEVER172	1308	Warning	Microsoft-Windows-ActiveDirectory_DomainService	Directory Service	5/17,
CORPSEVER172	1308	Warning	Microsoft-Windows-ActiveDirectory_DomainService	Directory Service	5/17,
CORPSEVER172	5002	Error	DFSR	DFS Replication	5/17,
CORPSEVER172	1308	Warning	Microsoft-Windows-ActiveDirectory_DomainService	Directory Service	5/17,

**FIGURE 3-13** Use the Events panels in Server Manager to track errors and warnings.

The columns on the Events panel can be adjusted by pressing and holding or right-clicking any column header and then tapping or 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 pressing and holding or right-clicking any column header, selecting **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. For more information, see “Managing Your Servers Remotely” in Chapter 2.

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, press and hold or right-click the server to which you want to connect.
3. Tap or 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 as shown in Figure 3-14.

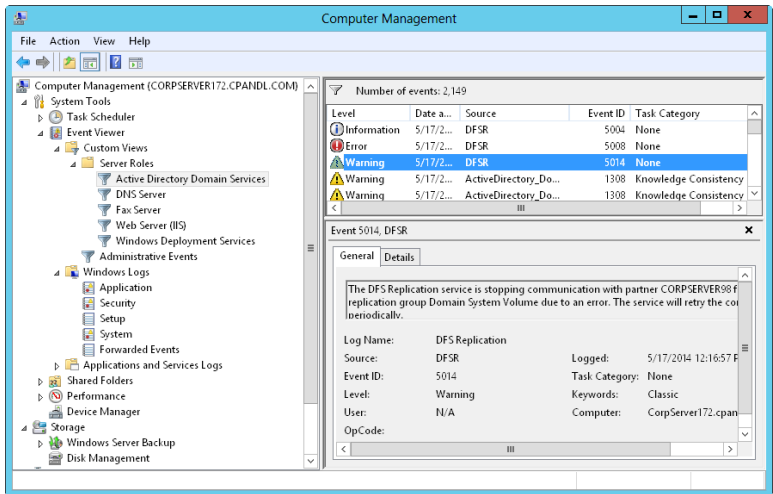


FIGURE 3-14 Event Viewer displays events for the selected log or custom view.

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, you should now see a list of all warning and error events for the server.
  - 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, you should now see a list of all events for the selected role.
  - 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.

As shown in Figure 3-14, entries in the main pane of Event Viewer provide a quick overview of when, where, and how an event occurred. To obtain detailed information on 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 shutting down 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:

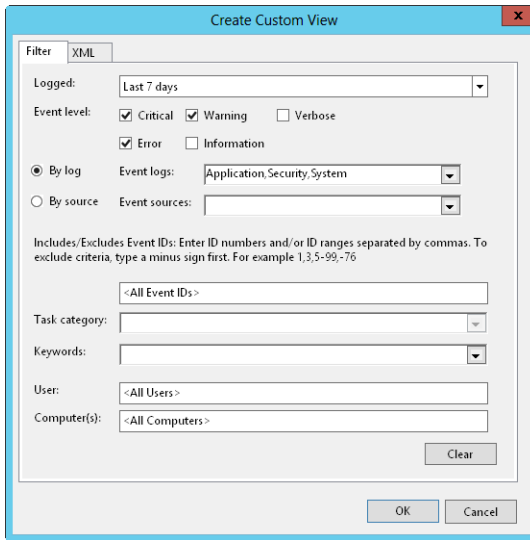
- **Source** The application, service, or component that logged the event
- **Event ID** Generally, a numeric identifier for the specific event, which could be helpful when searching knowledge bases
- **Task Category** The category of the event, which is almost always set to None, but is sometimes used 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, you'll see a list of all errors and warnings for all logs. When you expand the Server Roles node and then select a role-specific view, you'll see a list of all events for the selected role.

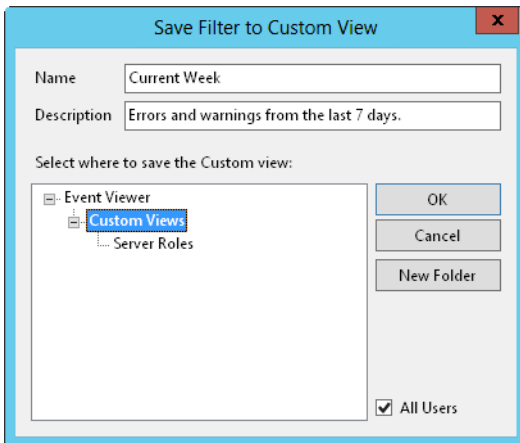
If you want to create a custom view of your own, you can do so in Computer Management by following these steps:

1. In the left pane, press and hold or right-click the Custom Views node and then tap or click Create Custom View. This opens the dialog box shown in Figure 3-15.



**FIGURE 3-15** You can filter logs so that only specific events are displayed.

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. You can 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 tap or click OK, Windows displays the Save Filter To Custom View dialog box, shown in Figure 3-16.



**FIGURE 3-16** Save the filtered view.

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 tapping or clicking New Folder, entering a name for the folder, and then tapping or clicking OK.
9. Tap or click OK to close the Save Filter To Custom View dialog box. You should now see a filtered list of events. Review these events carefully, and take steps to correct any problems that exist.

If you want to see 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. You should now see a list of event logs.
2. Press and hold or right-click the log you want to work with, and then tap or click Filter Current Log. This opens a dialog box similar to the one shown earlier in Figure 3-15.
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 Source 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.



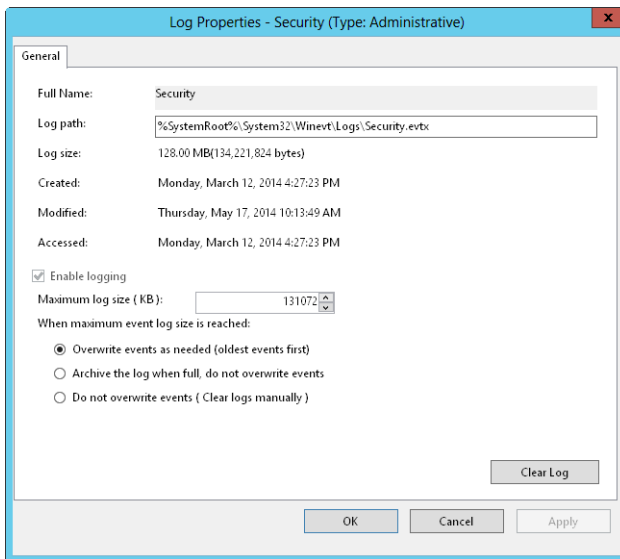
7. Tap or click OK. You should now see a filtered list of events. Review these events carefully, and take steps to correct any problems that exist. To clear the filter and see all events for the log, tap or click Clear Filter in the Actions pane or on the Action menu.

## Setting Event Log Options

Log options allow you to control the size of event logs as well as 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. You should now see a list of event logs.
2. Press and hold or right-click the event log whose properties you want to set, and then tap or click Properties on the shortcut menu. This opens the dialog box shown in Figure 3-17.



**FIGURE 3-17** Configure log settings according to the level of auditing on the system.

3. Type or set a maximum size in kilobytes (KB) in the Maximum Log Size text 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. Tap or click OK when you have 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.

## 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. You should now see a list of event logs.
2. Press and hold or right-click the event log whose properties you want to set, and then tap or click Clear Log on the shortcut menu.
3. Choose Save And Clear to save a copy of the log before clearing it. Choose Clear to continue without saving the log file.

## Archiving Event Logs

On key systems such as domain controllers and application servers, you'll want to keep several months' worth of logs. 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.

### Archive Log Formats

Logs can be archived in four formats:

- Event files (.evt) format for access in Event Viewer
- Tab-delimited text (.txt) format for access in text editors or word processors or to import into spreadsheets and databases
- Comma-delimited text (.csv) format for importing into spreadsheets or databases
- XML (.xml) format 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,07/21/14 3:43:24 PM,DNS Server,2,None,The DNS server has started.

Error,07/21/14 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 Archive The Log When Full, Do Not Overwrite Events. You can create a log archive manually 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. You should now see a list of event logs.
2. Press and hold or right-click the event log you want to archive, and then tap or click Save All Events As on the shortcut menu.
3. In the Save As dialog box, select a directory and type a log file name.
4. In the Save As Type list, Event Files (\*.evtx) is the default file type. Select the log format you want to use, and then choose Save. Note that you might not be able to use the .evtx format to save events from a remote computer to a local folder. In this case, you need to save the events to the local computer in a different file format, such as .xml. Otherwise, save the events in .evtx format on the remote computer.
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 tap or click OK. Otherwise, just tap or click OK to save the log without display information.

**NOTE** 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 and the period of the archive. For example, if you're archiving the system log file for January 2014, you might want to use the file name System Log January 2014.

**TIP** The best format to use for archiving is the .evtx format. Use this format 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 comma-delimited 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 in the .evtx format, 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.

## 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 press and hold or right-click the Event Viewer node. From the shortcut menu, select 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. Tap or click Open. If you are prompted about converting the log to the new event log format, tap or click Yes.
4. Windows displays the Open Saved Log dialog box. Type 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 tapping or clicking New Folder, entering a name for the folder, and then tapping or clicking OK.
6. Tap or click OK to close the Open Saved Log dialog box. You should now see the contents of the saved log.

**TIP** To remove the saved log from Event Viewer, tap or click Delete in the Actions pane or on the Action menu. When prompted to confirm, tap or click Yes. The saved log file still exists in its original location.

## Monitoring Server Performance and Activity

---

Monitoring a server isn't something you should do haphazardly. You need to have a clear plan—a set of goals you hope to achieve. Let's take a look at the reasons you might want to monitor a server and the tools you can use to do this.

### Why Monitor Your Server?

Troubleshooting server performance problems is a key reason for monitoring. For example, users might be having problems connecting to the server, and you might want to monitor the server to troubleshoot these problems. Your goal is to track down the problem by using the available monitoring resources and resolve it.

Another common reason for wanting to monitor a server is to improve server performance. You do this by improving disk I/O, reducing CPU usage, and cutting down the network traffic load on the server. Unfortunately, you often need to make trade-offs when it comes to resource usage. For example, as the number of users accessing a server grows, you might not be able to reduce the network traffic load, but you might be able to improve server performance through load balancing or by distributing key data files on separate drives.

## Getting Ready to Monitor

Before you start monitoring a server, you might want to establish baseline performance metrics for your server. 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.

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.

These procedures are examined later in this chapter. Although you should usually develop a monitoring plan, sometimes you might not want to go through all these steps to monitor your server. 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 your servers include the following:

- **Performance Monitor** Used to configure counters to watch resource usage over time. 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** Use information in the event logs to troubleshoot system-wide 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 tapping or clicking the Open Resource Monitor button on Task Manager's Performance tab. As shown in Figure 3-18, 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). If you expand the

CPU entry (by tapping or clicking the options button), you'll see 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. If you expand the Disk entry below the graph (by tapping or clicking the options button), you'll see a list of currently running executables that are performing or have performed I/O operations 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. If you expand the Network entry below the graph (by tapping or clicking the options button), you'll see a list of currently running executables that are transferring or have transferred data on the network 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. If you expand the Memory entry below the graph (by tapping or clicking the options button), you'll see 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.

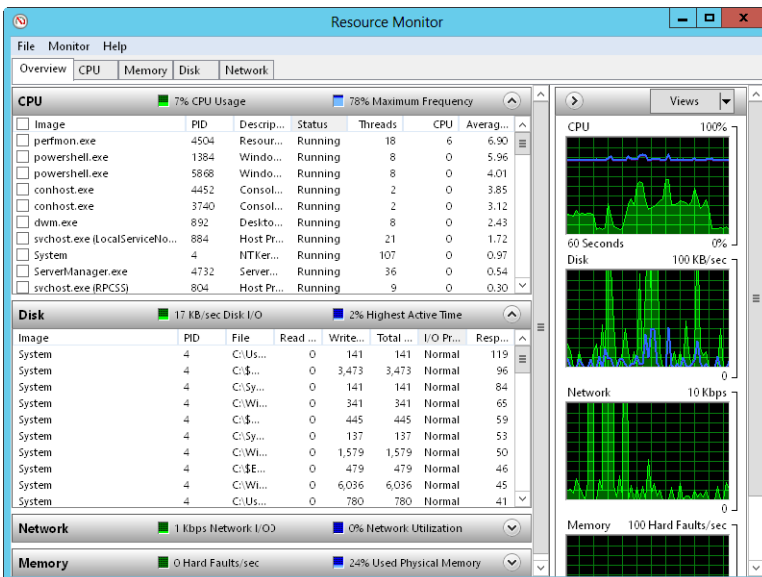
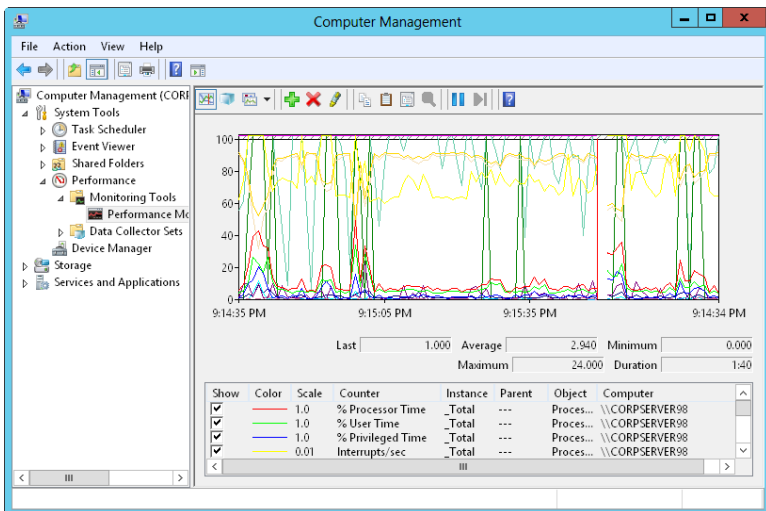


FIGURE 3-18 Review resource usage on the server.

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 can update these counters when you install additional services and add-ons for the application as well.

In Server Manager, you can access Performance Monitor in a standalone console by tapping or clicking Tools and then tapping or clicking Performance Monitor. In Computer Management, you can access the tool as a snap-in under the System Tools node. Expand System Tools, Performance, Monitoring Tools and then select Performance Monitor.

As Figure 3-19 shows, 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 you'll see when you work with Performance Monitor, the tracking information is most valuable when you record performance information in a log file so that it can be played back. Also, Performance Monitor is helpful when you configure alerts to send messages when certain events occur.



**FIGURE 3-19** Review performance measurements for the server.

Windows Server 2012 also includes Reliability Monitor. To access Reliability Monitor, follow these steps:

1. In Control Panel, tap or click Review Your Computer's Status under the System And Security heading.
2. In Action Center, expand the Maintenance panel, and then tap or click View Reliability History.

Alternatively, you can run Reliability Monitor by entering **perfmon /rel** at a command prompt or in the Apps Search box.

Reliability Monitor tracks changes to the server and compares them to changes in system stability. In this way, you can see a graphical representation of the relationship between changes in the system configuration and changes in system stability. By recording software installation, software removal, application failures, hardware failures, Windows failures, and key events regarding the configuration of the server, you can see a timeline of changes in both the server and its reliability and then use this information to pinpoint changes that are causing problems with stability. For example, if you see a sudden drop in stability, you can tap or click a data point and then expand the related data set to find the specific event that caused the drop in stability.

Although reliability monitoring is enabled by default for Windows clients, it might be disabled for Windows servers. When you open Reliability Monitor on a server where reliability monitoring is disabled, you'll see an information panel telling you to [click here](#) to see how to turn on or reconfigure RACTask. The RACTask is a scheduled task that runs in the background to collect reliability data.

## Choosing Counters to Monitor

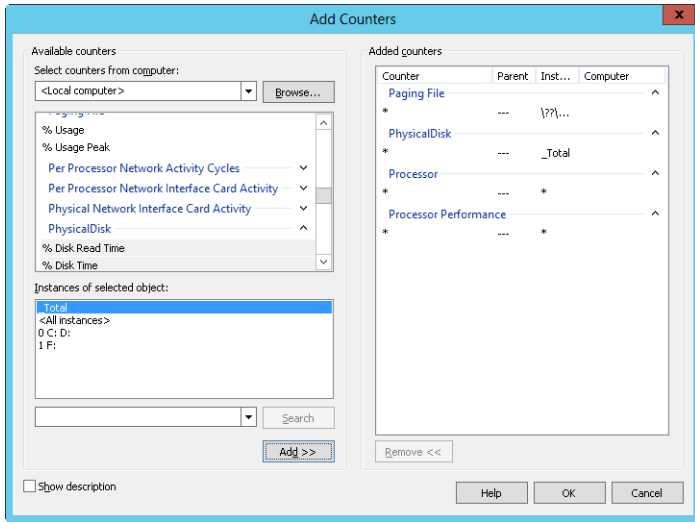
Performance Monitor displays information only for counters you're tracking. Several thousand counters are available, and you'll find counters related to just about every server role you've installed. The easiest way to learn about these counters is to read the explanations available in the Add Counters dialog box. Start Performance Monitor, tap or 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.

To select which counters you want to monitor, follow these steps:

1. Performance Monitor has several views and view types. Be sure that current activity is displayed by tapping or clicking View Current Activity on the toolbar or pressing Ctrl+T. You can switch between the view types (Line, Histogram Bar, and Report) by tapping or clicking Change Graph Type or pressing Ctrl+G.
2. To add counters, tap or click Add on the toolbar or press Ctrl+N. This displays the Add Counters dialog box, shown in Figure 3-20.





**FIGURE 3-20** Select the objects and counters you want to monitor.

3. 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.

**NOTE** 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.

4. In the Available Counters panel, performance objects are listed alphabetically. If you select an object entry by tapping or clicking it, all related counters are selected. If you expand an object entry, you can see all the related counters and then select individual counters by tapping or 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.
5. When you select an object or any of its counters, you see the related instances. 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.
6. When you've selected an object or a group of counters for an object as well as the object instances, tap or click Add to add the counters to the graph.
7. Repeat steps 4–6 to add other performance parameters.
8. Tap or click OK when you have finished.

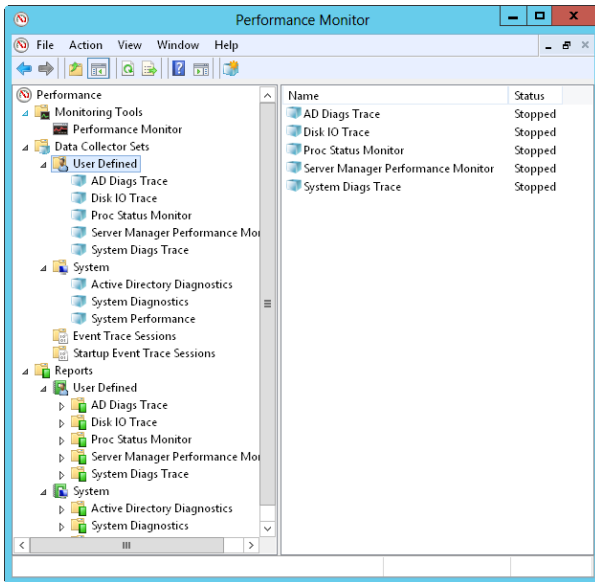
**TIP** Don't try to chart too many counters or counter instances at once. You'll make the display too difficult to read, and you'll use system resources—namely, CPU time and memory—that might affect server responsiveness.

## Performance Logging

Windows Server 2008 R2 introduced data collector sets and reports. Data collector sets allow you to specify sets of performance objects and counters you want to track. Once you've created 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 earlier 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

In the Performance tool, you can review currently configured data collector sets and reports under the Data Collector Sets and Reports nodes, respectively. As shown in Figure 3-21, 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.



**FIGURE 3-21** Access data collector sets and reports.

## Creating and Managing Data Collector Sets

To view the currently configured data collector sets, select the Performance Monitor option in the Administrative Tools program group and then expand the Data Collector Sets node. You can 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, you'll see the related data collectors 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 pre-determined 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 pressing and holding or right-clicking it and 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 pressing and holding or right-clicking a startup data collector and selecting Start As Event Trace Session. You can delete a startup data collector by pressing and holding or right-clicking it and then tapping or clicking Delete.
- You can save a data collector as a template that can be used as the basis of other data collectors by pressing and holding or right-clicking the data collector and selecting Save Template. In the Save As dialog box, select a directory, type a name for the template, and then tap or 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 pressing and holding or right-clicking it and selecting Delete. If a data collector is running, you need to stop collecting data first and then delete the collector. Deleting a collector deletes the related reports as well.

## Collecting Performance Counter Data

Data collectors can be used 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.

To collect performance counter data, follow these steps:

1. In Performance Monitor, under the Data Collector Sets node, press and hold or right-click the User Defined node in the left pane, point to New, and then choose Data Collector Set.
2. In the Create New Data Collector Set Wizard, type a name for the data collector, such as **System Performance Monitor** or **Processor Status Monitor**. Note that if you type an invalid name, such as one with a non-alphanumeric character, you won't be able to continue.

3. Select the Create Manually option, and then tap or 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 Performance Counter check box, and then tap or click Next.
5. On the Which Performance Counters Would You Like To Log page, tap or 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 have finished selecting counters, tap or click OK.
6. 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. Tap or click Next when you are ready to continue.
7. On the Where Would You Like The Data To Be Saved page, type the root path to use for logging collected data. Alternatively, tap or click Browse, and then use the Browse For Folder dialog box to select the logging directory. Tap or click Next when you are ready to continue.

**BEST PRACTICES** 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.

8. 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, tap or click Change. Type the user name and password for the account, and then tap or click OK. User names can be entered in domain\username format, such as cpandl\williams for the Williams account in the Cpandl domain.
9. Select the Open Properties For This Data Collector Set option, and then tap or click Finish. This saves the data collector set, closes the wizard, and then opens the related Properties dialog box.
10. By default, logging is configured to start manually. To configure a logging schedule, tap or click the Schedule tab and then tap or click Add. You can now set the Active Range, Start Time, and run days for data collection.
11. 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).
12. Tap or click OK when you've finished setting the logging schedule and stop conditions. You can manage the data collector as explained previously.

**NOTE** You can configure Windows to run a scheduled task when data collection stops. You configure tasks to run on the Tasks tab in the Properties dialog box.

## Collecting Performance 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, press and hold or right-click the User Defined node in the left pane, point to New, and then choose 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.
3. Select the Create Manually option, and then tap or 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 Event Trace Data check box, and then tap or click Next.
5. On the Which Event Trace Providers Would You Like To Enable page, tap or click Add. Select an event trace provider to track, and then tap or click OK. By selecting individual properties in the Properties list and tapping or 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. Tap or click Next when you are ready to continue.
6. Complete steps 7–12 from the procedure in the previous section, "Collecting Performance 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, press and hold or right-click the User Defined node in the left pane, point to New, and then choose Data Collector Set.
2. In the Create New Data Collector Set Wizard, type a name for the data collector, such as **AD Registry** or **Registry Adapter Info**.
3. Select the Create Manually option, and then tap or 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 tap or click Next.

5. On the Which Registry Keys Would You Like To Record page, tap or click Add. Type the registry path to track. Repeat this process to add other registry paths to track. Tap or click Next when you are ready to continue.
6. Complete steps 7–12 from the procedure in the section “Collecting Performance Counter Data.”

## Viewing Data Collector Reports

When you troubleshoot problems, 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 pressing and holding or right-clicking a data collector set and selecting 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, as shown in Figure 3-22.

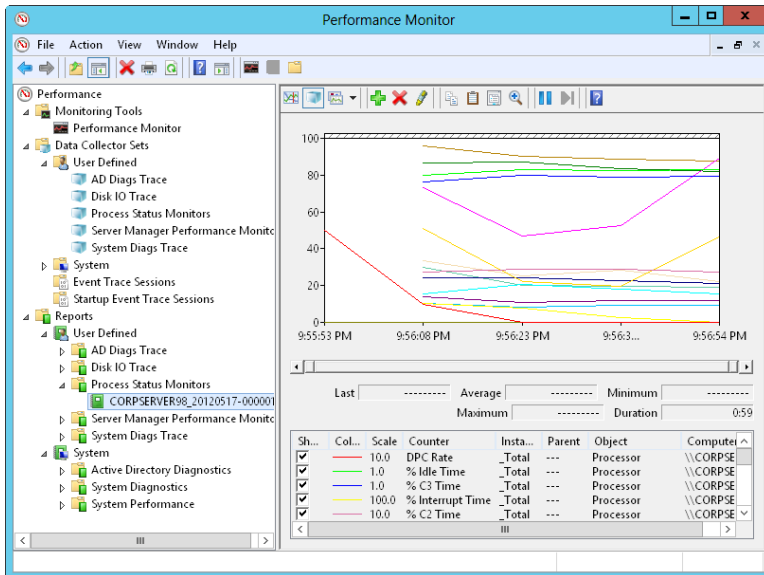


FIGURE 3-22 View data collector reports.

You can modify the report details using the following techniques:

1. In the monitor pane, press Ctrl+Q or tap or click the Properties button on the toolbar. This displays the Performance Monitor Properties dialog box.
2. Tap or click the Source tab.
3. Specify data sources to analyze. Under Data Source, tap or click Log Files and then tap or click Add to open the Select Log File dialog box. You can now select additional log files to analyze.
4. Specify the time window you want to analyze. Tap or 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.
5. Tap or click the Data tab. You can now select counters to view. Select a counter, and then tap or click Remove to remove it from the graph view. Tap or 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 you don't see a counter you want to work with, you need to modify the data collector properties, restart the logging process, and then check the logs at a later date.

6. Tap or click OK. In the monitor pane, tap or 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, press and hold or right-click the User Defined node in the left pane, point to New, and then choose Data Collector Set.
2. In the Create New Data Collector Set Wizard, type a name for the data collector, such as **Processor Alert** or **Disk IO Alert**.
3. Select the Create Manually option, and then tap or click Next.
4. On the What Type Of Data Do You Want To Include page, select the Performance Counter Alert option and then tap or click Next.
5. On the Which Performance Counters Would You Like To Monitor page, tap or 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. Tap or click OK when you have finished.

6. In the Performance Counters panel, select the first counter, and then use the Alert When Value Is text 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 95 percent, select Over, and then type 95. Repeat this process to configure other counters you've selected.
7. Complete steps 7–12 from the procedure in the section “Collecting Performance Counter Data.”

## 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. I'll examine the following areas:

- Memory usage and caching
- Processor utilization
- Disk I/O
- Network bandwidth and connectivity

## Monitoring and Tuning Memory Usage

Memory is often the source of performance problems, and you should always rule out memory problems before examining other areas of the system. Systems 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, as discussed in “Setting Application Performance” in Chapter 2.

Using the monitoring techniques discussed previously in this chapter, you can determine how the system is using memory and check for problems. Table 3-1 provides an overview of counters you'll want to track to uncover memory, caching, and virtual memory (paging) bottlenecks. The table is organized by issue category.



**TABLE 3-1** Uncovering Memory-Related Bottlenecks

ISSUE	COUNTERS TO TRACK	DETAILS
Physical and virtual memory usage	Memory\Available Kbytes Memory\Committed 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 need to add memory as well. In general, you want the committed bytes value to be no more than 75 percent of the total physical memory.
Memory page faults	Memory\Page Faults/sec Memory\Pages Input/sec Memory\Page Reads/sec	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 <i>soft page fault</i> . If the requested page must be retrieved from disk, the fault is called a <i>hard page fault</i> . 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.

ISSUE	COUNTERS TO TRACK	DETAILS
Memory paging	Memory\Pool Paged Bytes Memory\Pool Nonpaged Bytes	These counters 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.

## Monitoring and Tuning Processor Usage

The 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 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 the performance counters listed in Table 3-2. Be sure to monitor these counters for each CPU installed on the server.

**TABLE 3-2** Uncovering Processor-Related Bottlenecks

ISSUE	COUNTERS TO TRACK	DETAILS
Thread queuing	System\Processor Queue Length	This counter 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.

ISSUE	COUNTERS TO TRACK	DETAILS
CPU usage	Processor\% Processor Time	This counter 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.

## Monitoring and Tuning Disk I/O

With today's high-speed disks, the disk throughput rate is rarely the cause of a bottleneck. That said, accessing memory is much faster than accessing disks. So, if the 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 "Monitoring and Tuning Memory Usage."

In addition to memory tuning, you can monitor some counters to gauge disk I/O activity. Specifically, you should monitor the counters listed in Table 3-3.

**TABLE 3-3** Uncovering Drive-Related Bottlenecks

ISSUE	COUNTERS TO TRACK	DETAILS
Overall drive performance	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.
Disk I/O	PhysicalDisk\Disk Writes/sec, PhysicalDisk\Disk Reads/sec, PhysicalDisk\Avg. DiskWrite Queue Length, PhysicalDisk\Avg. DiskRead Queue Length, PhysicalDisk\CurrentDisk 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.

## Monitoring and Tuning Network Bandwidth and Connectivity

No other factor matters 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, 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 and the amount of bandwidth available to your servers 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.

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.

To determine the throughput and current activity on a server's network cards, you can check the following counters:

- Network\Bytes Received/sec
- Network\Bytes Sent/sec
- Network\Bytes Total/sec
- Network 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. Keep in mind that you should compare these values in conjunction with Physical-Disk\% 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.

# Automating Administrative Tasks, Policies, and Procedures

- Understanding Group Policies **138**
- Navigating Group Policy Changes **141**
- Managing Local Group Policies **143**
- Managing Site, Domain, and Organizational Unit Policies **147**
- Maintaining and Troubleshooting Group Policy **158**
- Managing Users and Computers with Group Policy **173**

Performing routine tasks day after day, running around policing systems, and walking users through the basics aren't efficient uses of your time. You'd be much more effective if you could automate these chores and focus on issues that are more important. Support services are all about increasing productivity and allowing you to focus less on mundane matters and more on what's important.

Microsoft Windows Server 2012 includes many roles, role services, and features that help you support server installations. You can easily install and use some of these components. If you need an administrative tool to manage a role or feature on a remote computer, you can select the tool to install as part of the Remote Server Administration Tools feature. If a server has a wireless adapter, you can install the Wireless LAN Service feature to enable wireless connections. Beyond these and other basic support components, you can use many other support features, including the following:

- **Automatic Updates** Ensures that the operating system is up to date and has the most recent security updates. If you update a server by using Microsoft Update instead of the standard Windows Updates, you can get updates for additional products. By default, Automatic Updates is installed but not enabled on servers running Windows Server 2012. You can configure Automatic Updates by using the Windows Update utility in Control

Panel. In Control Panel, System And Security, tap or click Turn Automatic Updating On Or Off. To learn how to configure Automatic Updates through Group Policy, see “Configuring Automatic Updates” later in this chapter.

- **BitLocker Drive Encryption** Provides an extra layer of security for a server’s hard disks. This protects the disks from attackers who have physical access to the server. BitLocker encryption can be used on servers with or without a Trusted Platform Module (TPM). When you add this feature to a server using the Add Roles And Features Wizard, you can manage it using the BitLocker Drive Encryption utility in Control Panel. In Control Panel, System And Security, tap or click BitLocker Drive Encryption. Windows Server 2008 R2 and later (like Windows 7 and later) include BitLockerToGo, which allows you to encrypt USB flash drives. If your server doesn’t have BitLocker, run the BitLocker To Go Reader, which is stored in an unencrypted area of the encrypted USB flash drive.
- **Remote Assistance** Provides an assistance feature that allows an administrator to send a remote assistance invitation to a more senior administrator. The senior administrator can then accept the invitation to view the user’s desktop and temporarily take control of the computer to resolve a problem. When you add this feature to a server using the Add Roles And Features Wizard, you can manage it by using options on the Remote tab of the System Properties dialog box. In Control Panel, System And Security, tap or click Allow Remote Access under the System heading to view the related options.
- **Remote Desktop** Provides a remote connectivity feature that allows you to connect to and manage a server from another computer. By default, Remote Desktop is installed but not enabled on servers running Windows Server 2012. You can manage the Remote Desktop configuration with the options on the Remote tab of the System Properties dialog box. In Control Panel, System And Security, tap or click Allow Remote Access under the System heading to view the related options. You can establish remote connections using the Remote Desktop Connection utility.
- **Task Scheduler** Allows you to schedule execution of one-time and recurring tasks, such as tasks for performing routine maintenance. Windows Server 2012 makes extensive use of the scheduled task facilities. You can view and work with scheduled tasks in Computer Management. Expand the System Tools, Task Scheduler, and Task Scheduler Library nodes to view configured scheduled tasks.
- **Desktop Experience** This subfeature of User Interfaces And Infrastructure installs Windows desktop functionality on the server. You can use this feature when you use Windows Server 2012 as your desktop operating system. When you add this feature using the Add Roles And Features Wizard, the server’s desktop functionality is enhanced, and the following programs are installed as well: Windows Media Player, Desktop Themes, Video for Windows (AVI support), Disk Cleanup, Sound Recorder, Character Map, and Snipping Tool.

- **Windows Firewall** Helps protect a computer from attack by unauthorized users. Windows Server includes a basic firewall called Windows Firewall and an advanced firewall called Windows Firewall With Advanced Security. By default, the firewalls are not enabled on server installations. To access the basic firewall, tap or click Windows Firewall in Control Panel. To access the advanced firewall, select Windows Firewall With Advanced Security on the Tools menu in Server Manager.
- **Windows Time** Synchronizes the system time with world time to ensure that the system time is accurate. You can configure computers to synchronize with a specific time server. The way Windows Time works depends on whether a computer is a member of a domain or a workgroup. In a domain, domain controllers are used for time synchronization, and you can manage this feature through Group Policy. In a workgroup, Internet time servers are used for time synchronization, and you can manage this feature through the Date And Time utility.

You can configure and manage these support components in the same way on Windows 8 and Windows Server 2012. You'll find extensive coverage of these support components in *Windows 8 Administration Pocket Consultant* (Microsoft Press, 2012).

Many other components provide support services. However, you need these additional support services only in specific scenarios. You can use IP Address Management (IPAM) servers when you want to manage your IP address space and track IP address usage trends. You can use Remote Desktop Services when you want to allow users to run applications on a remote server. You can use Windows Deployment Services when you want to enable automated deployment of Windows-based operating systems. The one always-on support service you must master to succeed with Windows Server 2012 is Group Policy.

**REAL WORLD** The Start screen's options panel has a Search option, which can have Apps, Settings, or Files as a focus. When you press the Windows key and type, the text is entered into the Search box. Because the default focus is for an Apps Search, this allows you to quickly search for a program installed on a server.

Throughout this text, when I refer to entering something in the Apps Search box, I'm referring to entering search text with Apps as a focus. As you enter text, matching results are displayed. When you press Enter, Windows runs the currently selected result. You can use the Apps Search to pass in commands with parameters and options as well. Simply type the command along with its parameters and options as you would at a command prompt.

Want to run Windows PowerShell commands from the Apps Search box? Simply type **powershell** and then enter your command.

# Understanding Group Policies

---

Group policies simplify administration by giving administrators centralized control over privileges, permissions, and capabilities of both users and computers. Through group policies, you can do the following:

- Control access to Windows components, system resources, network resources, Control Panel utilities, the desktop, and the Start screen. See “Using Administrative Templates to Set Policies” later in this chapter for more details.
- Create centrally managed directories for special folders, such as a user’s Documents folder. See “Centrally Managing Special Folders” later in the chapter for more details.
- Define user and computer scripts to run at specified times. This is covered in “User and Computer Script Management” later in the chapter.
- Configure policies for account lockout and passwords, auditing, user rights assignment, and security. Many of these topics are covered in “User Account Setup and Organization” in Chapter 8, “Creating User and Group Accounts.”

The sections that follow explain how you can work with and apply group policies.

## Group Policy Essentials

You can think of a policy as a set of rules that helps you manage users 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.

To understand group policies, you need to know a bit about the structure of Active Directory. 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* (OUs). 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 policies apply only to systems running Windows 2000 and later versions of Windows. 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 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 a top-level 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 is site, domain, 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.

As you might expect, you can override inheritance. To do this, you specifically assign a policy setting for a child container that is different from the policy setting for the parent. 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" later in this chapter.

## In What Order Are Multiple Policies Applied?

When multiple policies are in place, 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

If policy settings conflict, the policy settings applied later have precedence and overwrite previously set 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 "Blocking, Overriding, and Disabling Policies" later in this chapter.

## When Are Group Policies Applied?

As you'll discover when you start working with group policies, 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 and automatically within a 90 to 120 minute period. You can change this behavior by setting a Group Policy refresh interval, as discussed in "Refreshing Group Policy" later in this chapter. To do this, open a prompt and type **gpupdate**.

**REAL WORLD** 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 type **gpupdate /logoff** at a prompt or in the Apps 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 prompt or in the Apps Search box to restart the computer after the refresh.

## Group Policy Requirements and Version Compatibility

Group policies apply only to systems running professional and server versions of Windows. As you might expect, each new version of the Windows operating system has brought with it changes to Group Policy. Sometimes these changes have made older policies obsolete on newer versions of Windows. In this case, the policy works only on specific versions of Windows, such as only on Windows XP Professional and Windows Server 2003.

Generally speaking, most policies are forward compatible. This means that in most cases, policies introduced in Windows Server 2003 can be used on Windows 7 and later, as well as Windows Server 2008 and later. It also means that policies for Windows 8 and Windows Server 2012 usually aren't applicable to earlier releases of Windows. 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.

How will you know if a policy is supported on a particular version of Windows? Easy. The Properties dialog box for each policy setting has a Supported On text box. This text-only 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). You'll see a Requirements entry that lists compatibility.

You can also install new policies when you add a service pack, install Windows applications, or add Windows components. This means that you'll see a wide range of compatibility entries.

## Navigating Group Policy Changes

---

In an effort to streamline management of Group Policy, Microsoft removed management features from Active Directory–related tools and moved to a primary console called the Group Policy Management Console (GPMC) starting with Windows Vista and Windows Server 2008. The GPMC is a feature you can add to any installation of Windows Server 2008 or later by using the Add Roles And Features Wizard. The GPMC is available on Windows Vista and later when you install the Remote Server Administration Tools (RSAT). Once you add the GPMC to a computer, it is available on the Tools menu in Server Manager.

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 policy-management 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.

If you've worked with earlier versions of Windows, you might also be familiar with the Group Policy Object Editor (GPOE). With Windows Server 2003 and earlier versions of Windows, the GPOE is the primary editing tool for policy objects. The Group Policy Object Editor, 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.

You can manage policy settings for Windows Vista and later only from computers running Windows Vista or later. The reason for this is that the GPOE and the GPMC for Windows Vista and later releases were updated to work with the XML-based administrative templates format called ADMX.

**NOTE** You cannot use older versions of the policy editors with ADMX. You can edit GPOs using ADMX files only on a computer running Windows Vista or later.

Microsoft had many reasons for going to the ADMX format. The key reasons were to allow greater flexibility and extensibility. Because ADMX files are created using XML, the files are strictly structured and can be more easily and rapidly parsed during initialization. This can help to improve performance when the operating system processes Group Policy during the startup, logon, logoff, and shutdown phases, as well as during policy refreshes. Further, the strict structure of ADMX files makes it possible for Microsoft to continue in its internationalization efforts.

ADMX files are divided into language-neutral files ending with the .admx file extension and language-specific files ending with the .adml extension. The language-neutral 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 language-neutral 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 installed on computers running Windows Vista or later in the %SystemRoot%\PolicyDefinitions folder. Language-specific ADMX files are installed on computers running Windows 7 and Windows 8, as well as Windows Server 2008 R2 and Windows Server 2012 in the %SystemRoot%\PolicyDefinitions\*LanguageCulture* folder. Each subfolder is named 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. As long as you edit GPOs using Windows Vista or later, new GPOs will not contain ADM or ADMX files inside the GPO. For more information, see Chapter 2, “Deploying Group Policy,” in *Windows Group Policy Administrator's Pocket Consultant* (Microsoft Press, 2009).

When running in Windows Server 2008 or higher domain functional level, servers running Windows Server 2008 or later use Distributed File System (DFS) Replication Service for replicating Group Policy. With DFS replication, only the changes in GPOs are replicated, thereby eliminating the need to replicate an entire GPO after a change.

Unlike Windows XP and Windows Server 2003, Windows Vista and later releases use 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.

Computers running Windows Vista or later don't use the trace logging functionality in Userenv.dll and instead write Group Policy event messages to the System log. Further, the Group Policy operational log replaces previous Userenv logging. When you are troubleshooting Group Policy issues, you use the detailed event messages in the operational log rather than the Userenv log. In Event Viewer, you can access the operational log under Applications And Services Logs\Microsoft\Windows\GroupPolicy.

Computers running Windows Vista or later use Network Location Awareness instead of Internet Control Message Protocol (ICMP, or ping). With Network Location Awareness, a computer is aware of the type of network it is connected to and can be responsive to changes in the system status or network configuration. By using Network Location Awareness, the Group Policy client can determine the computer state, network state, and available network bandwidth for slow-link detection.

## Managing Local Group Policies

---

Computers running Windows Vista or later allow the use of multiple local Group Policy Objects on a single computer (as long as the computer is not a domain controller). Previously, computers had only one local GPO. Windows allows you to assign a different local GPO to each local user or general user type. This allows the application of policy to be more flexible and supports a wider array of implementation scenarios.

### Local Group Policy Objects

When computers are being used in a standalone 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.

Computers running Windows Vista or later have 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: Local Group Policy, Administrators and Non-Administrators local Group Policy, user-specific local Group Policy.

Because the available user-configuration settings are the same for all local GPOs, a setting in one GPO might conflict with a setting in another GPO. Windows resolves conflicts in settings by overwriting any previous setting with the last-read and most-current settings. The final setting is the one that Windows uses. When Windows resolves 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. To simplify domain administration, you can disable processing of local Group Policy Objects on computers running Windows Vista or later releases by enabling the Turn Off Local Group Policy Objects Processing policy setting in a domain GPO. In Group Policy, this setting is located under 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. Although a domain controller has a local Group Policy Object, you shouldn't edit its settings.

The quickest way to access the local GPO on a computer is to type the following command at a command prompt or in the Apps Search box:

```
gpedit.msc /gpcomputer: "%ComputerName%"
```

**NOTE** Because of the additional arguments passed in with the command, the command does not pass through to the command prompt properly from a PowerShell prompt. Enclose the arguments in single quotes to get them to pass through, as shown in the following example: `gpedit.msc '/gpcomputer: "%ComputerName%"'`.

This command starts the GPOE in a Microsoft Management Console (MMC) with its target set to the local computer. Here, `%ComputerName%` is an environment variable that sets the name of the local computer; it must be enclosed in double quotation marks as shown. To access the top-level local GPO on a remote computer, type the following command at a prompt or in the Apps Search box:

```
gpedit.msc /gpcomputer: "RemoteComputer"
```

Here `RemoteComputer` is the host name or fully qualified domain name (FQDN) of the remote computer. Again, the double quotation marks are required, as shown in the following example:

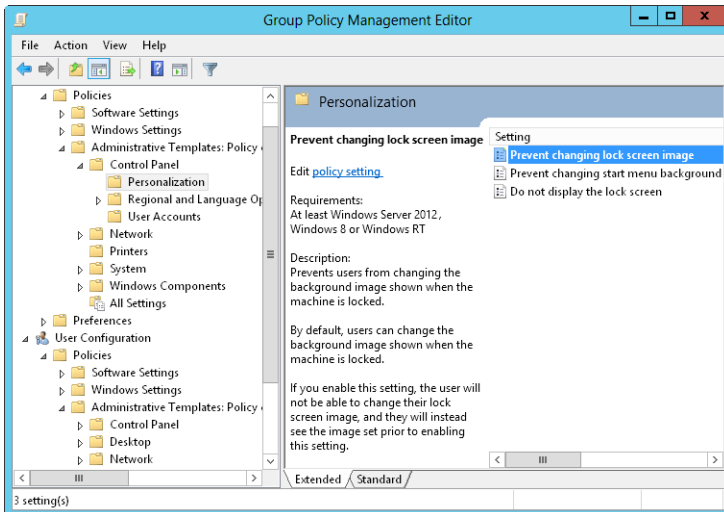
```
gpedit.msc /gpcomputer: "corpsvr82"
```

You can also manage the top-level local GPO on a computer by following these steps:

1. At a prompt or in the Apps Search box, type **mmc** and then press Enter.
2. In the Microsoft Management Console, tap or click File, and then tap or click Add/Remove Snap-In.
3. In the Add Or Remove Snap-Ins dialog box, tap or click Group Policy Object Editor and then tap or click Add.
4. In the Select Group Policy Object dialog box, tap or click Finish because the local computer is the default object. Tap or click OK.

As shown in Figure 4-1, you can now manage local policy settings using the options provided.

**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.



**FIGURE 4-1** Use the policy editor to manage local policy settings.

## Local Group Policy Object Settings

Local group policies are stored in the %SystemRoot%\System32\GroupPolicy folder on each Windows Server computer. In this folder, 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

**CAUTION** You shouldn'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, tap or click the View tab and then select Hidden Items. You also might want to select File Name Extensions.

## Accessing Administrator, Non-Administrator, and User-Specific Local Group Policy

By default, the only local policy object that exists on a computer is the Local Group Policy Object. You can create and manage other local objects as necessary (except on domain controllers). You can create or access the Administrator local Group Policy Object, the Non-Administrator local Group Policy Object, or a user-specific local Group Policy Object by following these steps:

1. At a prompt or in the Apps Search box, type **mmc** and then press Enter. In the Microsoft Management Console, tap or click File, and then tap or click Add/Remove Snap-In.
2. In the Add Or Remove Snap-Ins dialog box, tap or click Group Policy Object Editor, and then tap or click Add.
3. In the Select Group Policy Object dialog box, tap or click Browse. In the Browse For A Group Policy Object dialog box, tap or 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 Administrator local Group Policy Object.
  - Select Non-Administrators to create or access the Non-Administrator local Group Policy Object.
  - Select the local user whose user-specific local Group Policy Object you want to create or access.
5. Tap or 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 Windows Server computer. 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.



# Managing Site, Domain, and Organizational Unit Policies

---

When you deploy Active Directory Domain Services (AD DS), you can use Active Directory–based Group Policy. 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.

## Understanding Domain and Default Policies

When you work with Active Directory–based Group Policy, you’ll find that 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.

Typically, the Default Domain Policy GPO is the highest-precedence GPO linked to the domain level, and the Default Domain Controllers Policy GPO is the highest-precedence 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. These GPOs aren’t meant for general management of Group Policy.

The Default Domain Policy GPO is used only to manage the default Account Policies settings and, in particular, three specific areas of Account Policies: password policy, account lockout policy, and Kerberos policy. Several security options are managed through this GPO as well. 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. One way to override these settings is to create a GPO with the overriding settings and link it with a higher precedence to the domain container.

The Default Domain Controllers Policy GPO includes specific User Rights Assignment and Security Options settings that limit the ways domain controllers can be used. One way to override these settings is to create a GPO with the overriding settings and link it with a higher precedence to the Domain Controllers container.

To manage other areas of policy, you should create a GPO and link it to the domain or to an appropriate 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:

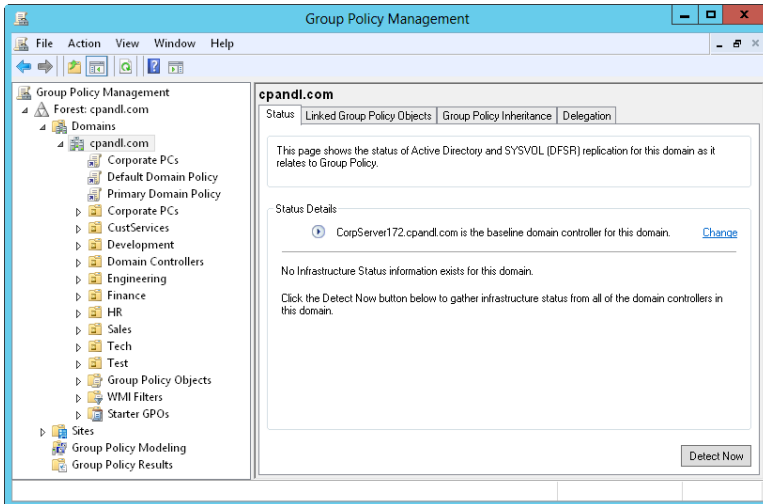
- **Machine** Stores computer scripts in the Script folder and registry-based policy information for HKEY\_LOCAL\_MACHINE (HKLM) in the Registry.pol file
- **User** Stores user scripts in the Script folder and registry-based policy information for HKEY\_CURRENT\_USER (HKCU) in the Registry.pol file

**CAUTION** Do not edit these folders and files directly. Instead, use the appropriate features of one of the Group Policy management tools.

## Using the Group Policy Management Console

You can run the GPMC from the Tools menu in Server Manager. At a prompt or in the Apps Search box, type **gpmc.msc** and then press Enter.

As shown in Figure 4-2, 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, press and hold or right-click the Group Policy Management node, and then tap or click Add Forest. In the Add Forest dialog box, type the name of the forest root domain in the Domain text box, and then tap or click OK.



**FIGURE 4-2** Use the GPMC to work with GPOs in sites, forests, and domains.

When you expand the Forest node, you see 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, press and hold or right-click the Domains node and then tap or click Show Domains. In the Show Domains dialog box, select the check boxes for the domains you want to add and then tap or 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, press and hold or right-click the Sites node and then tap or click Show Sites. In the Show Sites dialog box, select the check boxes for the sites you want to add and then tap or 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 OUs are available to work with in one location.

GPOs listed under the domain, site, and OU 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, while GPOs themselves do not.

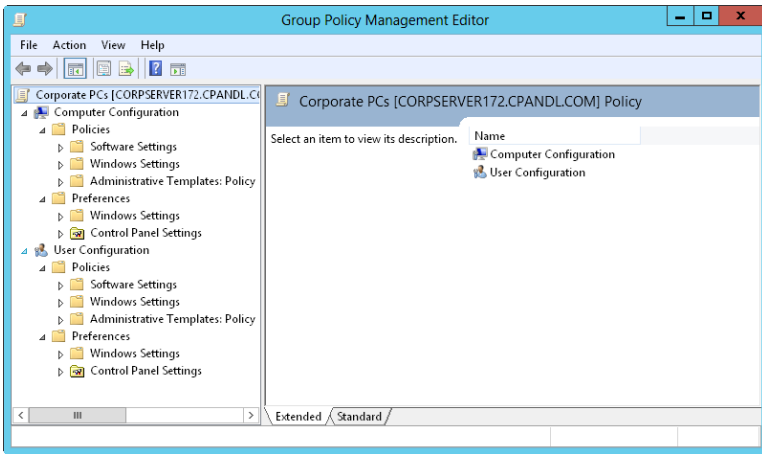
When you start the GPMC, the console connects to Active Directory running on the domain controller that is acting as the PDC emulator for your logon domain and obtains a list of all GPOs and OUs in that domain. It does this by using Lightweight Directory Access Protocol (LDAP) to access the directory store and the Server Message Block (SMB) protocol to access the SYSVOL directory. If the PDC emulator isn't available for some reason, 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. To change the domain controller you are connected to, press and hold or right-click the domain node for which you want to set the domain controller focus, and then tap or 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 tap or click OK.

## Getting to Know the Policy Editor

With the GPMC, you can edit a GPO by pressing and holding or right-clicking it and then selecting Edit on the shortcut menu. As Figure 4-3 shows, 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



**FIGURE 4-3** The configuration of the policy editor depends on the type of policy you're creating and the add-ons installed.

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.

**NOTE** When I reference settings under the Policies node, I'll sometimes use a shortcut such as User Configuration\Administrative Templates\Windows Components rather than User Configuration\Policies\Administrative Templates: Policy Definitions\Windows Components. This shortcut tells you that the policy setting being discussed is under User Configuration rather than Computer Configuration and can be found under Administrative Templates\Windows Components.

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.

**NOTE** A complete discussion of all the available options is beyond the scope of this book. The sections that follow focus on using folder redirection and administrative templates. Scripts are discussed in “User and Computer Script Management” later in this chapter. Security is covered in chapters later in this book.

## 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 add or remove administrative templates as well. 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-tap or 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 tap or 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

When you work with a policy object, creating an 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 OU. Then, as appropriate, you can link the GPO to a specific domain, site, or OU. You can also create a GPO and link it automatically to a domain, site, or OU. The technique you choose 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 OU, the GPO is applied to the user and computer objects in that site, domain, or OU according to the Active Directory options governing inheritance, the precedence order of GPOs, and other settings.

You can create and then link a GPO to a site, domain, or OU by following these steps:

1. In the GPMC, expand the entry for the forest you want to work with, and then expand the related Domains node by double-tapping or double-clicking each node in turn.
2. Press and hold or right-click Group Policy Objects, and then tap or click New. In the New GPO dialog box, type 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 tap or click OK, the new GPO is added to the Group Policy Objects container.

3. Press and hold or right-click the new GPO, and then tap or click Edit. In the policy editor, configure the necessary policy settings, and then close the policy editor.
4. In the GPMC, select the site, domain, or OU. Expand the Sites node you want to work with. In the right pane, the Linked Group Policy Objects tab shows the GPOs that are currently linked to the selected container (if any).
5. Press and hold or right-click the site, domain, or OU to which you want to link the GPO, and then tap or click Link An Existing GPO. In the Select GPO dialog box, select the GPO you want to link with and then tap or click OK. When Group Policy is refreshed for computers and users in the applicable site, domain, or OU, the policy settings in the GPO are applied.

You can create and link a GPO as a single operation by following these steps:

1. In the GPMC, press and hold or right-click the site, domain, or OU for which you want to create and link the GPO, and then tap or click Create A GPO In This Domain, And Link It Here.
2. In the New GPO dialog box, type 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 tap or click OK, the new GPO is added to the Group Policy Objects container and linked to the previously selected site, domain, or OU.
3. Press and hold or right-click the new GPO, and then tap or 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 OU, the policy settings in the GPO are applied.

## Creating and 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, expand the entry for the forest you want to work with, and then double-tap or double-click the related Domains node to expand it.
2. Press and hold or right-click Starter GPOs, and then tap or click New. In the New Starter GPO dialog box, type a descriptive name for the GPO, such as **General Management User GPO**. You can also enter comments describing the GPO's purpose. Tap or click OK.
3. Press and hold or right-click the new GPO, and then tap or click Edit. In the policy editor, configure the necessary policy settings, and then close the policy editor.

# Delegating Privileges for Group Policy Management

In Active Directory, 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)

In Active Directory, 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 tapping or clicking the Delegation tab. On the Delegation tab, you'll see a list of groups and users that can create GPOs in the domain. To grant GPO creation permission to a user or group, tap or click Add. In the Select User, Computer, Or Group dialog box, select the user or group and then tap or click OK.

In the GPMC, you have several ways to determine who has access permissions for Group Policy management. For domain, site, and OU permissions, select the domain, site, or OU you want to work with, and then tap or click the Delegation tab in the right pane, as shown in Figure 4-4. 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 OU
- **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

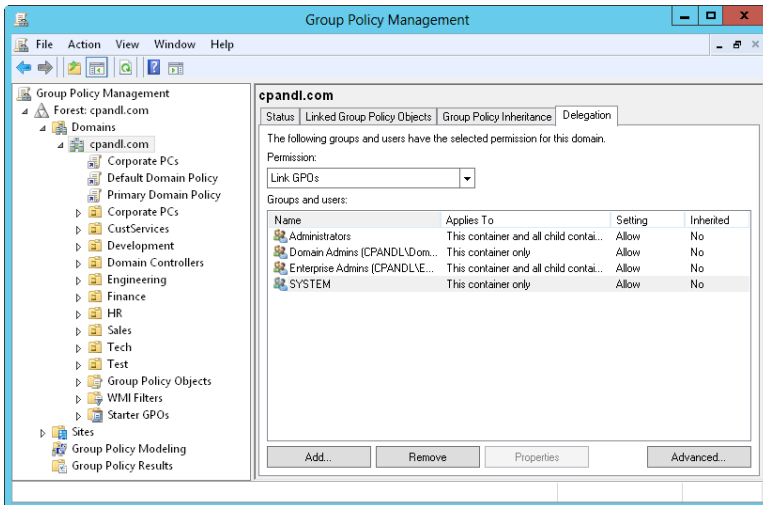


FIGURE 4-4 Review permissions for Group Policy management.



To grant domain, site, or OU permissions, complete the following steps:

1. In the GPMC, select the domain, site, or OU you want to work with and then tap or 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. Tap or click Add. In the Select User, Computer, Or Group dialog box, select the user or group and then tap or 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. Tap or click OK.

For individual GPO permissions, select the GPO you want to work with in the GPMC, and then tap or click the Delegation tab in the right pane. You then see one or more of the following permissions 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 and then tap or click the Delegation tab in the right pane. Tap or click Add.
2. To grant GPO creation permission to a user or group, tap or click Add. In the Select User, Computer, Or Group dialog box, select the user or group and then tap or click OK.
3. In the Add Group Or User dialog box, select the permission level and then tap or click OK.

## Blocking, Overriding, and Disabling Policies

Inheritance ensures that every computer and user object in a domain, site, or OU 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.

You can change the way inheritance works in four key ways:

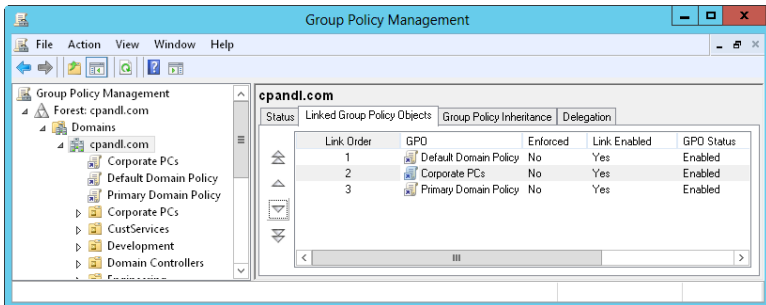
- Change the link order and precedence
- Override inheritance (as long as there is no enforcement)
- Block inheritance (to prevent inheritance completely)
- Enforce inheritance (to supersede and prevent overriding or blocking)

For Group Policy, the order of inheritance goes from the site level to the domain level and then to each nested OU level. Keep the following in mind:

- 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).

When multiple policy objects are linked at a specific level, 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 OU with which you want to work.
2. In the right pane, select the Linked Group Policy Objects tab (as shown in Figure 4-5). Tap or click the policy object you want to work with.



**FIGURE 4-5** Change the link order to modify processing order and precedence.

3. Tap or 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 done 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.

Overriding inheritance is a basic technique for changing 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.

Sometimes you will 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. OU 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 OU, you can ensure that domain or OU administrators have full control over the policies that apply to users and computers under their administration.

Using the GPMC, you can block inheritance by pressing and holding or right-clicking the domain or OU that should not inherit settings from higher-level containers and 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 OU has the Block Inheritance setting enabled.

To prevent administrators who have authority over a container from overriding or blocking inherited Group Policy settings, you can enforce 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 OU 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 OU administrators.

Using the GPMC, you can enforce policy inheritance by expanding the top-level container from which to begin enforcement, pressing and holding or right-clicking the link to the GPO, and then tapping or clicking Enforced. For example, if you want to ensure that a domain-level GPO is inherited by all OUs in the domain, expand the domain container, press and hold or right-click the domain-level GPO, and then tap or 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. Simply select a policy object anywhere in the GPMC, and then view the related Scope tab in the right pane. If the policy is enforced, the Enforced column under Links will display Yes, as shown in Figure 4-6.

After you select a policy object, you can press and hold or 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. Enable or disable enforcement by selecting or clearing the Enforced option.

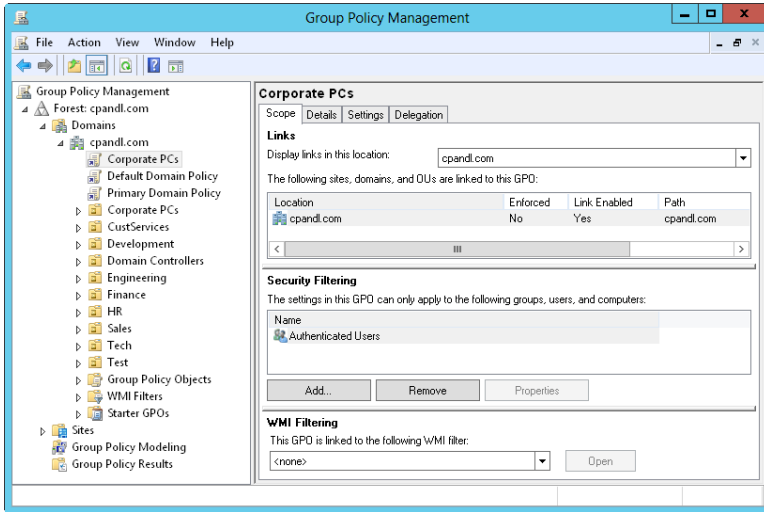


FIGURE 4-6 Enforce policy inheritance to ensure that settings are applied.

## Maintaining and Troubleshooting Group Policy

Group Policy is a broad area of administration that requires careful management. Like any area of administration, Group Policy must also be carefully maintained to ensure proper operation, and you must diagnose and resolve any problems that occur. 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.

### Refreshing Group Policy

When 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. Normally, if there is a conflict between computer and user settings, computer settings have priority and take precedence.

Once 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 non-domain-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.

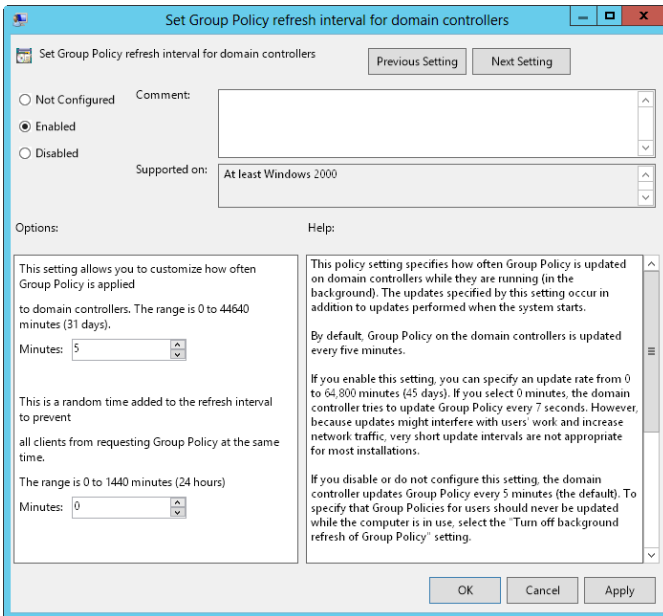
## Configuring the Refresh Interval

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, press and hold or right-click the Group Policy object you want to modify and then tap or 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-tap or double-click the Set Group Policy Refresh Interval For Domain Controllers policy. This displays a Properties dialog box for the policy, shown in Figure 4-7.
3. Define the policy by selecting Enabled. Set the base refresh interval in the first Minutes box. You usually want this value to be between 5 and 59 minutes.
4. 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. Tap or 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.



**FIGURE 4-7** Configure the refresh interval for Group Policy.

To set the refresh interval for member servers and workstations, follow these steps:

1. In the GPMC, press and hold or right-click the Group Policy Object you want to modify and then tap or 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-tap or double-click the Set Group Policy Refresh Interval For Computers policy. This displays a dialog box similar to the one in Figure 4-7.
3. Define the policy by selecting Enabled. In the first Minutes box, set the base refresh interval. You usually want this value to be between 60 and 240 minutes.
4. 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. Tap or click OK.

**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 might want to increase the policy refresh interval as well. 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.

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. Typing **gpupdate** at a prompt or in the Apps 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 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, type **gpupdate /target:computer** at the command prompt. To refresh only user configuration settings, type **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 Purposes

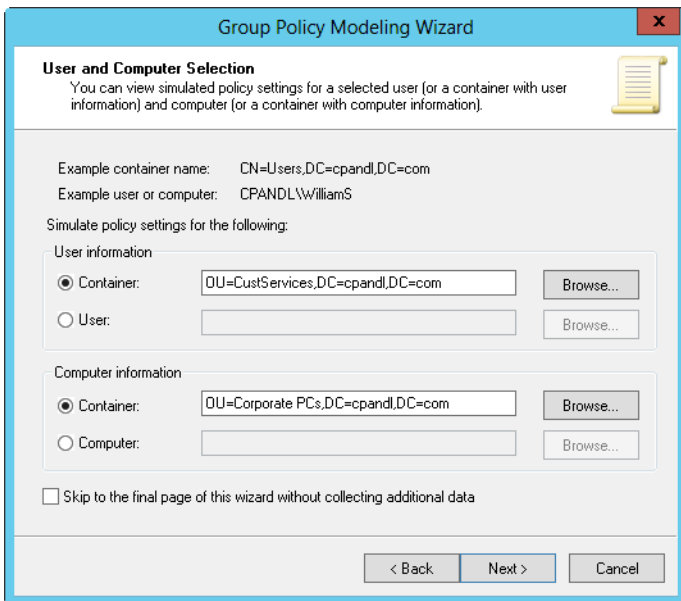
Modeling Group Policy for planning is useful when you want to test various implementation and configuration scenarios. For example, you might want to model the effect of loopback processing or slow-link detection. You can also model the effect of moving users or computers to another container in Active 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, press and hold or right-click the Group Policy Modeling node, select Group Policy Modeling Wizard, and then tap or 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 tap or click the domain controller to use. Tap or click Next.

3. On the User And Computer Selection page, shown in Figure 4-8, you have the option of simulating policy based on containers or individual accounts. Use one of the following techniques to choose accounts, and then tap or click Next:
  - Use containers to simulate changes for entire organizational units or other containers. Under User Information, select Container and then tap or 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, tap or 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, tap or click Browse to display the Select User dialog box, and then specify a user account. Under Computer Information, select Computer, tap or click Browse to display the Select Computer dialog box, and then specify a computer account.

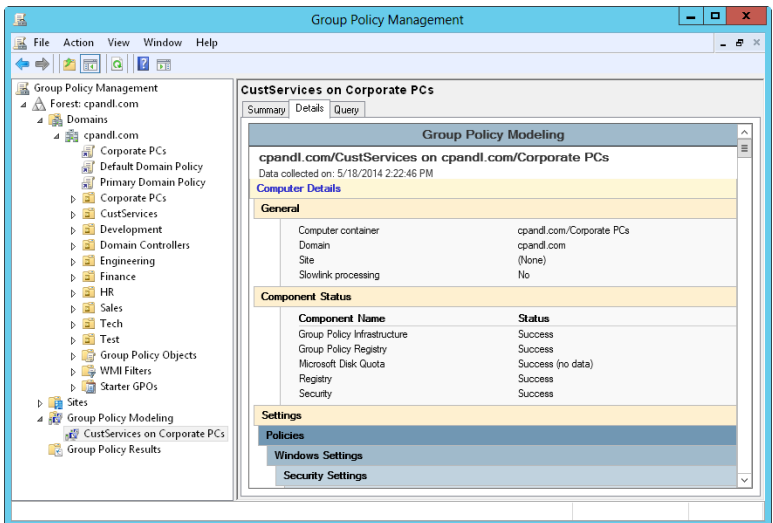


**FIGURE 4-8** Select containers or accounts to use in the simulation.

4. On the Advanced Simulation Options page, select any advanced options for Slow Network Connections, Loopback Processing, and Site as necessary, and then tap or click Next.
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, if you want to see 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. Tap or click Next.
- 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, if you want to see 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. Tap or click Next.
  - 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. Tap or click Next twice to accept the default options.
  - Review the selections you made, and then tap or click Next. After the wizard gathers policy information, tap or 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.
  - When you select the Details tab in the right pane as shown in Figure 4-9, 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.



**FIGURE 4-9** Review the report to determine the effects of modeling.

## Copying, Pasting, and Importing Policy Objects

The GPMC features built-in copy, paste, and import operations. Using the copy and paste features is fairly straightforward. The Copy and Paste options are available when you press and hold or right-click a GPO in the GPMC. You can copy a policy object and all its settings in one domain and then navigate to the domain into which you want to paste the copy of the policy object. 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, press and hold or right-click Group Policy Objects and then select New. In the New GPO dialog box, type a descriptive name for the new GPO and then tap or click OK.
2. The new GPO is now listed in the Group Policy Objects container. Press and hold or right-click the new policy object, and then tap or click Import Settings. This starts the Import Settings Wizard.
3. Tap or 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, tap or 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 tap or click OK. Tap or click Next to continue.
5. If multiple backups are stored in the designated backup folder, you'll see a list of them on the Source GPO page. Tap or click the one you want to use, and then tap or click Next.
6. The Import Settings Wizard scans the policy object for references to security principals and 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 tapping or clicking Next, and then tap or click Finish to begin the import process. When importing is complete, tap or click OK.

## 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, press and hold or right-click the Group Policy Objects node, and then tap or click Back Up All. If you want to back up a specific policy object in the domain, press and hold or right-click the policy object and select Back Up.
2. In the Back Up Group Policy Object dialog box, tap or click Browse. In the Browse For Folder dialog box, set the location where the GPO backup should be stored.
3. In the Description text box, type a description of the contents of the backup. Tap or click Back Up to start the backup process.
4. The Backup dialog box shows the progress and status of the backup. Tap or 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 once. 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, press and hold or right-click the Group Policy Objects node and then tap or click Manage Backups. This displays the Manage Backups dialog box.
2. In the Backup Location text box, tap or click Browse. In the Browse For Folder dialog box, find the backup folder and then tap or 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, tap or click View Settings, and then use Internet Explorer to verify that the settings are as expected. When you are ready to continue, tap or click Restore. Confirm that you want to restore the selected policy object by tapping or 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.

## Determining Current Group Policy 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, 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 pressing and holding or right-clicking the Group Policy Results node and selecting Group Policy Results Wizard. When the Group Policy Results Wizard starts, follow the prompts.

## Disabling an Unused Part of Group Policy

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 OU with which you want to work.
2. Select the policy object you want to work with, and then tap or click the Details tab in the right pane.
3. Choose one of the following status settings from the GPO Status list, and then tap or 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 Policy 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, press and hold or right-click the Group Policy object you want to modify and then tap or click Edit.
2. In the Administrative Templates for Computer Configuration under System\Group Policy, double-tap or 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 tapping or 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.

## Configuring Slow-Link Detection

Slow-link detection is used by Group Policy clients to detect 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. Once 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, which is stored in the Administrative Templates for Computer Configuration under System\Group Policy. 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. On the other hand, 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 using the Enable File Synchronization On Costed Networks policy found under Computer Configuration\Administrative Templates\Network\Offline Files.
- Control background transfers on costed networks 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 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 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

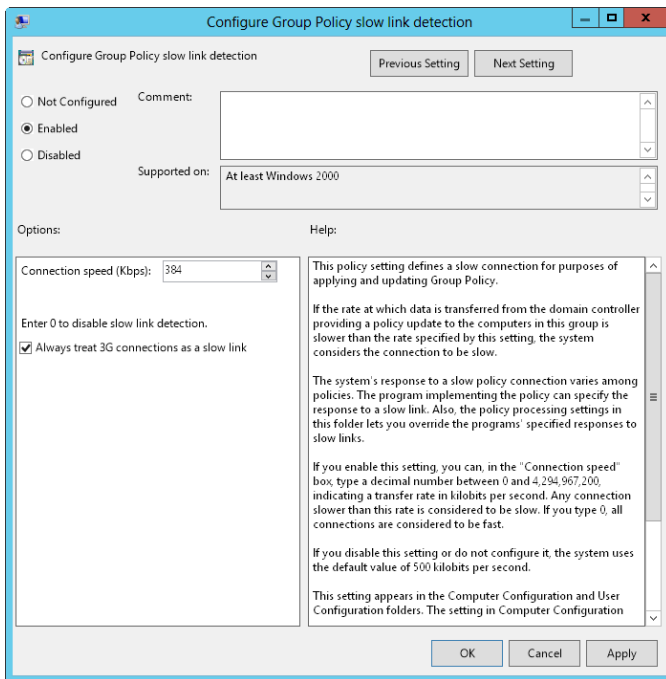
- Disk Quota Policy Processing
- EFS Recovery Policy Processing
- Folder Redirection Policy Processing
- Scripts Policy Processing
- Software Installation Policy Processing

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 if applications are failing during refresh operations.

You can configure slow-link detection and related policy processing by following these steps:

1. In the GPMC, press and hold or right-click the policy object you want to modify and then tap or click Edit.
2. In the Administrative Templates for Computer Configuration under System\Group Policy, double-tap or double-click the Configure Group Policy Slow Link Detection policy.
3. Select Enabled to define the policy, as shown in Figure 4-10. 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 3G connections should always be treated as slow links. Tap or click OK.



**FIGURE 4-10** Configure slow-link detection.

To configure slow-link and background policy processing of key areas of Group Policy, follow these steps:

1. In the GPMC, press and hold or right-click the policy object you want to modify and then tap or click Edit.
2. Expand Computer Configuration\Administrative Templates\System\Group Policy.

3. Double-tap or double-click the processing policy you want to configure. Select Enabled to define the policy, as shown in Figure 4-11, 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.

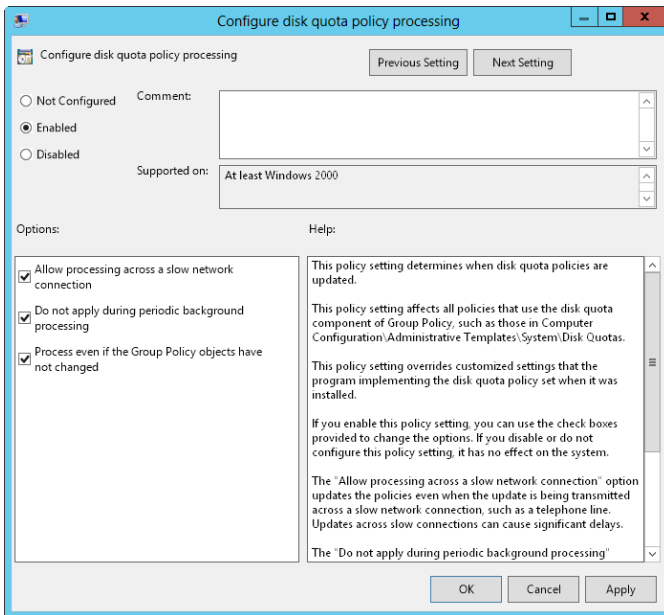


FIGURE 4-11 Configure policy processing for slow links.

4. Tap or click OK to save your settings.

## 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 OU from using the related policy settings but does not delete the GPO. Because of this, the GPO remains linked to other sites, domains, or OUs as appropriate. In the GPMC, you can remove a link to a GPO by pressing and holding or right-clicking the GPO link in the container that it is linked to and then selecting Delete. When prompted to confirm that you want



to remove the link, tap or click OK. If you remove all links to the GPO from sites, domains, and OUs, 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 OUs. 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. Press and hold or right-click the GPO, and then select Delete. When prompted to confirm that you want to remove the GPO and all links to it, tap or click Yes.

## Troubleshooting Group Policy

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, press and hold or right-click the Group Policy Results node and then tap or click Group Policy Results Wizard. When the wizard starts, tap or 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 tap or click Browse. In the Select Computer dialog box, type the name of the computer and then tap or click Check Names. After you select the correct computer account, tap or click OK and then tap or 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. Tap or click Next.
4. Review the selections you made, and then tap or click Next. After the wizard gathers policy information, tap or 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 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 view RSoP as well. Gpresult provides details on 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, as well as 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 Cpandl domain, you would type the following command:

```
gpresult /s corppc85 /user cpandl\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 using the */h* parameter or an XML report using the */x* parameter. The following examples use these parameters:

```
gpresult /s corppc85 /user cpandl\tedg /h gpreport.html  
gpresult /s corppc85 /user cpandl\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 be lost as well. 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 type **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, type **dcgpofix /target:domain**. If you want to fix only the Default Domain Controller Policy, type **dcgpofix /target:dc**.

## Managing Users and Computers with Group Policy

---

You can use Group Policy to manage users and computers in many different ways. In the sections that follow, I'll describe some specific management areas, including the following:

- Folder redirection
- Computer and user scripts
- Software deployment
- Computer and user certificate enrollment
- Automatic update settings

### 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 Windows XP Professional and earlier releases of Windows, the special folders you can centrally manage are Application Data, Start Menu, Desktop, My Documents, and My Pictures. For Windows Vista and later 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 Windows Vista and later store personal folders in slightly different ways, you manage the folders in the same way within Group Policy.

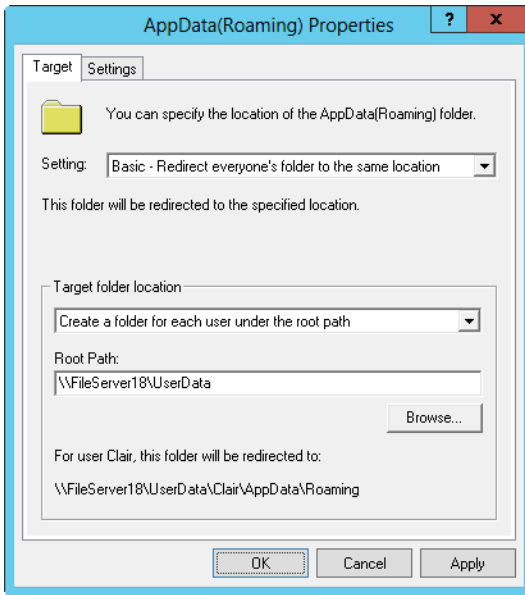
You have two general options for redirection. You can redirect a special folder to the same network location for all users, or you can designate locations based on user membership in security groups. In either case, you should make sure that the network location you plan to use is available as a network share. See Chapter 12, "Data Sharing, Security, and Auditing," 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 8 and Windows Server 2012 allow 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 "Local, Roaming, and Mandatory Profiles" in Chapter 9, "Managing User and Group Accounts."

## 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 GPMC, press and hold or right-click the GPO for the site, domain, or organizational unit you want to work with, and then tap or click Edit. This opens the policy editor for the GPO.
2. In the policy editor, expand the following nodes: User Configuration, Windows Settings, and Folder Redirection.
3. Under Folder Redirection, press and hold or right-click the special folder you want to work with, such as AppData(Roaming), and then tap or click Properties. This opens a Properties dialog box similar to the one shown in Figure 4-12.



**FIGURE 4-12** Set options for redirection using a special folder's Properties dialog box.

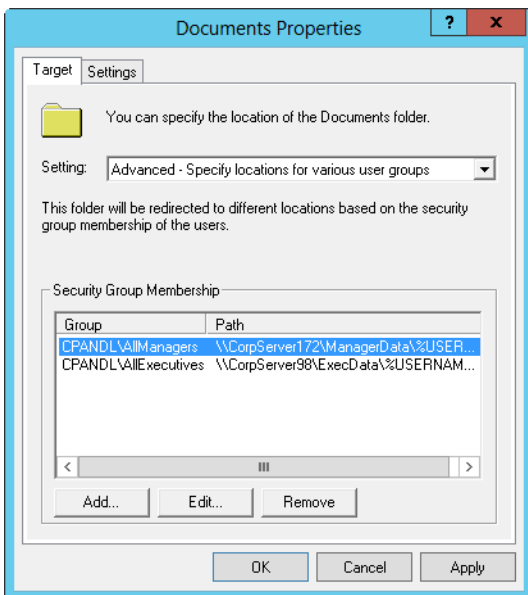
4. In the Setting list on the Target tab, choose Basic—Redirect Everyone's Folder To The Same Location.
5. Under Target Folder Location, you have several options. The options available depend on the folder you're working with and include the following:
  - **Redirect To The User's Home Directory** If 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.
6. Tap or click the Settings tab, configure the following additional options, and then tap or 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 FolderName To 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, press and hold or right-click the GPO for the site, domain, or organizational unit you want to work with, and then tap or click Edit. This opens the policy editor for the GPO.
2. In the policy editor, expand the following nodes: User Configuration, Windows Settings, and Folder Redirection.
3. Under Folder Redirection, press and hold or right-click the special folder you want to work with, such as AppData(Roaming), and then tap or click Properties.
4. On the Target tab, choose Advanced—Specify Locations For Various User Groups in the Setting list. As shown in Figure 4-13, a Security Group Membership panel is added to the Properties dialog box.



**FIGURE 4-13** Configure advanced redirection using the Security Group Membership panel.

5. Tap or click Add to open the Specify Group And Location dialog box. Or select a group entry, and then tap or click Edit to modify its settings.
6. In the Security Group Membership text box, type the name of the security group for which you want to configure redirection, or tap or click Browse to find a security group to add.
7. As with basic redirection, the options available depend on the folder you're working with 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. Tap or click OK. Repeat steps 5–7 for other groups you want to configure.
- 9. When you're done creating group entries, tap or click the Settings tab, configure the following additional options, and then tap or 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 FolderName To 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, press and hold or right-click the GPO for the site, domain, or organizational unit you want to work with. Then tap or click Edit to open the policy editor for the GPO.
2. In the policy editor, expand the following nodes: User Configuration, Windows Settings, and Folder Redirection.
3. Under Folder Redirection, press and hold or right-click the special folder you want to work with and then tap or click Properties.
4. Tap or click the Settings tab, and then make sure that an appropriate Policy Removal option is selected. 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. The contents aren't deleted from the previous location, however.
5. If you changed the Policy Removal option, tap or click Apply and then tap or click the Target tab. Otherwise, just tap or 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 tap or click Remove. Tap or click OK.

# User and Computer Script Management

With Windows Server you can configure four types of scripts:

- **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 2000 and later releases support 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 lets you 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.

Windows 7 and Windows 8, as well as Windows Server 2008 R2 and Windows Server 2012 also support Windows PowerShell scripts. If you installed Windows PowerShell 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.

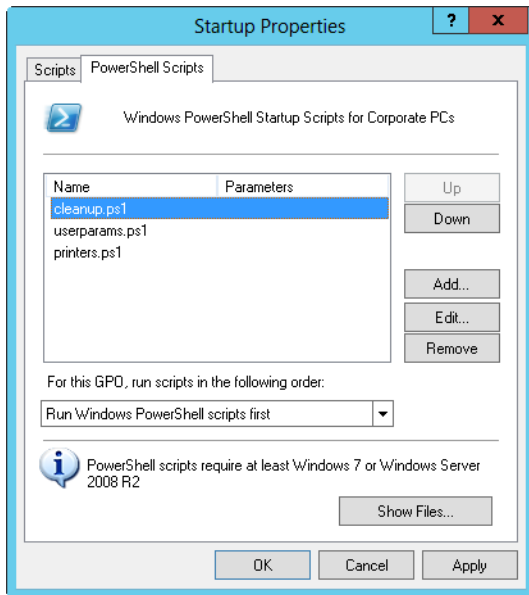
## Assigning Computer Startup and Shutdown Scripts

Computer startup and shutdown scripts are assigned as part of a GPO. In this way, all computers that are members of the site, domain, or organizational unit—or all three—execute scripts automatically when they're booted or shut down.

To assign a computer startup or shutdown script, follow these steps:

1. Open the folder containing the script or scripts you want to use in File Explorer.
2. In the GPMC, press and hold or right-click the GPO for the site, domain, or organizational unit you want to work with, and then tap or click Edit. This opens the policy editor for the GPO.
3. In the Computer Configuration node, double-tap or double-click the Windows Settings folder and then tap or click Scripts.
4. To work with startup scripts, press and hold or right-click Startup and then tap or click Properties. To work with shutdown scripts, press and hold or right-click Shutdown and select Properties. This opens a dialog box similar to the one shown in Figure 4-14.





**FIGURE 4-14** Add, edit, and remove computer startup scripts using the Startup Properties dialog box.

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, tap or 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. Tap or click Add to assign a script. This opens the Add A Script dialog box. In the Script Name text box, type 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 tap or click Edit. To delete a script, select the script in the Script For list and tap or click Remove.
10. To save your changes, tap or click OK.


## Assigning User Logon and Logoff Scripts

You can assign user scripts in one of three ways:

- You can assign logon and logoff scripts as part of a GPO. In this way, 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. In this way, you can assign each user or group a separate logon script. For details, see “Configuring the User’s Environment Settings” in Chapter 9.
- You can also assign individual logon scripts as scheduled tasks. You schedule tasks using the Scheduled Task Wizard.

To assign a logon or logoff script in a GPO, follow these steps:

1. Open the folder containing the script or scripts you want to use in File Explorer.
2. In the GPMC, press and hold or right-click the GPO for the site, domain, or organizational unit you want to work with, and then tap or click Edit. This opens the policy editor for the GPO.
3. Double-tap or double-click the Windows Settings folder in the User Configuration node, and then tap or click Scripts.
4. To work with logon scripts, press and hold or right-click Logon and then tap or click Properties. To work with logoff scripts, press and hold or right-click Logoff and then tap or click Properties. This opens a dialog box similar to the one shown in Figure 4-15.



**FIGURE 4-15** Add, edit, and remove user logon scripts using the Logon Properties dialog box.

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, tap or 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. Tap or click Add to assign a script. This opens the Add A Script dialog box. In the Script Name text box, type 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 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 tap or click Edit. To delete a script, select the script in the Script For list and then tap or click Remove.
10. To save your changes, tap or 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 Systems Management Server (SMS), you can use it to automate the deployment and maintenance of software in just about any size organization, provided that your computers are running business editions of Windows 2000 or later.

### Getting to Know Software Installation Policy

In Group Policy, you can deploy software on a per-computer or per-user basis. Per-computer applications are available to all users of a computer and configured under Computer Configuration\Software Settings\Software Installation. Per-user applications are available to individual users and configured under User Configuration\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 is 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, such as Microsoft Office, you prepare the distribution point by performing an administrative installation to the share. With Microsoft Office, you can do this by running the application's Setup program with the */a* parameter and designating the share as the install location. 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) files. When you use computer assignment, user assignment, or user publishing, you can deploy software using Windows Installer Packages. When you use user publishing, you can deploy software 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.

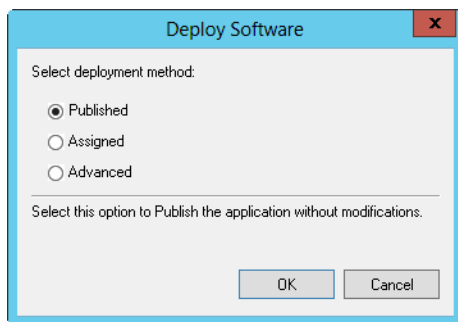
Because Software Installation policy is applied only during foreground processing of policy settings, per-computer application deployments are processed at startup and per-user application deployments are processed at logon. You can customize installation using transform (.mst) files. Transform files 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, press and hold or right-click the GPO you want to use for the deployment and then tap or click Edit.
2. In the policy editor, open Computer Configuration\Software Settings\Software Installation or User Configuration\Software Settings\Software Installation as appropriate for the type of software deployment.
3. Press and hold or right-click Software Installation. On the shortcut menu, tap or click New, and then tap or click Package.
4. In the Open dialog box, navigate to the network share where your package is located, tap or click the package to select it, and then tap or 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, shown in Figure 4-16, select one of the following deployment methods and then tap or click OK:
  - **Published** To publish the application without modifications
  - **Assigned** To assign the application without modifications
  - **Advanced** To deploy the application using advanced configuration options



**FIGURE 4-16** Select the deployment method.

## Configuring Software Deployment Options

You can view and set general options for a software package by following these steps:

1. In the GPMC, press and hold or right-click the GPO you want to use for the deployment and then tap or click Edit.
2. In the policy editor, access Computer Configuration\Software Settings\Software Installation or User Configuration\Software Settings\Software Installation as appropriate for the type of software deployment.
3. Double-tap or double-click the Software Installation package. In the Properties dialog box, review or modify software deployment options.

4. On the Deployment tab, shown in Figure 4-17, 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 sees all setup screens and messages during installation. With the Basic option, the user sees only error and completion messages during installation.

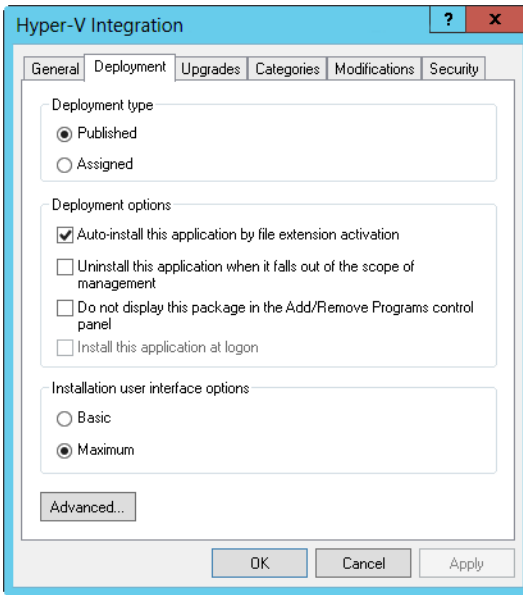


FIGURE 4-17 Review and modify the deployment options as necessary.

5. Tap or 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 (patch) 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, press and hold or right-click the GPO you want to use for the deployment and then tap or click Edit.
3. In the policy editor, access Computer Configuration\Software Settings\Software Installation or User Configuration\Software Settings\Software Installation as appropriate for the type of software deployment.
4. Press and hold or right-click the package you want to work with. On the shortcut menu, tap or click All Tasks, and then tap or click Redeploy Application.
5. When prompted to confirm the action, tap or click Yes. The application is then redeployed to all users and computers as appropriate for the GPO you are working with.

When an application uses a non-Windows Installer package, you can update a deployed application or apply a service pack by following these steps:

1. In the GPMC, press and hold or right-click the GPO you want to use for the deployment and then tap or click Edit.
2. In the policy editor, access Computer Configuration\Software Settings\Software Installation or User Configuration\Software Settings\Software Installation as appropriate for the type of software deployment.
3. Press and hold or right-click the package. On the shortcut menu, tap or click All Tasks, and then tap or click Remove. Tap or 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, press and hold or right-click the GPO you want to use for the deployment and then tap or click Edit.
3. In the policy editor, access Computer Configuration\Software Settings\Software Installation or User Configuration\Software Settings\Software Installation as appropriate for the type of software deployment.

4. Press and hold or right-click Software Installation. On the shortcut menu, tap or click New, and then tap or click Package. Create an assigned or published application by using the Windows Installer file for the new software version.
5. Press and hold or right-click the upgrade package, and then tap or click Properties. On the Upgrades tab, tap or 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, tap or 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. Tap or 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 tap or click OK to close the upgrade package's Properties dialog box.

## 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 autoenrollment. 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 autoenrollment works. When you install enterprise CAs, autoenrollment policies for users and computers are enabled automatically. The policy for computer certificate enrollment is Certificate Services Client—Auto-Enrollment Settings under Computer Configuration\Windows Settings\Security Settings\Public Key Policies. The policy for user certificate enrollment is Certificate Services Client—Auto-Enrollment under User Configuration\Windows Settings\Security Settings\Public Key Policies.

You can configure autoenrollment by following these steps:

1. In the GPMC, press and hold or right-click the GPO you want to work with and then tap or click Edit.
2. In the policy editor, access User Configuration\Windows Settings\Security Settings\Public Key Policies or Computer Configuration\Windows Settings\



Security Settings\Public Key Policies as appropriate for the type of policy you want to review.

3. Double-tap or double-click Certificate Services Client—Auto-Enrollment. To disable automatic enrollment, select Disabled from the Configuration Model list, tap or 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. Tap or click OK to save your 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 8 and Windows Server 2012 use Windows Update to download Windows Components as well as 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. For more information, see "Server Manager Essentials and Binaries" in Chapter 2.

### 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 then 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 later.
- **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 later.

- **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, press and hold or right-click the GPO you want to work with and then tap or click Edit.
2. In the policy editor, access Computer Configuration\Administrative Templates\Windows Components\Windows Update.
3. Double-tap or 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, tap or click OK, and then skip the remaining steps.
4. Choose an update configuration from the options in the Configure Automatic Updating list.
5. If you select Auto Download And Schedule The Install, you can schedule the installation day and time by using the lists provided. Tap or click OK to save your settings.

## Optimizing Automatic Updates

Generally, most automatic updates are installed only when a computer is shut down and restarted. Some automatic updates can be installed immediately without interrupting system services or requiring system restart. To ensure that some updates can be installed immediately, follow these steps:

1. In the GPMC, press and hold or right-click the GPO you want to work with and then tap or click Edit.
2. In the policy editor, access Computer Configuration\Administrative Templates\Windows Components\Windows Update.
3. Double-tap or double-click Allow Automatic Updates Immediate Installation. In the Properties dialog box, select Enabled and then tap or click OK.

By default, only users with local administrator privileges receive notifications about updates. You can allow any user logged on to a computer to receive update notifications by following these steps:

1. In the GPMC, press and hold or right-click the GPO you want to work with and then tap or click Edit.
2. In the policy editor, access Computer Configuration\Administrative Templates\Windows Components\Windows Update.
3. Double-tap or double-click Allow Non-Administrators To Receive Update Notifications. In the Properties dialog box, select Enabled and then tap or click OK.

Another useful policy is Remove Access To Use All Windows Update Features. This policy prohibits access to all Windows Update features. If enabled, all Automatic Updates features are removed and can't be configured. This includes the Automatic Updates tab in the System utility and driver updates from the Windows Update website in Device Manager. This policy is located in User Configuration\Administrative Templates\Windows Components\Windows Update.

## Using Intranet Update Service Locations

On networks with hundreds or thousands of computers, the Automatic Updates process 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, consider using the Specify Intranet Microsoft Update Service Location policy, which tells individual computers to check a designated internal server for updates.

The designated update server must run Windows Server Update Services (WSUS), be configured as a web server running Microsoft Internet Information Services (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 you want to work with for editing. In the policy editor, access Computer Configuration\Administrative Templates\Windows Components\Windows Update.
2. Double-tap or double-click Specify Intranet Microsoft Update Service Location. In the Properties dialog box, select Enabled.
3. In the Set The Intranet Update Service For Detecting Updates text box, type the URL of the update server. In most cases, this is `http://servername`, such as `http://CorpUpdateServer01`.

4. Type 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. Tap or click OK. After the applicable Group Policy object 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.

# Enhancing Computer Security

- Using Security Templates **191**
- Using the Security Configuration Wizard **206**

Sound security practices and settings are essential to successful system administration. Two key ways to configure security settings are to use security templates and security policies. Both of these features manage system settings that you would typically manage through Group Policy otherwise.

## Using Security Templates

---

Security templates provide a centralized way to manage security-related settings for workstations and servers. You use security templates to apply customized sets of Group Policy definitions to specific computers.

These policy definitions generally affect the following policies:

- **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

**NOTE** Security templates are available in all Microsoft Windows Server installations and can be imported into any Group Policy object. Security templates apply only to the Computer Configuration area of Group Policy. They do not apply to the User Configuration area. In Group Policy, you'll find applicable settings under Computer Configuration\Windows Settings\Security Settings. Some security settings are not included, such as those that apply to wireless networks, public keys, software restrictions, and IP security.

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 you are working with 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, as shown in Figure 5-1.

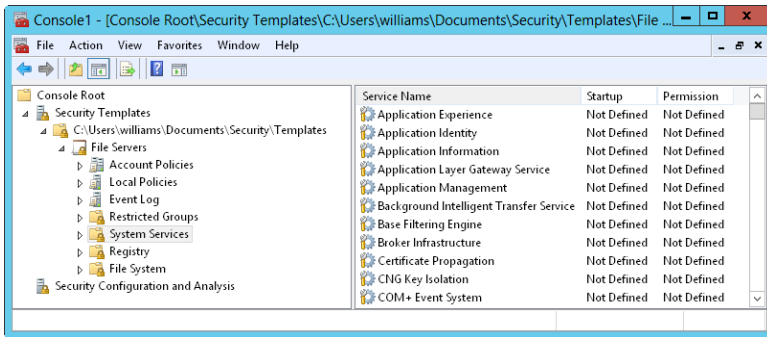


FIGURE 5-1 View and create security templates with the Security Templates snap-in.

**TIP** 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 make sense, you should modify it appropriately or delete it.

You don't use the Security Templates snap-in to apply templates. You use the Security Configuration And Analysis snap-in to apply templates. 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 where the current settings don't match those in the template. This is useful to determine whether security settings have changed over time.

# Using the Security Templates and Security Configuration And Analysis Snap-ins

You can open the security snap-ins by following these steps:

1. Start the Microsoft Management Console (MMC). One way to do this is by pressing the Windows key, typing **mmc.exe**, and then pressing Enter.
2. In the Microsoft Management Console, tap or click File and then tap or click Add/Remove Snap-In.
3. In the Add Or Remove Snap-Ins dialog box, tap or click Security Templates and then tap or click Add.
4. Tap or click Security Configuration And Analysis, and then tap or click Add. Tap or 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. With the Security Templates snap-in selected in the MMC, choose New Template Search Path from the Action menu.
2. In the Browse For Folder dialog box, select the template location to add, such as %SystemRoot%\Security\Templates\Policies. Tap or click OK.

Now that you've located the template search path you want to work with, you can select a template and expand the related notes to review its settings.

You can create a template by following these steps:

1. In the Security Templates snap-in, either press and hold or right-click the search path where the template should be created and then tap or click New Template.
2. Type a name and description for the template in the text boxes provided.
3. Tap or click OK to create the template. The template will have no settings configured, so you need to modify the settings carefully before the template is ready for use.
4. After you modify the template, save the changes by pressing and holding or right-clicking the template in the Security Templates snap-in and selecting Save. Alternatively, you can select Save As to assign a different name to the modified template.

## Reviewing and Changing Template Settings

The sections that follow discuss how to work with template settings. As you'll learn, you manage each type of template setting in a slightly different way.

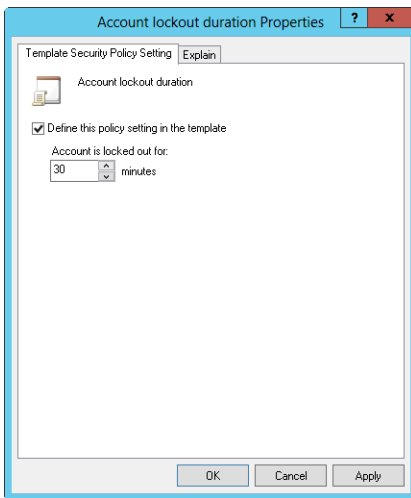
### Changing Settings for Account, Local, and Event Log Policies

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 and local policy settings, see

Chapter 8, "Creating User and Group Accounts." For detailed information on configuring event logging, see Chapter 3, "Monitoring Processes, Services, and Events."

With account, local, and event log policies, you can change template settings by following these steps:

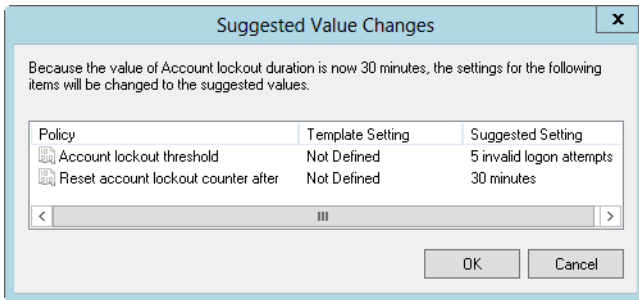
1. In the Security Templates snap-in, expand the Account Policies or Local Policies node as necessary and then select a related subnode, such as Password Policy or Account Lockout Policy.
2. In the right pane, policy settings are listed alphabetically. The value in the Computer Setting column shows the current setting. If the template changes the setting so that it is no longer defined, the value is listed as Not Defined.
3. Double-tap or double-click a setting to display its Properties dialog box, shown in Figure 5-2. To determine the purpose of the setting, tap or 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.



**FIGURE 5-2** Change template settings for account and local policies.

4. If you enable the policy setting, specify how the policy setting is to be used by configuring any additional options.
5. Tap or click OK to save your changes. You might see the Suggested Value Changes dialog box, shown in Figure 5-3. 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, as shown in the figure.



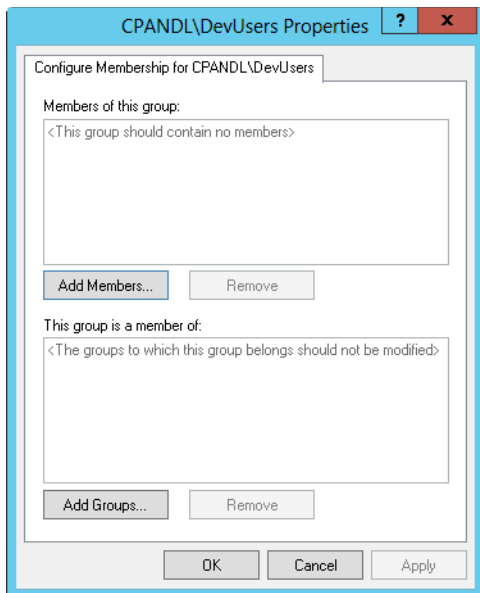


**FIGURE 5-3** Review the suggested value changes.

## Configuring Restricted Groups

Restricted groups policy settings control the list of members of groups as well as 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. 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.
2. You can add a restricted group by pressing and holding or right-clicking the Restricted Groups node in the left pane and then tapping or clicking Add Group. In the Add Group dialog box, tap or click Browse.
3. In the Select Groups dialog box, type the name of a group you want to restrict and then tap or click Check Names. If multiple matches are found, select the account you want to use and then tap or click OK. If no matches are found, update the name you entered and try searching again. Repeat this step as necessary, and then tap or click OK.
4. In the Properties dialog box, shown in Figure 5-4, you can use the Add Members option to add members to the group. Tap or click Add Members, and then specify the members of the group. If the group should not have any members, remove all members by tapping or clicking Remove. Any members who are not specified in the policy setting for the restricted group are removed when the security template is applied.
5. In the Properties dialog box, tap or 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 (provided that 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.
6. Tap or click OK to save your settings.



**FIGURE 5-4** Review the suggested value changes.

You can remove a restriction on a group by following these steps:

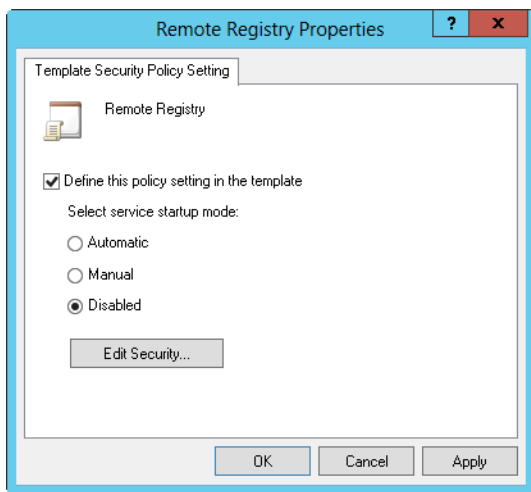
1. In the Security Templates snap-in, select the Restricted Groups node. In the right pane, any currently restricted groups are listed by name. Members of the group are listed along with the groups of which the restricted group is a member.
2. Press and hold or right-click the group that should not be restricted, and then tap or click Delete. When prompted to confirm the action, tap or click Yes.

## Enabling, Disabling, and Configuring 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. In the right pane, all currently installed services on the computer you are working with 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.
2. Double-tap or double-click the entry for a system service to display its Properties dialog box, shown in Figure 5-5. 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.



**FIGURE 5-5** Change template settings for system services.

3. 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 that the template is being applied to.
  - 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.
4. If you know the security configuration that the service should use, tap or 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.
5. Tap or click OK.

## Configuring Security Settings for 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 you want work with. In the right pane, all currently secured paths are listed.
2. Double-tap or double-click a registry or file path to view its current settings, as shown in Figure 5-6.

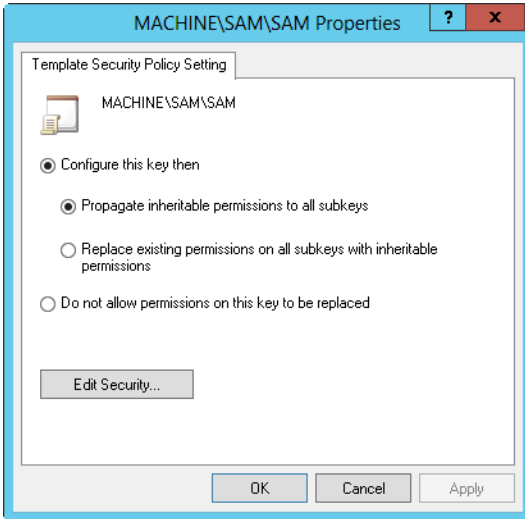


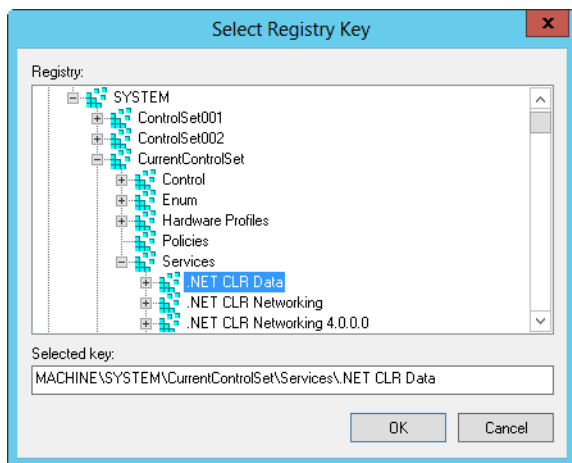
FIGURE 5-6 Change template settings for paths and keys.

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 tap or 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. Tap or 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. See Chapter 12, “Data Sharing, Security, and Auditing,” for details on permissions, auditing, and ownership.
6. Tap or 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 press and hold or right-click the Registry node and then tap or click Add Key. This displays the Select Registry Key dialog box, shown in Figure 5-7.



**FIGURE 5-7** Select the registry path or value to secure.

2. In the Select Registry Key dialog box, select the registry path or value you want to work with and then tap or 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. See Chapter 12 for details on permissions, auditing, and ownership.
4. Tap or 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 tap or 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

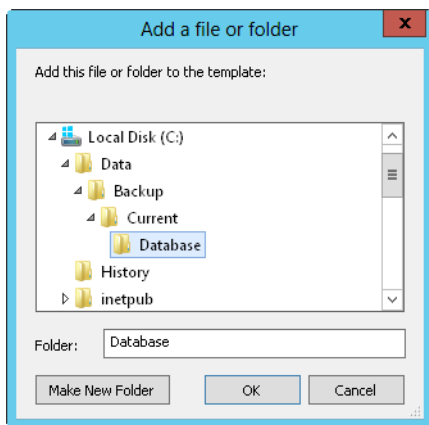
path. Existing permissions are replaced only if they conflict with a security permission set for this path.

- 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.** Tap or click **OK**.

You can define security settings for file paths by following these steps:

- 1.** In the **Security Templates** snap-in, select and then press and hold or right-click the **File System** node, and then tap or click **Add File**. This displays the **Add A File Or Folder** dialog box, shown in Figure 5-8.



**FIGURE 5-8** Select the file or folder path to secure.

- 2.** In the **Add A File Or Folder** dialog box, select the file or folder path or value you want to work with and then tap or 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. See Chapter 12 for details on permissions, auditing, and ownership.
- 4.** Tap or 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 tap or 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. Tap or click OK.

## Analyzing, Reviewing, and Applying Security Templates

As stated previously, you use the Security Configuration And Analysis snap-in 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.

**REAL WORLD** The key drawback to using the Security Configuration And Analysis snap-in is that you cannot configure multiple computers at once. 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, 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) and then deploy the security configuration to multiple computers. For more information, see "Deploying Security Templates to Multiple Computers" later in this chapter.

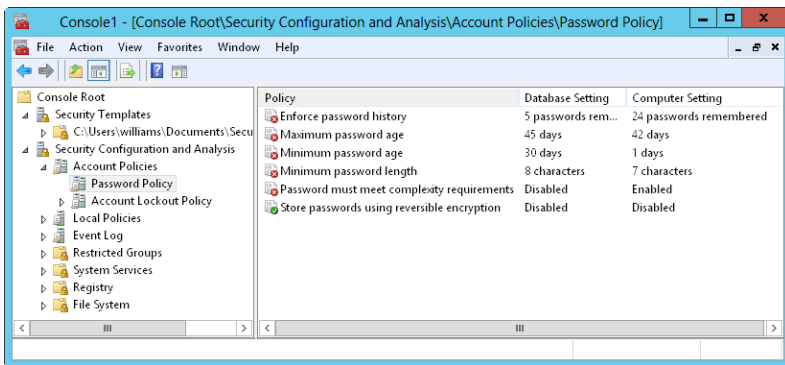
The Security Configuration And Analysis snap-in uses a working database to store template security settings and then applies the settings from this database. For analysis and comparisons, the template settings are listed as the effective database settings and the current computer settings are listed 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. Open the Security Configuration And Analysis snap-in.
2. Press and hold or right-click the Security Configuration And Analysis node, and then tap or click Open Database. This displays the Open Database dialog box.
3. By default, the Open Database dialog box's search path is set to %System-Drive%\Users\%UserName%\Documents\Security\Database. As necessary, select options in the Open Database dialog box to navigate to a new save location. In the File Name text box, type a descriptive name for the database, such as **Current Config Comparison**, and then tap or click Open. The security database is created in the Security Database Files format with the .sdb file extension.

4. The Import Template dialog box is displayed with the default search path set to %SystemDrive%\Users\%UserName%\Documents\Security\Templates. As necessary, select options in the Import Template dialog box to navigate to a new template location. Select the security template you want to use, and then tap or click Open. Security template files end with the .inf file extension.
5. Press and hold or right-click the Security Configuration And Analysis node, and then tap or click Analyze Computer Now. When prompted to set the error log path, type a new path or tap or click OK to use the default path.
6. 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 pressing and holding or right-clicking the Security Configuration And Analysis node and choosing View Log File.

When you are working with the Security Configuration And Analysis snap-in, you can review the differences between the template settings and the current computer settings. As Figure 5-9 shows, 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.

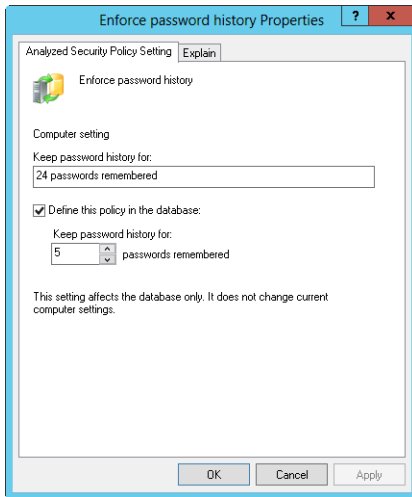


**FIGURE 5-9** Review the differences between the template settings and the current computer settings.

You can make changes to a setting stored in the database by following these steps:

1. In the Security Configuration And Analysis snap-in, double-tap or double-click the setting you want to work with.
2. In the Properties dialog box, shown in Figure 5-10, note the current computer setting. If information about the purpose of the setting is available, you can view this by tapping or clicking the Explain tab.





**FIGURE 5-10** Change a policy setting in the database before applying the template.

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, press and hold or right-click the Security Configuration And Analysis node and then tap or 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. Type the following:

```
secedit /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, press and hold or right-click the Security Configuration And Analysis node, and then tap or click Configure Computer Now. When prompted to set the error log path, tap or click OK because the default path should be sufficient. To view the configuration error log, press and hold or right-click the Security Configuration And Analysis node and then tap or 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 their previous state. To apply a rollback template, follow these steps:

1. In the Security Configuration And Analysis snap-in, press and hold or right-click the Security Configuration And Analysis node and then tap or 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 tap or click Open.
4. Press and hold or right-click the Security Configuration And Analysis node, and then tap or click Configure Computer Now. Tap or click OK.

The only settings that cannot be restored are for access control lists on file system and registry paths. Once 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.

## Deploying Security Templates to 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 that the template settings should apply to. Then, when policy is refreshed, all computers within the scope of the GPO receive the security configuration.

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.
- **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** You need to be extra 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 is extremely important to create and link a new GPO rather than use an existing GPO.

To deploy a security template to a computer GPO, follow 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. Press and hold or right-click Security Settings, and then tap or click Import Policy.
3. In the Import Policy From dialog box, select the security template to import and then tap or click Open. Security templates end with the .inf file extension.
4. Check the configuration state of the security settings to verify that the settings were imported as expected, and then close the policy editor. Repeat this 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.

# Using the Security Configuration Wizard

---

The Security Configuration Wizard can help you create and apply a comprehensive security policy. A security policy is an XML file you can use to configure services, network security, registry values, and audit policies. Because security policies are role-based and feature-based, you generally need to create a separate policy for each of your standard server configurations. For example, if your organization uses domain controllers, file servers, and print servers, you might want to create a separate policy for each of these server types. If your organization has mail servers, database servers, and combined file/print servers as well as domain controllers, you should create separate policies tailored to these server types.

You can use the Security Configuration Wizard to do the following:

- Create a security policy.
- Edit a security policy.
- Apply a security policy.
- Roll back the last-applied security policy.

Security policies can incorporate one or more security templates. Much like you can with security templates, you can apply a security policy to the currently logged-on computer using the Security Configuration Wizard. Through Group Policy, you can apply a security policy to multiple computers as well. By default, security policies created with the Security Configuration Wizard are saved in the %System-Root%\security\msscw\Policies folder.

The command-line counterpart to the graphical wizard is the Scwcmd (Scwcmd.exe) utility. At an elevated administrator prompt, you can use Scwcmd Analyze to determine where a computer is in compliance with a security policy and Scwcmd Configure to apply a security policy.

## Creating Security Policies

The Security Configuration Wizard allows you to configure policy only for roles and features that are installed on a computer when you run the wizard. The precise step-by-step process for creating security policies depends on the server roles and features available on the currently logged-on computer. That said, the general configuration sections presented in the wizard are the same regardless of the computer configuration.

The Security Configuration Wizard has the following configuration sections:

- **Role-Based Service Configuration** Configures the startup mode of system services based on a server's installed roles, installed features, installed options, and required services.
- **Network Security** Configures inbound and outbound security rules for Windows Firewall With Advanced Security based on installed roles and installed options.
- **Registry Settings** Configures protocols used to communicate with other computers based on installed roles and installed options.

- **Audit Policy** Configures auditing on the selected server based on your preferences.
- **Save Security Policy** Allows you to save and view the security policy. You can also include one or more security templates.

With this in mind, you can create a security policy by following these steps:

1. Start the Security Configuration Wizard. In Server Manager, you can do this by tapping or clicking Tools, Security Configuration Wizard. On the Welcome page of the wizard, tap or click Next.
2. On the Configuration Action page, review the actions you can perform. (See Figure 5-11.) Create A New Security Policy is selected by default. Tap or click Next.

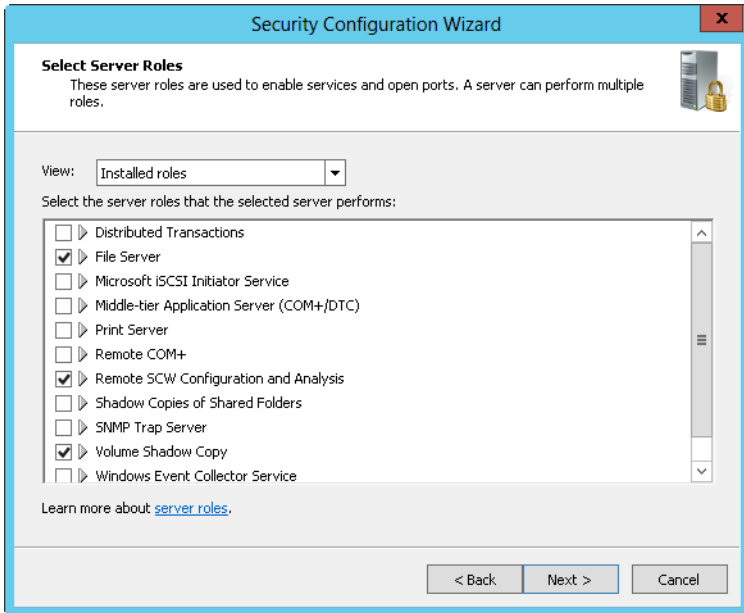


**FIGURE 5-11** Review the actions you can perform.

3. On the Select Server page, select the server you want to use as a baseline for this security policy. The baseline server is the server on which the roles, features, and options you want to work with are installed. The currently logged-on computer is selected by default. To choose a different computer, tap or click Browse. In the Select Computer dialog box, type the name of the computer and then tap or click Check Names. Select the computer account you want to use, and then tap or click OK.
4. When you tap or click Next, the wizard collects the security configuration and stores it in a security configuration database. On the Processing Security Configuration Database page, tap or click View Configuration Database to

view the settings in the database. After you review the settings in the SCW Viewer, return to the wizard and tap or click Next to continue.

- Each configuration section has an introductory page. The first introductory page is the one for Role-Based Service Configuration. Tap or click Next.
- The Select Server Roles page, shown in Figure 5-12, lists the installed server roles. Select each role that should be enabled. Clear the check box for each role that should be disabled. Selecting a role enables services, inbound ports, and settings required for that role. Clearing a role disables services, inbound ports, and settings required for that role, provided that they aren't required by an enabled role. Tap or click Next.



**FIGURE 5-12** Select the server roles to enable.

- On the Select Client Features page, you'll see the installed client features used to enable services. Select each feature that should be enabled. Clear each feature that should be disabled. Selecting a feature enables services required for that feature. Clearing a feature disables services required for that feature, provided that they aren't required by an enabled feature. Tap or click Next.
- On the Select Administration And Other Options page, you'll see the installed options used to enable services and open ports. Select each option that should be enabled. Clear each option that should be disabled. Selecting an option enables services required for that option. Clearing an option disables services required for that option, provided that they aren't required by an enabled option. Tap or click Next.

9. On the Select Additional Services page, you'll see a list of additional services found on the selected server while processing the security configuration database. Select each service that should be enabled. Clear each service that should be disabled. Selecting a service enables services required for that service. Clearing a service disables services required for that service, provided that they aren't required by an enabled service. Tap or click Next.
10. On the Handling Unspecified Services page, indicate how unspecified services should be handled. Unspecified services are services that are not installed on the selected server and are not listed in the security configuration database. By default, the startup mode of unspecified services is not changed. To disable unspecified services instead, select Disable The Service. Tap or click Next.
11. On the Confirm Service Changes page, review the services that will be changed on the selected server if the security policy is applied. Note the current startup mode and the startup mode that will be applied by the policy. Tap or click Next.
12. On the introductory page for Network Security, tap or click Next. On the Network Security Rules page, you'll see a list of firewall rules needed for the roles, features, and options you previously selected. You can add, edit, or remove inbound and outbound rules using the options provided. Tap or click Next when you are ready to continue.
13. On the introductory page for Registry Settings, tap or click Next. On the Require SMB Security Signatures page, review the server message block (SMB) security signature options. By default, minimum operating system requirements and digital signing are used, and you won't want to change these settings. Tap or click Next.
14. For domain controllers and servers with LDAP, on the Require LDAP Signing page, you can set minimum operating system requirements for all directory-enabled computers that access Active Directory.
15. On the Outbound Authentication Methods page, choose the methods that the selected server uses to authenticate with remote computers. Your choices set the outbound LAN Manager authentication level that will be used. If the computer communicates only with domain computers, select Domain Accounts but do not select the other options. This will ensure that the computer uses the highest level of outbound LAN Manager authentication. If the computer communicates with both domain and workgroup computers, select Domain Accounts and Local Accounts On The Remote Computers. In most cases, you won't want to select the file-sharing option because this will result in a substantially lowered authentication level. Tap or click Next.
16. The outbound authentication methods you choose determine what additional Registry Settings–related pages are displayed. Keep the following in mind:
  - If you don't select any outbound authentication methods, the outbound LAN Manager authentication level is set as Send NTLMv2 Response Only and an additional page is displayed to allow you to set the inbound

authentication method. On the Inbound Authentication Using Domain Accounts page, choose the types of computers from which the selected server will accept connections. Your choices set the inbound LAN Manager authentication level that will be used. If the computer communicates only with Windows XP Professional or later computers, clear both options. This ensures that the computer uses the highest level of inbound LAN Manager authentication. If the computer communicates with older PCs, accept the default selections. Tap or click Next.

- If you select domain accounts, local accounts, or both, you'll have additional related pages that allow you to set the LAN Manager authentication level used when making outbound connections. You'll also be able to specify that you want to synchronize clocks with this server's clock. Inbound authentication is set as Accept All.
- If you allow file sharing passwords for early releases of Windows, the outbound LAN Manager authentication level is set as Send LM & NTLM Only and the inbound authentication level is set as Accept All. Because of this, when you tap or click Next, the Registry Settings Summary page is displayed.

- 17.** On the Registry Settings Summary page, review the values that will be changed on the selected server if the security policy is applied. Note the current value and the value that will be applied by the policy. Tap or click Next.
- 18.** On the introductory page for Audit Policy, tap or click Next. On the System Audit Policy page, configure the level of auditing you want. To disable auditing, select Do Not Audit. To enable auditing for successful events, select Audit Successful Activities. To enable auditing for all events, select Audit Successful And Unsuccessful Activities. Tap or click Next.
- 19.** On the Audit Policy Summary page, review the settings that will be changed on the selected server if the security policy is applied. Note the current setting and the setting that will be applied by the policy. Tap or click Next.
- 20.** On the introductory page for Save Security Policy, tap or click Next. On the Security Policy File Name page, you can configure options for saving the security policy and adding one or more security templates to the policy. To view the security policy in the SCW Viewer, tap or click View Security Policy. When you have finished viewing the policy, return to the wizard.
- 21.** To add security templates to the policy, tap or click Include Security Templates. In the Include Security Templates dialog box, tap or click Add. In the Open dialog box, select a security template to include in the security policy. If you add more than one security template, you can prioritize them in case any security configuration conflicts occur between them. Settings from templates higher in the list have priority. Select a template, and then tap or click the Up and Down buttons to prioritize the templates. Tap or click OK.
- 22.** By default, the security policy is saved in the %SystemRoot%\Security\Msscw\Policies folder. Tap or click Browse. In the Save As dialog box, select a different save location for the policy if necessary. After you type a name for the security policy, tap or click Save. The default or selected folder path and file name are then listed in the Security Policy File Name text box.



23. Tap or click Next. On the Apply Security Policy page, you can choose to apply the policy now or later. Tap or click Next, and then tap or click Finish.

## Editing Security Policies

You can use the Security Configuration Wizard to edit a security policy by following these steps:

1. Start the Security Configuration Wizard. In Server Manager, you can do this by tapping or clicking Tools, Security Configuration Wizard. When the wizard starts, tap or click Next.
2. On the Configuration Action page, select Edit An Existing Security Policy and then tap or click Browse. In the Open dialog box, select the security policy you want to work with and then tap or click Open. Security policies end with the .xml extension. Tap or click Next.
3. Follow steps 3–23 of the procedure in the section “Creating Security Policies” to edit the configuration of the security policy.

## Applying Security Policies

You can use the Security Configuration Wizard to apply a security policy by following these steps:

1. Start the Security Configuration Wizard. In Server Manager, you can do this by tapping or clicking Tools, Security Configuration Wizard. When the wizard starts, tap or click Next.
2. On the Configuration Action page, select Apply An Existing Security Policy and then tap or click Browse. In the Open dialog box, select the security policy you want to work with and then tap or click Open. Security policies end with the .xml extension. Tap or click Next.
3. On the Select Server page, select the server you want to apply the security policy to. The currently logged-on computer is selected by default. To choose a different computer, tap or click Browse. In the Select Computer dialog box, type the name of the computer and then tap or click Check Names. Select the computer account you want to use, and then tap or click OK.
4. Tap or click Next. On the Apply Security Policy page, tap or click View Security Policy to view the security policy in the SCW Viewer. When you have finished viewing the policy, return to the wizard.
5. Tap or click Next to apply the policy to the selected server. When the wizard finishes applying the policy, tap or click Next, and then tap or click Finish.

## Rolling Back the Last-Applied Security Policy

You can use the Security Configuration Wizard to roll back the last security policy you applied by following these steps:

1. Start the Security Configuration Wizard. In Server Manager, you can do this by tapping or clicking Tools, Security Configuration Wizard. When the wizard starts, tap or click Next.

2. On the Configuration Action page, select Rollback The Last Applied Security Policy and then tap or click Next.
3. On the Select Server page, select the server on which you want to roll back the last security policy you applied. The currently logged-on computer is selected by default. To choose a different computer, tap or click Browse. In the Select Computer dialog box, type the name of the computer and then tap or click Check Names. Select the computer account you want to use, and then tap or click OK.
4. Tap or click Next. On the Rollback Security Configuration page, tap or click View Rollback File to view the details of the last-applied security policy in the SCW Viewer. When you finish viewing the policy, return to the wizard.
5. Tap or click Next to roll back the policy to the selected server. When the wizard finishes the rollback process, tap or click Next, and then tap or click Finish.

## Deploying a Security Policy to Multiple Computers

In an organization with many computers, you probably won't want to apply a security policy to each computer separately. As discussed in "Deploying Security Templates to Multiple Computers" earlier in this chapter, you might want to apply a security policy through Group Policy, and you might want to create computer OUs for this purpose.

Once you've created the necessary OUs, you can use the `Scwcmd` utility's `transform` command to create a GPO that includes the settings in the security policy (and any security templates attached to the policy). You then deploy the settings to computers by linking the new GPO to the appropriate OU or OUs. By default, security policies created with the Security Configuration Wizard are saved in the `%SystemRoot%\security\msscw\Policies` folder.

Use the following syntax to transform a security policy:

```
scwcmd transform /p:FullFilePathToSecurityPolicy /g:GPOName
```

where *FullFilePathToSecurityPolicy* is the full file path to the security policy's .xml file, and *GPOName* is the display name for the new GPO. Consider the following example:

```
scwcmd transform /p:"c:\users\wrs\documents\fspolicy.xml"  
/g: "FileServer GPO"
```

When you create the GPO, you can link the GPO by following these steps:

1. In the Group Policy Management Console (GPMC), select the OU you want to work with. In the right pane, the Linked Group Policy Objects tab shows the GPOs that are currently linked to the selected OU (if any).
2. Press and hold or right-click the OU to which you want to link the previously created GPO, and then select Link An Existing GPO. In the Select GPO dialog box, select the GPO you want to link to and then tap or click OK.

When Group Policy is refreshed for computers in the applicable OU, the policy settings in the GPO are applied.

Because you created a new GPO and linked the GPO to the appropriate level in the Active Directory structure, you can restore the computers to their original state by removing the link to the GPO. To remove a link to a GPO, follow these steps:

- 1.** In the GPMC, select and then expand the OU you want to work with. In the right pane, the Linked Group Policy Objects tab shows the GPOs that are currently linked to the selected OU.
- 2.** Press and hold or right-click the GPO. On the shortcut menu, the Link Enabled option should have a check mark to show it is enabled. Clear this option to remove the link.



PART II

# Windows Server 2012 Directory Services Administration

- CHAPTER 6**    Introducing Active Directory    **217**
- CHAPTER 7**    Core Active Directory Administration    **249**
- CHAPTER 8**    Creating User and Group Accounts    **295**
- CHAPTER 9**    Managing User and Group Accounts    **347**



# Using Active Directory

- Introducing Active Directory 217
- Working with Domain Structures 223
- Working with Active Directory Domains 230
- Understanding the Directory Structure 237
- Using the Active Directory Recycle Bin 244

**A**ctive Directory Domain Services (AD DS) is an extensible and scalable directory service you can use to efficiently manage network resources. As an administrator, you need to be deeply familiar with how Active Directory technology works, and that's exactly what this chapter is about. If you haven't worked with Active Directory technology before, you'll notice immediately that the technology is fairly advanced and has many features.

## Introducing Active Directory

---

Since Windows 2000, Active Directory has been the heart of Microsoft 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.

## Active Directory and DNS

Active Directory uses Domain Name System (DNS). 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.microsoft.com*. Here, *zeta* represents the name of an individual computer, *microsoft* 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 *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, which you can then use for different offices, divisions, or geographic locations. For example, the FQDN for a computer at Microsoft's Seattle office could be designated as *jacob.seattle.microsoft.com*. Here, *jacob* is the computer name, *seattle* is the subdomain, and *microsoft.com* is the parent domain. Another term for a subdomain is a *child domain*.

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. Working with DNS is covered in Chapter 16, "Optimizing DNS."

With Windows Server 2012, you install Active Directory in a two-part process. First you start the process in Server Manager by tapping or clicking Manage and then 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. This 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 2012, not only can you uninstall a role or feature, but you also can uninstall and remove the payload for that feature or role using the `-Remove` parameter of the `Uninstall-WindowsFeature` cmdlet.

You can restore a removed payload using the `Install-WindowsFeature` cmdlet. By default, payloads are restored via Windows Update. 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 tapping or clicking the Promote This Server To A Domain Controller link on the Installation Progress page and then use this wizard to configure the role. This wizard replaces `Dcpromo.exe`, which was used previously for promoting domain controllers. The wizard also will run `Adprep.exe` to prepare schema as appropriate. If you do not run `Adprep.exe` separately beforehand and you are installing the first domain controller that runs Windows Server 2012 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 RODC in a forest, you need to provide credentials for a member of the Enterprise Admins group.

If DNS isn't already installed, you are prompted to install it. 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'll 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 2012 and the forest functional level is Windows Server 2003 or higher, 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 it is compatible with Windows Server 2012 as necessary.

## Read-Only Domain Controller Deployment

When the domain and forest are operating at the Windows Server 2003 functional level or higher and your primary domain controller (PDC) emulator for a domain is running Windows Server 2008 or later, you can deploy read-only domain controllers (RODCs). Any domain controller running Windows Server 2008 R2 or later can be configured as an 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.

The first Windows Server 2008 R2 or later domain controller installed in a forest or domain cannot be an RODC. However, you can configure subsequent domain controllers as read-only.

**MORE INFO** The domain and forest must have the correct schema level to support RODCs and must also be prepared to work with RODCs. Previously, in some cases, this required that you prepare the forest and domain schemas for Windows Server 2008 R2 and then update the forest schema again for RODCs. When you use Server Manager,

Windows Server 2012, and the Windows Server 2003 or higher forest functional level, any necessary preparations are done automatically as part of DC and RODC deployment.

## Active Directory Features for Windows Server 2008 R2

If you are upgrading to Windows Server 2012 but haven't yet deployed Windows Server 2008 R2, you'll want to know about related features for Active Directory. When you are using Windows Server 2008 R2 and Windows Server 2012 and have deployed these operating systems on all domain controllers throughout the domains in your Active Directory forest, your domains can operate at the Windows Server 2008 R2 domain functional level, and the forest can operate at the Windows Server 2008 R2 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 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 this chapter.
- **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 8, "Creating User and Group Accounts."
- **Managed virtual accounts** Introduces a special type of local computer account for managed services that provides the ability to access the network with a computer identity in a domain environment. For more information, see "Using Virtual Accounts" in Chapter 8.

**REAL WORLD** Technically, you can use managed service accounts and managed virtual accounts in a mixed-mode domain environment. However, you have to manually manage SPNs for managed service accounts and the Active Directory schema must be compatible with Windows Server 2008 R2 and higher.

- **Authentication Mechanism Assurance** Improves the authentication process by allowing administrators to control resource access based on whether a user logs on using a certificate-based logon method. Thus, an administrator can specify that a user has one set of access permissions when logged on using a smart card and a different set of access permissions when not logged on using a smart card.

Other improvements don't require that you raise domain or forest functional levels, but they do require that you use Windows Server 2012. These improvements include the following:

- **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.

- **Active Directory module for Windows PowerShell** Provides cmdlets for managing Active Directory when you are working with Windows PowerShell. Import the Active Directory module by typing **import-module activedirectory** at the PowerShell prompt.
- **Active Directory Administrative Center** Provides a task-orientated interface for managing Active Directory. In Server Manager, tap or click Tools and then tap or click Active Directory Administrative Center.
- **Active Directory Web Services** Introduces a web service interface for Active Directory domains.

These features are discussed in more detail in Chapter 7, “Core Active Directory Administration.”

## Active Directory Features for Windows Server 2012

Active Directory Domain Service in Windows Server 2012 has many additional features that give administrators additional options for implementing and managing Active Directory. Table 6-1 lists key features. At the least, these features require that you update the Active Directory schema in your forests and domains for Windows Server 2012. You also might need to update the domain, forest, or both functional levels to the new Windows Server 2012 operating level.

**TABLE 6-1** Key Active Directory Features for Windows Server 2012

FEATURE	BENEFITS	REQUIREMENTS
Active Directory–based activation	Allows you to use AD to automatically activate clients running Windows 8 and Windows Server 2012. Any client connected to the service is activated.	Volume Licensing; Active Directory schema must be updated for Windows Server 2012; key is set using Volume Activation server role or command line.
Claims-based policy controls	Allows access and audit policies to be defined flexibly.	Claims policy must be enabled for Default Domain Controllers Policy; file servers must run Windows Server 2012; domain must have at least one Windows Server 2012 domain controller.
Deferred index creation	Allows deferring of index creation within the directory until <i>UpdateSchemaNow</i> is received or the domain controller is rebooted.	The domain controller must run Windows Server 2012.

FEATURE	BENEFITS	REQUIREMENTS
Enhanced Fine-Grained Password Policy	Allows administrators to use Active Directory Administrative Center for Windows Server 2012 to create and manage password-settings objects (PSOs).	Windows Server 2008 or higher domain functional level.
Enhanced Recycle Bin	Allows administrators to recover deleted objects using Active Directory Administrative Center for Windows Server 2012.	Domain must have Recycle Bin enabled and Windows Server 2008 R2 or higher forest functional level.
Group Managed Service Accounts	Allows multiple services to share a single managed service account.	Active Directory schema must be updated for Windows Server 2012; must have at least one Windows Server 2012 domain controller; services must run on Windows Server 2012.
Kerberos constrained delegation across domains	Allows managed service accounts to act on behalf of users across domains and forests.	Each affected domain must have at least one Windows Server 2012 domain controller; front-end server must run Windows Server 2012; back-end server must run Windows Server 2003 or later; and other requirements as well.
Kerberos with Armoring	Improves domain security; allows a domain-joined client and domain controller to communicate over a protected channel.	Windows Server 2012 domain controllers; Windows Server 2012 domain functional level; 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.	Domain must be Direct Access-enabled, and domain controllers must run Windows Server 2012.

FEATURE	BENEFITS	REQUIREMENTS
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.	A domain controller with RID role must run Windows Server 2012, and domain controllers must run Windows Server 2012.
Server Manager integration	Allows you to perform all the steps required to deploy local and remote domain controllers.	Windows Server 2012; forest functional level of Windows Server 2003 or higher.
Virtual domain controller cloning	Allows you to safely deploy virtualized replicas of domain controllers. Also helps maintain domain controller state.	A domain controller with PDC Emulator role must run Windows Server 2012, and virtual domain controllers must run Windows Server 2012 as well.

## 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.

# Understanding Domains

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, as shown in Figure 6-1. Within a domain's directory database, you'll find objects defining accounts for users, groups, and computers as well as shared resources such as printers and folders.



**FIGURE 6-1** This network diagram depicts a wide area network (WAN) with multiple sites and subnets.

**NOTE** User and group accounts are discussed in Chapter 8. Computer accounts and the various types of computers used in Windows Server domains are discussed in "Working with Active Directory Domains" later in this chapter.

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

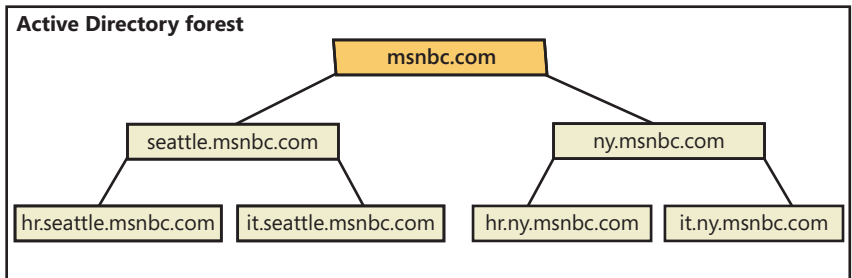
- **Windows Server 2003** Supports domain controllers running Windows Server 2003 and later.
- **Windows Server 2008** Supports domain controllers running Windows Server 2008 and later.
- **Windows Server 2008 R2** Supports domain controllers running Windows Server 2008 R2 and Windows Server 2012.
- **Windows Server 2012** Supports domain controllers running Windows Server 2012.

For further discussion of domain functional levels, see “Working with Domain Functional Levels” later in this chapter.

## Understanding Domain Forests and Domain Trees

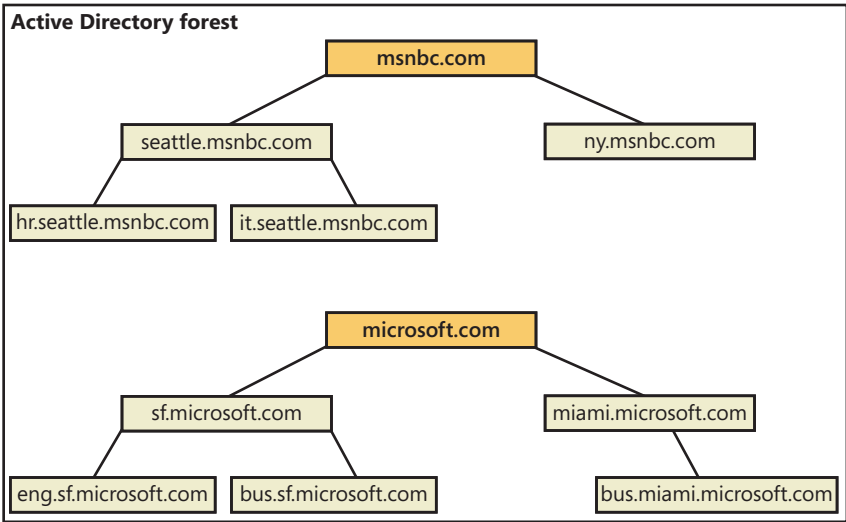
Each Active Directory domain has a DNS domain name, such as microsoft.com. One or more domains sharing the same directory data are referred to as a *forest*. The domain names within this forest can be noncontiguous or contiguous in the DNS naming hierarchy.

When domains have a contiguous naming structure, they’re said to be in the same domain tree. Figure 6-2 shows an example of a domain tree. In this example, the root domain msnbc.com has two child domains: seattle.msnbc.com and ny.msnbc.com. These domains, in turn, have subdomains. All the domains are part of the same tree because they have the same root domain.



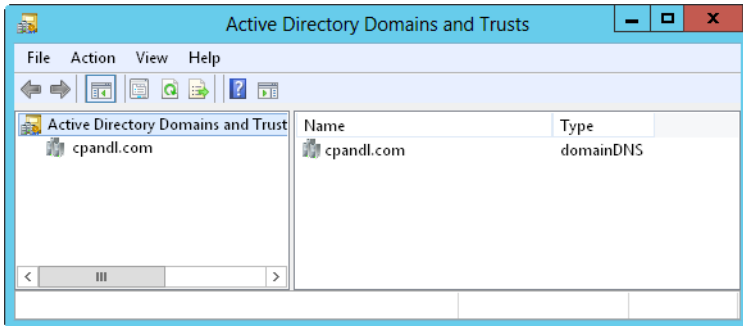
**FIGURE 6-2** Domains in the same tree share a contiguous naming structure.

If the domains in a forest have noncontiguous DNS names, they form separate domain trees within the forest. As shown in Figure 6-3, a domain forest can have one or more domain trees. In this example, the msnbc.com and microsoft.com domains form the roots of separate domain trees in the same forest.



**FIGURE 6-3** Multiple trees in a forest with noncontiguous naming structures.

You can access domain structures by using Active Directory Domains And Trusts, shown in Figure 6-4. Active Directory Domains And Trusts is a snap-in for the Microsoft Management Console (MMC). You can also start it from the Tools menu in Server Manager. You'll find separate entries for each root domain. In Figure 6-4, the active domain is cpandl.com.



**FIGURE 6-4** Use Active Directory Domains And Trusts to work with domains, domain trees, and domain forests.

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

- **Windows Server 2003** Supports domain controllers running Windows Server 2003 and later.



- **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 Windows Server 2012.
- **Windows Server 2012** Supports domain controllers running Windows Server 2012.

When all domains within a forest are operating in Windows Server 2003 forest functional level, you'll see improvements over earlier implementations in global catalog replication and replication efficiency. Because link values are replicated, you might see improved intersite replication as well. You can deactivate schema class objects and attributes; use dynamic auxiliary classes; rename domains; and create one-way, two-way, and transitive forest trusts.

The Windows Server 2008 forest functional level offers incremental improvements over the Windows Server 2003 forest functional level in Active Directory performance and features. When all domains within a forest are operating in this mode, you'll see improvements in both intersite and intrasite replication throughout the organization. Domain controllers can use Distributed File System (DFS) replication rather than File Replication Service (FRS) replication as well. In addition, Windows Server 2008 security principals are not created until the primary domain controller (PDC) emulator operations master in the forest root domain is running Windows Server 2008.

The Windows Server 2008 R2 forest functional level has several additional features. These features include the Active Directory Recycle Bin, managed service accounts, and Authentication Mechanism Assurance.

Although Active Directory for Windows Server 2012 has many enhancements, most of these enhancements require using only Windows Server 2012 domain controllers and schema. The main exception is for Kerberos with Armoring, which requires the Windows Server 2012 domain functional level.

Generally, you cannot lower the forest functional level once you raise it. However, when you raise the forest functional level to Windows Server 2012, you can lower it to Windows Server 2008 R2. If Active Directory Recycle Bin has not been enabled, you can also lower the forest functional level from Windows Server 2012 to Windows Server 2008 R2 or Windows Server 2008 or from Windows Server 2008 R2 back to Windows Server 2008. You cannot roll the domain functional level back to Windows Server 2003 or lower.

## Understanding Organizational Units

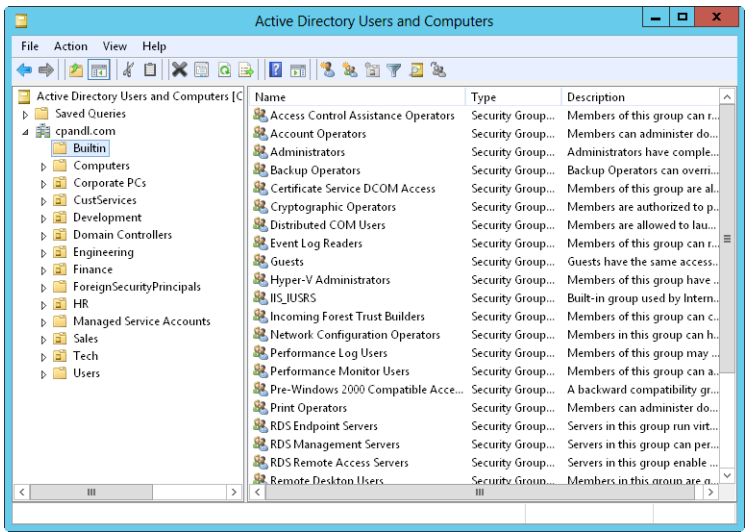
Organizational units (OUs) are subgroups within domains that often mirror an organization's functional or business structure. You can also think of OUs as logical containers into which you place accounts, shared resources, and other OUs. For example, you could create OUs named HumanResources, IT, Engineering, 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 only 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, as shown in Figure 6-5. This utility is a snap-in for the MMC, and you can also start it from the Tools menu in Server Manager.



**FIGURE 6-5** Use Active Directory Users And Computers to manage users, groups, computers, and organizational units.

## Understanding Sites and Subnets

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. Here, the network address 192.168.19.9 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 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, shown in Figure 6-6. 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.

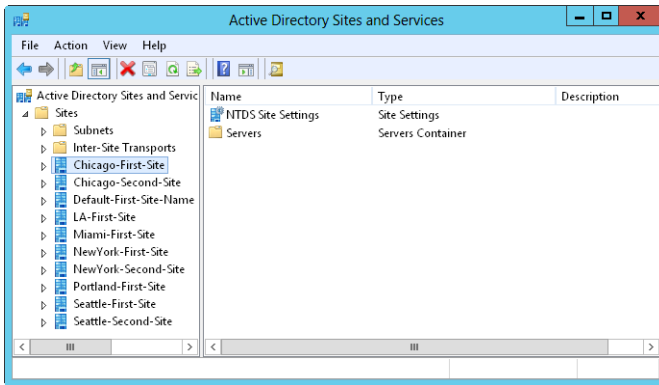


FIGURE 6-6 Use Active Directory Sites And Services to manage sites and subnets.

## 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.microsoft.com`, to numeric TCP/IP addresses, such as `172.16.18.8`. To learn more about DNS and DNS domains, see Chapter 16.

## Using Computers with 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. You promote member servers to domain controllers by installing Active Directory Domain Services. You demote domain controllers to member servers by uninstalling Active Directory Domain Services. You use the Add Role And Features and Remove Role And Features wizards to add or remove Active Directory Domain Services. You promote or demote a server through the Active Directory Installation Wizard (`Dcpromo.exe`).

Domains can have one or more domain controllers. When a domain has multiple domain controllers, the controllers automatically replicate directory data with one another 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 “Understanding Operations Master Roles” later in this chapter.

Every Windows 2000 or later 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 “Understanding the Directory Structure” later in this chapter.

## Working with Domain Functional Levels

To support domain structures, Active Directory includes support for the following domain functional levels:

- **Windows Server 2003 mode** When the domain is operating in Windows Server 2003 mode, the directory supports domain controllers running Windows Server 2008 R2, Windows Server 2008, and Windows Server 2003. A domain operating in Windows Server 2003 mode can use universal groups, group nesting, group type conversion, easy domain controller renaming, update logon time stamps, and Kerberos KDC key version numbers.
- **Windows Server 2008 mode** When the domain is operating in Windows Server 2008 mode, the directory supports Windows Server 2008 and Windows Server 2008 R2 domain controllers. Windows Server 2003 domain controllers are no longer supported. A domain operating in Windows Server 2008 mode can use additional Active Directory features, including the DFS replication service for enhanced intersite and intrasite replication.
- **Windows Server 2008 R2 mode** When the domain is operating in Windows Server 2008 R2 mode, the directory supports only Windows Server 2008 R2 domain controllers. Windows Server 2003 and Windows Server 2008 domain controllers are no longer supported. A domain operating in Windows Server 2008 R2 mode can use Active Directory Recycle Bin,

managed service accounts, Authentication Mechanism Assurance, and other important Active Directory enhancements.

- **Windows Server 2012 mode** When the domain is operating in Windows Server 2012 mode, the directory supports only Windows Server 2012 domain controllers. Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 domain controllers are no longer supported. Active Directory schema for Windows Server 2012 includes many enhancements, but only the Kerberos with Armoring feature requires this mode.

Generally, you cannot lower the domain functional level once you raise it. However, when you raise the domain functional level to Windows Server 2008 R2 or Windows Server 2012 and the forest functional level is Windows Server 2008 or lower, you have the option of rolling the domain functional level back to Windows Server 2008 or Windows Server 2008 R2. You cannot roll the domain functional level back to Windows Server 2003 or lower.

## Using Windows Server 2003 Functional Level

Every domain in your enterprise should be operating at the Windows Server 2003 functional level or higher, if possible, which will ensure computers in your domains can take advantage of many of the most recent enhancements to Active Directory. After you decommission Windows NT structures and upgrade the Windows 2000 structures in your organization, you can change the functional level to Windows Server 2003 mode operations.

Before updating Windows 2000 domain controllers, you should prepare the domain for upgrade. To do this, you need to update the forest and the domain schema so that they are compatible with Windows Server 2003 domains. A tool called `Adprep.exe` is provided to automatically perform the update for you. All you need to do is run the tool on the schema operations master in the forest and then on the infrastructure operations master for each domain in the forest. As always, you should test any procedure in a lab before performing it in a production environment.

On the Windows Server 2003 installation media, you'll find `Adprep` and related files in the `i386` subfolder. Follow these steps to perform the upgrade:

1. On the schema operations master in the forest, run **`<cdrom>:\i386\adprep.exe /forestprep`**. You need to use an administrator account that is a member of Enterprise Admins, Schema Admins, or Domain Admins in the forest root domain.
2. On the infrastructure operations master for each domain in the forest, run **`<cdrom>:\i386\adprep.exe /domainprep`**. You need to use an account that is a member of the Domain Admins group in an applicable domain.

**NOTE** To determine which server is the current schema operations master for the domain, open a command prompt and type **`dsquery server -hasfsmo schema. A`**

directory service path string is returned containing the name of the server, such as "CN=CORPSEVER01,CN=Servers,CN=Default-First-Site-Name,CN=Sites,CN=Configuration,DC=microsoft,DC=com." This string tells you that the schema operations master is CORPSEVER01 in the microsoft.com domain.

**NOTE** To determine which server is the current infrastructure operations master for the domain, start a command prompt and type **dsquery server -hasfsmo infr**.

**REAL WORLD** Generally speaking, anything you can type at a command prompt can be typed at the PowerShell prompt as well. This is possible because PowerShell looks for external commands and utilities as part of its normal 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, and 5) external commands, utilities, and files. Thus, if any element in steps 1 through 4 of the execution order has the same name as a command, that element will run instead of the expected command.

After upgrading your servers, you can raise the domain and forest functionality to take advantage of the additional Active Directory features of the Windows Server 2003 functional level. Keep in mind that once you upgrade, you can use only Windows Server 2003 and later resources in the domain and you can't go back to any other mode. You should use Windows Server 2003 mode only when you're certain that you don't need old Windows NT domain structures, Windows NT backup domain controllers (BDCs), or Windows 2000 domain structures.

## Using Windows Server 2008 Functional Level

After you upgrade the Windows 2000 and Windows Server 2003 structures in your organization, you can change the functional level to Windows Server 2008 mode operations.

Before updating Windows Server 2003 domain controllers, you should prepare the domain for Windows Server 2008. To do this, you need to use Adprep.exe to update the forest and the domain schema so that they are compatible with Windows Server 2008 domains. Follow these steps:

1. On the schema operations master in the forest, copy the contents of the Sources\Adprep folder from the Windows Server 2008 installation media to a local folder, and then run **adprep /forestprep**. If you plan to install any read-only domain controllers, you should also run **adprep /rodcprep**. You need to use an administrator account that is a member of Enterprise Admins, Schema Admins, or Domain Admins in the forest root domain.

2. On the infrastructure operations master for each domain in the forest, copy the contents of the Sources\Adprep folder from the Windows Server 2008 installation media to a local folder, and then run **adprep /domainprep**. You need to use an account that is a member of the Domain Admins group in an applicable domain.
3. If you haven't previously run **adprep /domainprep /gpprep** in each domain, you need to manually perform this task. Server Manager for Windows Server 2012 will not prepare Group Policy for you. Note that Group Policy needs to be prepared only the first time you deploy domain controllers running Windows Server 2003 SP1 or later. **Adprep /gpprep** modifies the access control entries (ACEs) for all Group Policy Object (GPO) folders in the SYSVOL directory to grant read access to all enterprise domain controllers. This level of access is required to support Resultant Set of Policy (RSOP) for site-based policy and causes the NT File Replication Service (NTFRS) to resend all GPOs to all domain controllers.

As always, you should test any procedure in a lab before performing it in a production environment.

**NOTE** To determine which server is the current schema operations master for the domain, start a command prompt and type **dsquery server -hasfsmo schema**. To determine which server is the current infrastructure operations master for the domain, start a command prompt and type **dsquery server -hasfsmo infr**.

After upgrading all domain controllers to Windows Server 2008, you can raise the domain and forest level functionality to take advantage of additional Active Directory features. If you do this, you can use only Windows Server 2008 or later resources in the domain and you can't go back to any other mode. You should use Windows Server 2008 mode only when you're certain that you don't need old Windows NT domain structures, Windows NT BDCs, or Windows 2000 or Windows Server 2003 domain structures.

## Using Windows Server 2008 R2 Functional Level

Windows Server 2008 R2 and Windows Server 2012 run only on 64-bit hardware. You'll likely need to install Windows Server 2008 R2 and Windows Server 2012 on new hardware rather than hardware designed for earlier releases of Windows Server.

Before updating Windows Server 2008 domain controllers, you should prepare the domain for Windows Server 2008 R2. To do this, you need to use Adprep.exe to update the forest and the domain schema so that they are compatible with Windows Server 2008 R2 domains. Follow these steps:

1. On the schema operations master in the forest, copy the contents of the Support\Adprep folder from the Windows Server 2008 R2 installation media to a local folder, and then run **adprep /forestprep**. If you plan to install any read-only domain controllers, you should also run **adprep /rodcprep**. You need to use an administrator account that is a member of Enterprise Admins, Schema Admins, or Domain Admins in the forest root domain.



2. On the infrastructure operations master for each domain in the forest, copy the contents of the Support\Adprep folder from the Windows Server 2008 R2 installation media to a local folder, and then run **adprep /domainprep**. You need to use an account that is a member of the Domain Admins group in an applicable domain.

As always, you should test any procedure in a lab before performing it in a production environment.

**NOTE** To determine which server is the current schema operations master for the domain, start a command prompt and type **dsquery server -hasfsmo schema**. To determine which server is the current infrastructure operations master for the domain, start a command prompt and type **dsquery server -hasfsmo infr**.

After upgrading all domain controllers to Windows Server 2008 R2, you can raise the domain and forest level functionality to take advantage of the latest Active Directory features. If you do this, you can use only Windows Server 2008 R2 resources in the domain. You should use Windows Server 2008 R2 mode only when you're certain that you don't need old Windows NT domain structures; Windows NT BDCs; or Windows 2000, Windows Server 2003, or Windows Server 2008 domain structures.

## Using Windows Server 2012 Functional Level

Like Windows Server 2008 R2, Windows Server 2012 runs only on 64-bit hardware and you'll likely need to install Windows Server 2012 on new hardware rather than on hardware designed for earlier releases of Windows Server. Unlike earlier releases of Windows Server, the domain and forest preparations required for updating Active Directory schema don't need to be performed manually. Instead, when you use Server Manager for Windows Server 2012 and the forest functional level is Windows Server 2003 or higher, any necessary preparations are done automatically when you deploy a domain controller running Windows Server 2012. This means the Configuration Wizard automatically updates forest and domain schema.

You also have the option of manually preparing for Windows Server 2012. To do this, you can use Adprep.exe to update the forest and the domain schema so that they are compatible with Windows Server 2012 domains. The steps are similar to those discussed in the previous section.

After upgrading all domain controllers to Windows Server 2012, you can raise the domain and forest level functionality to take advantage of the latest Active Directory features. If you do this, you can use only Windows Server 2012 resources in the domain.

## Raising or Lowering Domain and Forest Functionality

Domains operating in a Windows Server 2003 or higher functional level can use universal groups, group nesting, group type conversion, update logon time stamps, and Kerberos KDC key version numbers. In this mode or higher, administrators can do the following:

- Rename domain controllers without having to demote them first.

- Rename domains running on Windows Server 2003 or higher domain controllers.
- Create extended two-way trusts between two forests.
- Restructure domains in the domain hierarchy by renaming them and putting them at different levels.
- Take advantage of replication enhancements for individual group members and global catalogs.

As compared to earlier implementations, forests operating in a Windows Server 2003 or higher functional level have better global catalog replication and intrasite and intersite replication efficiency, as well as the ability to establish one-way, two-way, and transitive forest trusts.

**REAL WORLD** The domain and forest upgrade process can generate a lot of network traffic as information is being replicated around the network. 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) before performing this operation.

You can raise the domain level functionality by following these steps:

1. Open Active Directory Domains And Trusts. In the console tree, press and hold or right-click the domain you want to work with, and then tap or 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 tap or click Raise.
3. Tap or 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, press and hold or right-click the Active Directory Domains And Trusts node, and then tap or 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 tap or click Raise.
3. Tap or 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, tap or click Manage and then tap or click Add Navigation Nodes. In the Add Navigation Nodes dialog box, select the domain you want to work with and then tap or click OK.
2. Select the domain you want to work with by tapping or clicking it in the left pane. In the Tasks pane, tap or click Raise 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 and then tap or click Raise.
4. Tap or 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 tapping or clicking it in the left pane. In the Tasks pane, tap or 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 tap or click Raise.
3. Tap or 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.

Generally, you cannot lower the forest or domain functional level once you raise it. However, there are specific exceptions as discussed previously in this chapter. Keep in mind that if you enabled Active Directory Recycle Bin, you won't be able to lower the forest functional level.

## 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. A feature of Windows Server 2008 or later called *application directory partitions* also changes 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 running Windows Server 2008 or later, including 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 Ntlds.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 2008 or later. 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, as well as 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 on operations master roles, see "Understanding Operations Master Roles" 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 "Configuring Global Catalogs" in Chapter 7. 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. Additionally, users in the site must

be part of a Windows domain running in a Windows Server 2003 or higher functional mode. To learn how to configure caching, see “Configuring Universal Group Membership Caching” in Chapter 7.

## Replication and Active Directory

Regardless of whether you use FRS or DFS replication, the three types of information stored in the directory are 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 GC hosts
- 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 as well as 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 Interface (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 Windows Script Host to create and manage objects in Active Directory.

## Understanding Operations Master Roles

Operations master roles accomplish tasks that are impractical to perform in multi-master 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 type **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 type **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 type **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 type **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 type **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 as well. 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. Here, 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 usually want the RID master and PDC emulator to be on the same server as well. 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 "Exploring Global Catalogs" earlier in this chapter for details.

## 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 non-link-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. This differs substantially from the previously available technique, which used an authoritative restore to recover deleted objects from the Deleted Objects container. Previously, when you deleted an object, most of its non-link-valued attributes were cleared and all of its link-valued attributes were removed, which meant that although you could recover a deleted object, it was not restored to its previous state.

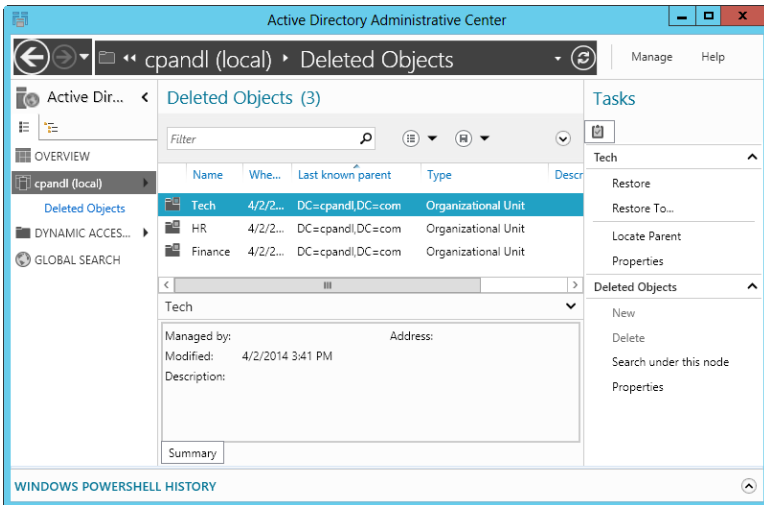
### Preparing Schema for the Recycle Bin

Before you can make the recycle bin available, you must update Active Directory schema with the required recycle bin attributes. You do this by preparing the forest and domain for the Windows Server 2008 R2 functional level or higher. When you do this, the schema is updated, and then every object in the forest is updated with the recycle bin attributes as well. This process is irreversible once 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 to enable the enhanced recycle bin.

After these operations, you can enable and access the recycle bin. Once 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, shown in Figure 6-7. 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.



**FIGURE 6-7** Deleted objects remain in the Deleted Objects container for the deleted object lifetime value.

## 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 previous 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 *Ldp.exe* administration tool or the Active Directory cmdlets for Windows PowerShell. If you updated the Active Directory schema in your forests and domains for Windows Server 2012, you also can enable the enhanced recycle bin, which allows you to recover deleted objects using Active Directory Administrative Center.

Keep in mind that Active Directory blocks access to an object for a short while after it 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 Ldp.exe for Basic Recovery

You can use Ldp.exe to display the Deleted Objects container and recover a deleted object by following these steps:

1. Type **Ldp.exe** in the Apps Search box, and then press Enter.
2. On the Options menu, tap or click Controls. In the Controls dialog box, select Return Deleted Objects in the Load Predefined list, and then tap or click OK.
3. Bind to the server that hosts the forest root domain by choosing Bind from the Connection menu. Select the Bind type, and then tap or click OK.
4. On the View menu, tap or click Tree. In the Tree View dialog box, use the BaseDN list to select the appropriate forest root domain name, such as DC=Cpandl,DC=Com, and then tap or click OK.
5. In the console tree, double-tap or double-click the root distinguished name and locate the CN=Deleted Objects container.
6. Locate and press and hold or right-click the Active Directory object you want to restore, and then tap or click Modify. This displays the Modify dialog box.
7. In the Edit Entry Attribute text box, type **isDeleted**. Do not enter anything in the Values text box.
8. Under Operation, tap or click Delete, and then tap or click Enter.
9. In the Edit Entry Attribute text box, type **distinguishedName**. In Values, type the original distinguished name of this Active Directory object.
10. Under Operation, tap or click Replace. Select the Extended check box, tap or click Enter, and then tap or click Run.

## Using Windows PowerShell for Basic and Advanced Recovery

The Active Directory cmdlets for Windows PowerShell allow you to recover deleted objects using scripts or by typing commands at a 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.

**NOTE** The Active Directory module is not imported into Windows PowerShell by default. Import the Active Directory module by typing **import-module activedirectory** at the PowerShell prompt. For more information, see “Active Directory Administrative Center and Windows PowerShell” in Chapter 7.

To use the Active Directory cmdlets for recovery, you need to open an elevated, administrator PowerShell prompt by pressing and holding or right-clicking the Windows PowerShell entry on the menu and tapping or 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=Cpandl,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=Cpandl,DC=com" -Filter  
{lastKnownParent -eq "OU=Corporate_Services,DC=Cpandl,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. Once you updated the Active Directory schema in your forests and domains for Windows Server 2012, you enable the enhanced recycle bin for use by following these steps:

1. In Active Directory Administrative Center, the local domain is opened for management by default. If you want to work with a different domain, tap or click Manage and then tap or click Add Navigation Nodes. In the Add Navigation Nodes dialog box, select the domain you want to work with and then tap or click OK.

2. Select the domain you want to work with by tapping or clicking it in the left pane. In the Tasks pane, tap or click Enable Recycle Bin and then tap or click OK in the confirmation dialog box.
3. Active Directory will begin replicating the change to all domain controllers in the forest. Once the change is replicated, the enhanced recycle bin will be available for use. If you then tap or click Refresh in Active Directory Administrative Center, you'll see that a Deleted Object container is now available for domains using 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, you'll see a list of deleted objects. 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 tapping or 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 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.cpandl.com domain, you could restore it to the Devs OU in the eng.cpandl.com domain.

# Core Active Directory Administration

- Tools for Managing Active Directory **249**
- Managing Computer Accounts **259**
- Managing Domain Controllers, Roles, and Catalogs **270**
- Managing Organizational Units **281**
- Managing Sites **282**
- Maintaining Active Directory **289**
- Troubleshooting Active Directory **292**

Core Active Directory administration focuses on key tasks you perform routinely with Active Directory Domain Services (AD DS), such as creating computer accounts or joining computers to a domain. In this chapter, you'll learn about the tools you can use to manage Active Directory as well as about specific techniques for managing computers, domain controllers, and organizational units.

## Tools for Managing Active Directory

---

Several sets of tools are available for managing Active Directory, including graphical administration tools, command-line tools, support tools, and Microsoft Windows PowerShell cmdlets.

### Active Directory Administration Tools

Active Directory administration tools are provided as snap-ins for the Microsoft Management Console (MMC). You use the following key tools to manage Active Directory:

- **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.
- **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.

**SECURITY ALERT** 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 Remote Administration Exception policy setting within Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile. Alternatively, type the following at a command prompt on the remote computer: **netsh firewall set portopening tcp 445 smb enable**. See Microsoft Knowledge Base Article 840634 for more details ([support.microsoft.com/default.aspx?scid=kb;en-us;840634](http://support.microsoft.com/default.aspx?scid=kb;en-us;840634)).

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 for AD DS.

## Active Directory Command-Line Tools

Several tools are provided to let you manage Active Directory from the command line:

- **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** As discussed in Chapter 6, "Using Active Directory," Server Manager for Windows Server 2012 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 server running a 64-bit version of Windows Server 2008 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. If either of these operations masters is running Windows Server 2003, the server from which you are running Adprep should be domain joined and you won't be able to use smart card credentials.

- **Dsadd** Adds computers, contacts, groups, organizational units, and users to Active Directory. Type **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. Type **dsget objectname /?** at a command prompt to display help information about using the command, such as **dsget subnet /?**
- **Dsmod** Modifies properties of computers, contacts, groups, organizational units, users, and servers that exist in Active Directory. Type **dsmod object-name /?** at a command prompt to display help information about using the command, such as **dsmod server /?**
- **Dsmove** Moves a single object to a new location within a single domain or renames the object without moving it. Type **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. Type **dsquery /?** at a command prompt to display help information about using the command.
- **Dsrm** Removes objects from Active Directory. Type **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. Type **ntdsutil /?** at a command prompt to display help information about using the command.

While Adprep is located in the \support\adprep folder on the Windows Server 2012 installation media, the other tools become available when you install the Remote Server Management Tools for AD DS.

## Active Directory Support Tools

Many support tools for Active Directory are included in the management tools for AD DS. Table 7-1 lists some of the most useful support tools for configuring, managing, and troubleshooting Active Directory.

**TABLE 7-1** Quick Reference for Active Directory Support Tools

SUPPORT TOOL	EXECUTABLE NAME	DESCRIPTION
ADSI Edit	Adsiedit.msc	Opens and edits the Active Directory Services Interface for domain, schema, and configuration containers
Active Directory Administration Tool	Ldp.exe	Performs Lightweight Directory Access Protocol (LDAP) operations on Active Directory
Directory Services Access Control Lists Utility	Dsacls.exe	Manages access control lists (ACLs) for objects in Active Directory
Distributed File System Utility	Dfsutil.exe	Manages the Distributed File System (DFS) and displays DFS information
DNS Server Troubleshooting Tool	Dnscmd.exe	Manages properties of Domain Name System (DNS) servers, zones, and resource records
Replication Diagnostics Tool	Repadmin.exe	Manages and monitors replication using the command line
Windows Domain Manager	Netdom.exe	Allows domain and trust relationships management from the command line

## Using Active Directory Users 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.

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, as shown in Figure 7-1. 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 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, you'll see the following standard set of folders:

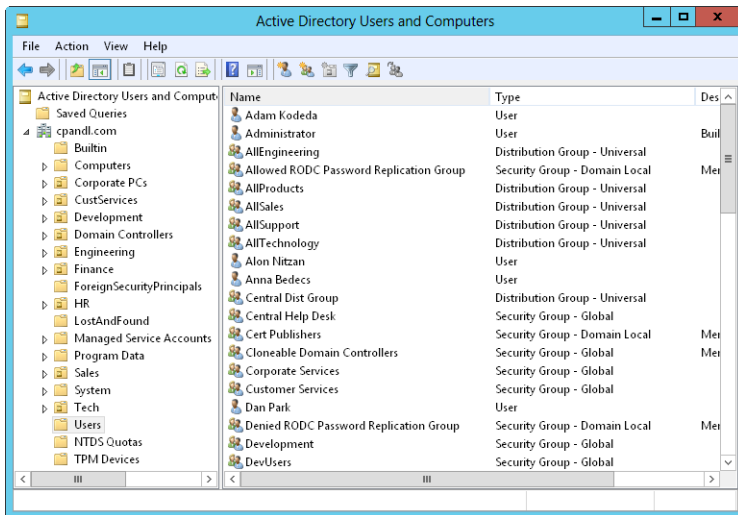
- **Builtin** The list of built-in user accounts and groups.
- **Computers** The default container for computer accounts.
- **Domain Controllers** The default container for domain controllers.

- **ForeignSecurityPrincipals** Contains information on 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, tap or click View and then select Advanced Features. You now see the following additional folders:

- **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. In Figure 7-1, there are multiple administrator-created organizational units in the cpandl.com domain. These include Corporate PCs, CustServices, Development, Engineering, and Finance.



**FIGURE 7-1** When you're working with Active Directory Users And Computers, you can access computer and user objects through the console tree.

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, press and hold or right-click Active Directory Users And Computers and then tap or click Change Domain.
2. The Change Domain dialog box displays the current (or default) domain. Type a new domain name, or tap or click Browse, select a domain in the Browse For Domain dialog box, and then tap or 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 tap or click OK. Otherwise, just tap or click OK.

If you start Active Directory Users And Computers and no objects are available, it may 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, press and hold or right-click Active Directory Users And Computers and then tap or click Change Domain Controller.  
You'll see the current domain and domain controller you're working with in the Change Directory Server dialog box.
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 tap or click OK. Otherwise, just tap or click OK.

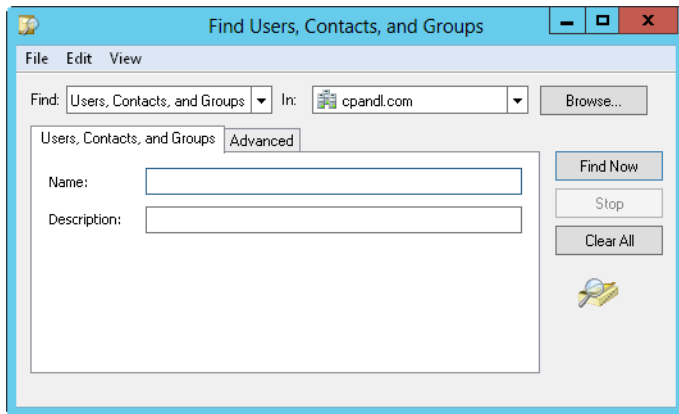
**NOTE** The Change Directory Server dialog box also shows you the site associated with domain controllers as well as the domain controller type, version, and status. If the domain controller type is listed as GC, the domain controller is also hosting a global catalog.

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.

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, press and hold or right-click the current domain or a specific container that you want to search, and then tap or click Find. This opens a Find dialog box similar to the one shown in Figure 7-2.



**FIGURE 7-2** In the Find dialog box, you can search for resources in Active Directory.

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, as well as 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** Allows you to 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, tap or click Entire Directory.
4. Enter your search parameters, and then tap or click Find Now. As shown in Figure 7-3, any matching entries are displayed in the search results. Double-tap or double-click an object to view or modify its property settings. Press and hold or 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 type the name of the object you're looking for in the Name text 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.



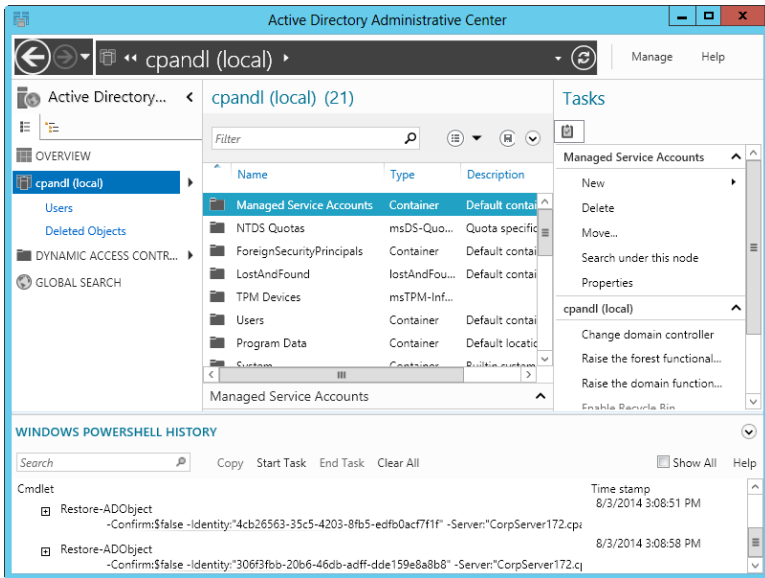
**FIGURE 7-3** Objects that match search criteria are displayed in the search results; you can manage them by pressing and holding or right-clicking their entries.

## Active Directory Administrative Center and Windows PowerShell

Active Directory Administrative Center, shown in Figure 7-4, provides a task-oriented 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
- Perform global searches of Active Directory
- Raise forest and domain functional levels
- Recover deleted objects from the Active Directory Recycle Bin

Active Directory Administrative Center is installed by default on Windows Server 2012 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 Active Directory Administrative Center.

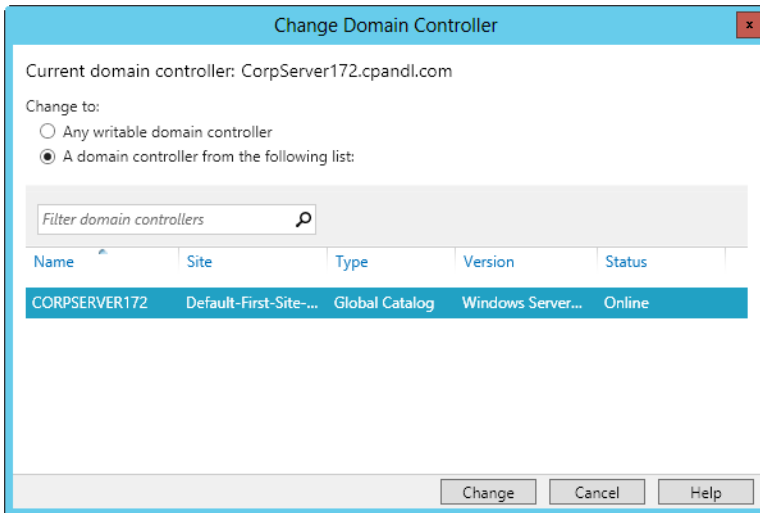


**FIGURE 7-4** Perform task-oriented management of Active Directory.

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

In the **Change Domain Controller** dialog box, you'll see the current domain and domain controller you're working with, as shown in Figure 7-5. Select a domain controller to use and then tap or click **Change**.



**FIGURE 7-5** Change the domain controller.

Like Active Directory Users And Computers, Active Directory 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 type 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 using the same filtering technique. When you select the Global Search node, you can search the names of all container-level objects as well as users, groups, computers, and so on for the currently selected container node.

With global searches, you can change the associated container node by tapping or 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, universal principal name (UPN), and group SAM account name are included.



Additionally, Active Directory Administrative Center makes use of the web services provided by Active Directory Web Services (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 with Active Directory by using the Active Directory module for Windows PowerShell. The module is automatically imported when you select the related option on the Tools menu in Server Manager. Otherwise, this module is not imported into Windows PowerShell by default, and you need to import it before you can work with any Active Directory cmdlets.

At the Windows PowerShell prompt, you can import the Active Directory module by entering **Import-Module ActiveDirectory**. Once 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. Alternatively, you can select the Active Directory Module For Windows PowerShell option on the Tools menu in Server Manager to import the module when Windows PowerShell starts.

At the Windows PowerShell prompt, you can list all available cmdlets by typing **get-command**. 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 on a specific cmdlet, type **get-help** followed by the cmdlet name. 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 Windows PowerShell prompt.

**NOTE** The Active Directory module for Windows PowerShell is installed by default on Windows Server 2012, and it's available on client computers when you install the Remote Server Administration Tools and select the related options. Windows PowerShell relies on the .NET Framework and Windows Remote Management (WinRM) to perform administrative tasks.

## Managing 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.

### Creating Computer Accounts on a Workstation or Server

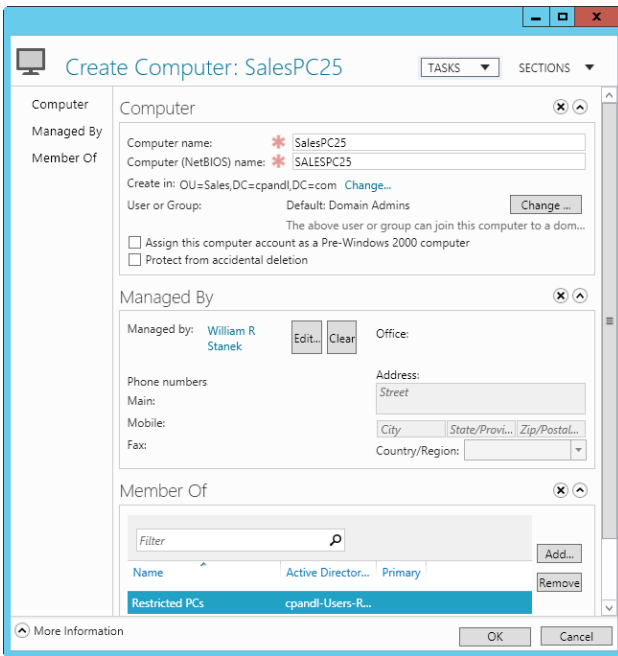
The easiest way to create a computer account is to log on to the computer you want to configure and then join a domain, as described in "Joining a Computer to a

Domain or Workgroup” later in this chapter. When you do this, the necessary computer account is created automatically and placed in the Computers folder or the Domain Controllers folder, as appropriate. 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 Computer Accounts in Active Directory 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. In the Active Directory Administrative Center console tree, press and hold or right-click the container in which you want to place the computer account, tap or click New, and then tap or click Computer. This opens the Create Computer dialog box shown in Figure 7-6.



**FIGURE 7-6** Create new computer accounts, and set their managed by and member of properties.

2. Type the computer name.
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, tap or 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.

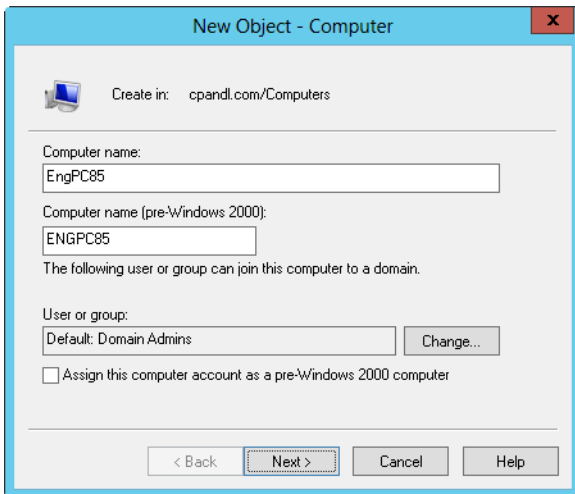
4. If this account will be used with applications written for legacy operating systems, select Assign This Computer Account As A Pre–Windows 2000 Computer.
5. Optionally, select Protect From Accidental Deletion to mark the account as protected in Active Directory. Protected accounts can be deleted only if you remove the Protect flag prior to attempting to delete the account.
6. Optionally, assign a security principal as the manager of the computer by tapping or 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.
7. The computer account is added to the appropriate default computer group automatically. Typically, this is Domain Computers. You can add the computer account to other groups by tapping or 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.
8. Tap or click OK to create the computer account.

## Creating Computer Accounts in Active Directory Users And Computers

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, press and hold or right-click the container in which you want to place the computer account, tap or click New, and then tap or click Computer. This starts the New Object—Computer Wizard shown in Figure 7-7.



**FIGURE 7-7** Create new computer accounts using the New Object—Computer Wizard.

2. Type the computer name.
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, tap or 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, tap or click OK to create the computer account. Otherwise, tap or click Next twice, and then tap or 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 to 4 in the previous procedure. Tap or click Next to display the Managed page.
2. Select the This Is A Managed Computer check box, and then type the computer's globally unique identifier/universally unique identifier (GUID/UUID). Tap or 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, tap or click Find Now to display a list of all remote installation servers in the organization. Tap or click the host server you want to use, and then tap or click OK to close the Find dialog box.

4. Tap or click Next, and then tap or 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 using the Win32\_ComputerSystemProduct class of the Windows Management Instrumentation (WMI) interface. The following example returns the 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-tap or double-click the computer account to open its Properties dialog box.
3. On the Object tab, select Protect Object From Accidental Deletion and then tap or click OK.

## Viewing and Editing Computer Account Properties

Using Active Directory Users And Computers or Active Directory Administrative Center, you can view and edit computer account properties 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-tap or 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.

## Deleting, Disabling, and Enabling 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. Press and hold or right-click the computer account, and then do one of the following:
  - Tap or click Delete to delete the account permanently. Tap or click Yes to confirm the deletion.
  - Tap or click Disable Account to temporarily disable the account, and tap or click Yes to confirm the action. A red circle with an X indicates that the account is disabled.
  - Tap or 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 it. Double-tap or double-click the account to open its Properties dialog box. Clear the Protect Object From Accidental Deletion check box and then tap or click OK. With Properties dialog boxes for Active Directory Users And Computers, this check box is on the Object tab. With 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.

## 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 press and hold or right-click the computer account in Active Directory Users And Computers and 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 press and hold or 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` as well as `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 `Test-ComputerSecureChannel` to test the secure connection between a local computer and a domain. You'll want to log on locally to the computer, open a 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. See Microsoft Knowledge Base Article 325850 for more details ([support.microsoft.com/default.aspx?scid=kb;en-us;325850](http://support.microsoft.com/default.aspx?scid=kb;en-us;325850)).

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, type **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. Type 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.

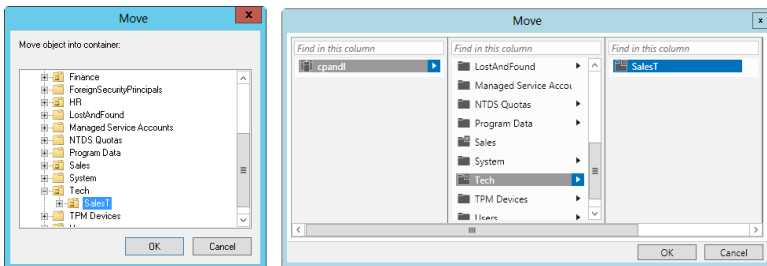
For domain controllers, you must perform additional steps. After you log on locally, you must stop the Kerberos Key Distribution Center service and set its startup type to Manual. After you restart the computer and verify that the password has been successfully reset, you can restart the Kerberos Key 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. Press and hold or right-click the computer account you want to move, and then tap or click Move. This displays the Move dialog box, shown in Figure 7-8.



**FIGURE 7-8** In the Move dialog box, you can move computer accounts to different containers.

3. In the Move dialog box, use the options provided to select the container to which you want to move the computer. Navigate to subcontainers or child OUs as necessary. Tap or click OK.



## Managing Computers

As its name indicates, you use Computer Management to manage computers. 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 pressing and holding or right-clicking the computer entry and selecting Manage from the shortcut menu. This launches Computer Management and automatically connects to the selected computer.

## 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. You might also want to refer to Chapter 14, "Managing TCP/IP Networking," for details on configuring TCP/IP connections. 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 2012 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 2012 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 Windows Vista or later, as well as Windows Server 2008 or later, computer to a domain, see "The Computer Name Tab" in Chapter 2, "Managing Servers Running Microsoft Windows Server 2012." The process is nearly identical for configuring Windows 2000 Professional, Windows 2000 Server, Windows XP Professional, and Windows Server 2003 computers as well. A key difference is that

tapping or clicking System And Security, System in Control Panel opens the System Properties dialog box directly.

If the name change is unsuccessful, you'll see a message informing you that the change was unsuccessful or a message telling you 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. Close 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, as discussed in Chapter 14.

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 while 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. Press and hold or right-click the Computers container, and then tap or click Properties.
3. On the Security tab, tap or click Advanced. In the Advanced Security Settings For Computers dialog box, tap or click Add to open the Permission Entry For Computers dialog box.
4. Tap or 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 tap or click OK. Click OK again.

## Using Offline Domain Join

Computers running editions of Windows 7 and Windows 8 designed for workplaces support offline domain join, as do servers running Windows Server 2008 R2 or later. 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 new remote domain join feature, Windows Server 2012 provides the capability for computers running Windows 8 to join a domain and receive domain settings remotely from the Internet.

You run Djoin.exe at an elevated, 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 elevated, 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 cpand1 /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 2012, add the */downlevel* command.

3. On the new computer, use Djoin.exe to import the .txt file. At an elevated, administrator command prompt, type **djoin /requestODJ /loadfile *FileName* /windowspath %SystemRoot% /localosCaution**, 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 reboot 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.

## Managing Domain Controllers, Roles, and Catalogs

---

Domain controllers perform many important tasks in Active Directory domains. Many of these tasks are discussed in Chapter 6.

### Installing and Demoting Domain Controllers

You install a domain controller by configuring Active Directory Domain Services on a server. Later, if you don't want the server to handle controller tasks, you can demote the server. It will then act as a member server again. You follow a similar procedure to install or demote servers, but before you do, you should consider the impact on the network and read "Understanding the Directory Structure" in Chapter 6.

As that section explains, when you install a domain controller, you might need to transfer operations master roles and reconfigure the global catalog structure. Also, before you can install Active Directory Domain Services, DNS must be working on the network. When you install AD DS, you can include DNS server installation, if it is needed. When you create a new domain, a DNS delegation is created automatically during the installation process and to do this requires credentials that have permissions to update the parent DNS zones.

To add the first domain controller that runs Windows Server 2012 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. For RODC prep, you must use an account that is a member of the Enterprise Admins group.

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.

**REAL WORLD** Note that with Windows Server 2012 and later, all AD DS installation and configuration tasks are performed via Server Manager. You no longer have to run an installation wizard and a separate command-line promotion task. You also might not need to manually prepare Active Directory for Windows Server 2012.

Also note that in Windows Server 2003 and later, you no longer have to demote a domain controller to rename it. You can rename a domain controller at any time. The only problem is that the server is unavailable to users during the renaming process and you might need to force a directory refresh to reestablish proper communications with the server. You can't, however, move a domain controller to a different domain. You must demote the domain controller, update the domain settings for the server and its computer account, and then promote the server to be a domain controller once more.

To install a domain controller, follow 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 using the Add Servers option. Using Server Manager for remote management requires the configuration discussed 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. To install a new Active Directory forest, you must be logged on as the local Administrator for the computer. To install a new child domain or new domain tree, you must be logged on as a member of the Enterprise Admins group. To install an additional domain controller in an existing domain, you must be logged on as a member of the Domain Admins group.
2. In Server Manager, tap or click Manage and then tap or 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 tap or click Next.
3. On the Select Installation Type page, select Role-Based Or Feature-Based Installation and then tap or click Next.
4. On the Select Destination Server page, the server pool shows servers you've added for management. Tap or click the server you are configuring and then tap or click Next.

5. On the Select Server Roles page, select Active Directory Domain Services and then tap or click Next twice. Tap or click Install. This runs the Active Directory Domain Services Installation Wizard.
6. When the initial installation task completes, you need to tap or 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 tap or click the Notifications icon and then tap or click Promote This Server To A Domain Controller.

**MORE INFO** If the installation fails, note the error and take appropriate corrective action before restarting this 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.

7. If the computer is currently a member server, the 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 2012. You need to specify whether this is a domain controller for a new domain or an additional domain controller for an existing domain. 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, and verify replication of changes to Active Directory.

To demote a domain controller, follow these steps:

1. In Server Manager, tap or click Manage and then tap or 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 tap or click Next.
2. On the Select Destination Server page, the server pool shows servers you've added for management. Tap or click the server you are configuring and then tap or click Next.
3. On the Select 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.
5. On the Credentials page, note your current logon account. If necessary, provide alternate credentials with permissions to remove the domain controller. Click Next.
6. If the Warnings page is displayed, note the warnings, select Proceed With Removal, and then click Next.

7. Enter and then confirm a new password for the server's local Administrator account. The passwords must match. Click Next.
8. On the Confirm Removal Selections page, you have the option of selecting the Restart The Destination Server Automatically If Required check box. Because a restart of the server is required to complete the removal, you might want to choose this option and then confirm it by tapping or clicking Yes. When you are ready to proceed, click Remove.

**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 select the Force The Removal Of This Domain Controller check box as part of the removal process. Here, the flexible single master operation (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.

**REAL WORLD** An alternative technique for installing domain controllers is to use backup media. This option was introduced in Windows Server 2003. To install a domain controller from backup media, create a backup of the system state data of a domain controller and restore it on a different server running Windows Server 2003 or later. 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 has thousands of entries.


## Viewing and Transferring Domainwide Roles

You can use Active Directory Users And Computers to view or change the location of domainwide operations master roles. At the domain level, you can work with roles for relative ID (RID) masters, primary domain controller (PDC) emulator masters, and infrastructure masters.

**NOTE** Operations master roles are discussed in "Understanding Operations Master Roles" in Chapter 6. You use Active Directory Domains And Trusts to set the domain naming master role and Active Directory Schema to change the schema master role. The fastest way to determine the current FSMO for all roles is to type **netdom query fsmo** at a command prompt.

You can view the current operations master roles by following these steps:

1. In Active Directory Users And Computers, press and hold or right-click Active Directory Users And Computers in the console tree. On the shortcut menu, point to All Tasks and then tap or click Operations Masters. This opens the Operations Masters dialog box, shown in Figure 7-9.



**FIGURE 7-9** In the Operations Masters dialog box, transfer operations masters to new locations or simply view their current locations.

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. Start Active Directory Users And Computers. In the console tree, press and hold or right-click Active Directory Users And Computers and then tap or click Change Domain Controller.
2. In the Change Directory Server dialog box, tap or click This Domain Controller Or AD LDS Instance, select the domain controller to which you want to transfer an operations master role, and then tap or click OK.
3. In the console tree, press and hold or right-click Active Directory Users And Computers. On the shortcut menu, point to All Tasks and then tap or click Operations Masters.
4. In the Operations Masters dialog box, tap or click the RID, PDC, or Infrastructure tab as appropriate for the type of role you want to transfer.
5. Tap or click Change to transfer the role to the previously selected domain controller. Tap or click OK.



## Viewing and Transferring the Domain Naming Master Role

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.

**TIP** If you need to connect to a different domain, connect to a domain controller following steps similar to those described in “Using Active Directory Users And Computers” earlier in this chapter. The only difference is that you press and hold or right-click Active Directory Domains And Trusts in the console tree.

To transfer the domain naming master role, follow these steps:

1. Start Active Directory Domains And Trusts. In the console tree, press and hold or right-click Active Directory Domains And Trusts and then tap or click Change Active Directory Domain Controller.
2. In the Change Directory Server dialog box, select the This Domain Controller Or AD LDS Instance option, and then select the domain controller to which you want to transfer the domain naming master role. Tap or click OK.
3. In the console tree, press and hold or right-click Active Directory Domains And Trusts and then tap or click Operations Master. This opens the Operations Master dialog box.
4. The Domain Naming Operations Master box displays the current domain naming master. Tap or click Change to transfer this role to the previously selected domain controller.
5. Tap or click Close.

## Viewing and Transferring Schema Master Roles

You use Active Directory Schema to view or change the schema master's location. Type **regsvr32 schmmgmt.dll** at an elevated, administrator command prompt to register Active Directory Schema. You can then transfer the schema master role by following these steps:

1. Add the Active Directory Schema snap-in to an MMC.
2. In the console tree, press and hold or right-click Active Directory Schema and then tap or 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. Tap or click OK. In the console tree, press and hold or right-click Active Directory Schema and then tap or click Operations Master.
5. Tap or click Change in the Change Schema Master dialog box. Tap or click OK, and then tap or click Close.

## Transferring Roles 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 typing **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 prompt. One way to do this is by pressing the Windows key, typing **cmd.exe**, and then pressing Enter.
4. At the prompt, type **ntdsutil**. This starts the Directory Services Management Tool.
5. At the ntdsutil prompt, type **roles**. This puts the utility in Operations Master Maintenance mode.
6. At the fsmo maintenance prompt, type **connections**. At the server connections prompt, type **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.adatum.com
```
7. After you've established a successful connection, type **quit** to exit the server connections prompt. At the fsmo maintenance prompt, type **transfer**, and then type the identifier for the role to transfer. The identifiers are as follows:
  - **pdcc** 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. Type **quit** at the fsmo maintenance prompt, and then type **quit** at the ntdsutil prompt.

## Seizing Roles 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, it will recognize the 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 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 2012 includes several tools for working with Active Directory replication. One tool you can use is Repadmin.

You can display the status of the last inbound replication for a domain controller 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 Server252 in the Cpandl.com domain:

```
repadmin /showrepl server252.cpandl.com dc=cpandl,dc=com
```

**NOTE** PowerShell parses commands differently than the command prompt. Normally, you can type commands at the PowerShell prompt in the same way you type them at the command prompt. Here, though, PowerShell will misinterpret `dc=cpandl,dc=com` as two separate parameters. To prevent this, enclose the value in quotes: `"dc=cpandl,dc=com"`.

To display the highest sequence number for a specified naming context on each replication partner of a designated domain controller, type 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 Server252 in the Cpandl.com domain:

```
repadmin /showutdvec server252.cpandl.com dc=cpandl,dc=com
```

The output shows the highest USN on replication partners for the default configuration partition:

```
Default-First-Site-Name\SERVER252 @ USN    45164 @ Time 2014-03-30 11:35:24
Default-First-Site-Name\SERVER147 @ USN    45414 @ Time 2014-03-30 11:42:16
```

If Server252 was the previous role owner and the domain controller you are examining has an equal or larger USN for Server252, the domain controller is up to date. However, if Server252 was the previous role owner and the domain controller you are examining has a lower USN for Server252, 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 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-ADReplicationUpToDateenessVectorTable** 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 using the following syntax:

```
get-adreplicationpartnermetadata -target corpserver98 -partition *
```

Like Repadmin /Showutdvec, Get-ADReplicationUpToDateenessVectorTable displays the highest sequence numbers for partitions being replicated and can help you troubleshoot replication issues. Here is the basic syntax:

```
Get-ADReplicationUpToDateenessVectorTable -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-adreplicationuptodateneessvectortable -target corpserver98
```

The output shows the highest USN on replication partners for the default configuration partition:

```
LastReplicationSuccess : 3/30/2014 1:45:57 PM
Partition               : DC=cpandl,DC=com
PartitionGuid           : c39cfdbd-e1a1-4c4c-9355-85d7ea05c10a
Partner                 : CN=NTDS Settings,CN=CORPSEVER172,CN=Servers,
CN=Default-First-Site-Name,CN=Sites,CN=Configuration,DC=cpandl,DC=com
PartnerInvocationId     : fb32931c-e319-473a-8069-d781f980057b
Server                  : CorpServer98.cpandl.com
UsnFilter                : 82656
```

```
LastReplicationSuccess : 3/30/2014 1:48:44 PM
Partition               : DC=cpandl,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=cpandl,DC=com
PartnerInvocationId     : d8bf2da2-b08d-4d36-bc53-1b7f62643437
Server                  : CorpServer98.cpandl.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 Cpandl.com domain by entering the following:

```
get-adreplicationfailure -Target "cpandl.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 CorpServer172
```

Follow these steps to seize a server role:

1. Type **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 window.
5. At the command prompt, type **ntdsutil**. This starts the Directory Services Management Tool.
6. At the ntdsutil prompt, type **roles**. This puts the utility in Operations Master Maintenance mode.

7. At the fsmo maintenance prompt, type **connections**. At the server connections prompt, type **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.adatum.com
```
8. After you've established a successful connection, type **quit** to exit the server connections prompt. At the fsmo maintenance prompt, type **seize**, and then type the identifier for the role to seize. The identifiers are as follows:
  - **pdcc** 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. Type **quit** at the fsmo maintenance prompt, and then type **quit** at the ntdsutl prompt.

## Configuring Global Catalogs

Global catalogs have an important role on the network. This role is discussed in “Understanding the Directory Structure” in Chapter 6. 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, press and hold or right-click NTDS Settings and then tap or 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 operating in 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 “Universal Group Membership Caching” in Chapter 6, 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, expand and then select the site you want to work with.
2. In the details pane, press and hold or right-click NTDS Site Settings, and then tap or 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. Tap or click OK.

## Managing Organizational Units

---

As discussed in Chapter 6, organizational units help you organize objects, 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, press and hold or right-click the domain node or existing organizational unit folder in which you want to add an organizational unit. Tap or click New on the shortcut menu, and then tap or click Organizational Unit.
2. Type the name of the organizational unit, and then tap or click OK.
3. You can now move accounts and shared resources to the organizational unit. See “Moving Computer Accounts” earlier in this chapter for an example.

## Viewing and Editing Organizational Unit Properties

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. Press and hold or right-click the organizational unit you want to work with, and then tap or click Properties. This displays a Properties dialog box that lets you view and edit settings.

## Renaming and Deleting Organizational Units

You can rename or delete an organizational unit by following these steps:

1. In Active Directory Users And Computers, press and hold or right-click the organizational unit folder you want to work with.
2. To delete the organizational unit, tap or click Delete. Then confirm the action by tapping or clicking Yes.
3. To rename the organizational unit, tap or click Rename. Type a new name for the organizational unit, and then press Enter.

In Active Directory Administrative Center, you delete organizational units in the same way, but to rename an organizational unit you open its Properties dialog box, type the new name, and then tap or click OK.

## 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 either Active Directory Users And Computers or Active Directory Administrative Center, you can also follow these steps to move organizational units:

1. Press and hold or right-click the organizational unit folder you want to move, and then tap or click Move.
2. In the Move dialog box, expand the domain and then select the container to which you want to move the organizational unit. Tap or click OK.

## Managing Sites

---

The Active Directory Domain Services Installation Wizard creates a default site and a default site link when you install Active Directory Domain Services on the first domain controller in a site. 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.

Configuring a site is a multipart process that includes the following steps:

1. Creating the site
2. Creating one or more subnets and associating them with the site
3. Associating a domain controller with the site



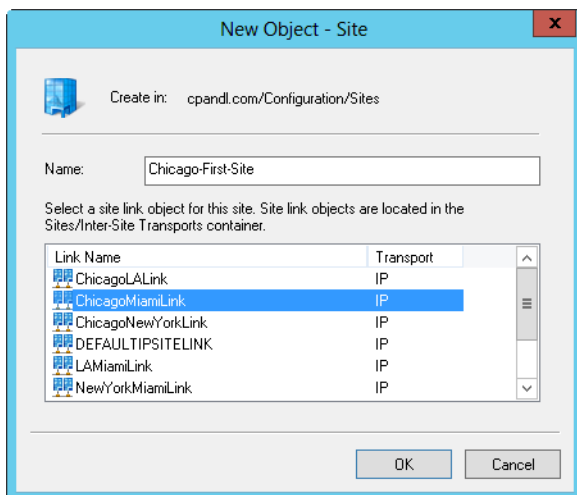
4. Linking the site to other sites using site links and, if necessary, creating site link bridges

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, press and hold or right-click the Sites container in the console root and then tap or click New Site.
2. In the New Object—Site dialog box, shown in Figure 7-10, type a name for the site, such as **Chicago-First-Site**. Site names cannot contain spaces or any special characters other than a dash.



**FIGURE 7-10** Create the site by setting the site name and a related site link.

3. Tap or 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. Tap or click OK. A prompt is displayed detailing the steps you must complete to finish site configuration. Tap or click OK again.
5. To complete site configuration, you must complete the remaining configuration tasks.

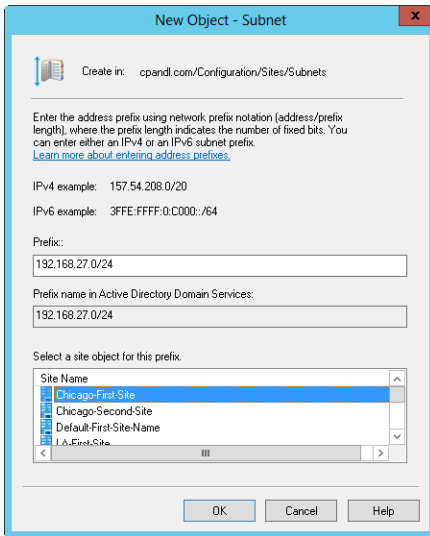
**TIP** You can rename a site at any time. In Active Directory Sites And Services, press and hold or right-click the site and then select Rename. Type the new name for the site, and then press Enter.

# 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, press and hold or right-click the Subnets container in the console tree and then tap or click New Subnet. This displays the New Object—Subnet dialog box, shown in Figure 7-11.



**FIGURE 7-11** Create the subnet by entering the network prefix and selecting an associated site.

2. In the Prefix text box, type the IPv4 or IPv6 network address prefix using the network prefix notation. In network prefix notation, you type 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 192.168.27.0 and the first 24 bits identify the network ID, you enter **192.168.27.0/24** as the network prefix notation.
3. Select the site with which the subnet should be associated, and then tap or click OK.

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

## 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 type 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.cpan1.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. Press and hold or right-click the domain controller, and then tap or click Move. In the Move Server dialog box, tap or click the site that should contain the server and then tap or click OK.

**NOTE** Don't move a domain controller to a site if it 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

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.

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.

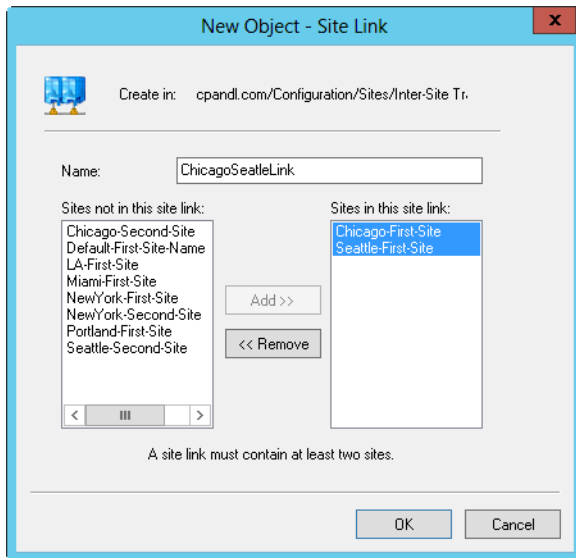
You can configure site links to use either RPC over IP or Simple Mail Transfer Protocol (SMTP) as the transport protocol. With IP as the transport, domain controllers establish an RPC over IP connection with a single replication partner at a time and replicate Active Directory changes synchronously. 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. Press and hold or right-click the container for the transport protocol you want to use (either IP or SMTP), and then tap or click New Site Link.
3. In the New Object—Site Link dialog box, shown in Figure 7-12, type a name for the site link, such as **ChicagotoSeattleLink**. Site link names cannot contain spaces or special characters other than a dash.
4. In the Sites Not In This Site Link list, tap or click the first site that should be included in the link, and then tap or 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. Tap or click OK.



**FIGURE 7-12** Create the site link by entering a name for the link and selecting the associated sites.

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, press and hold or right-click the site link in the details pane and then tap or 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.
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, tap or click Change Schedule and then set the replication schedule in the Schedule For dialog box. Tap or 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, press and hold or right-click the site link in the details pane and then tap or click Properties.
2. In the Properties dialog box, the General tab is selected by default. In the Sites Not In This Site Link list, tap or click the first site that should be included in the link, and then tap or 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, tap or click the first site that should not be included in the link, and then tap or 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. Tap or click OK.

## 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 they don't. 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. Press and hold or right-click the container for the transport protocol you want to work with (either IP or SMTP), and then tap or click Properties.
3. To enable site link transitivity, select Bridge All Site Links and then tap or 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 tap or click OK. When site link transitivity is disabled, you must configure site link bridges for the affected protocol.

Once 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. Press and hold or right-click the container for the transport protocol you want to work with (either IP or SMTP), and then tap or click New Site Link Bridge.

3. In the New Object—Site Link Bridge dialog box, type 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 tap or 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. Tap or 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, press and hold or right-click the container for the transport protocol you want to work with and then tap or 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, tap or click the first site link that should be included in the bridge, and then tap or 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, tap or click the first site link that should not be included in the bridge, and then tap or 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. Tap or click OK.

## Maintaining Active Directory

---

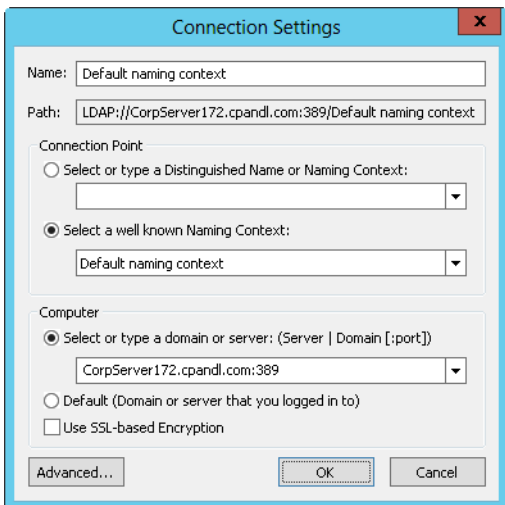
To ensure proper operations of Active Directory, you need to perform periodic monitoring and maintenance. In your monitoring and maintenance efforts, you'll find that some tools are instrumental to your success. In this section, I'll introduce these tools as well as some general maintenance tasks.

### Using ADSI Edit

When you are diagnosing problems and troubleshooting, the Active Directory administration tool you should use is 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:

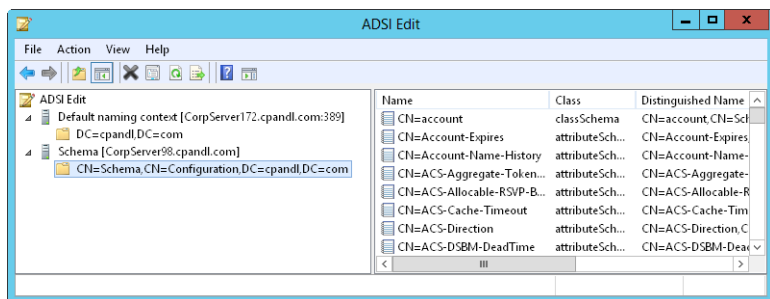
1. Press and hold or right-click the ADSI Edit node in the console tree, and then tap or click Connect To. This displays the Connection Settings dialog box, shown in Figure 7-13.



**FIGURE 7-13** Connect to a naming context in ADSI Edit.

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 tap or click OK, you are connected to any available domain controller in your logon domain. To connect to a different domain or server, select 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 FileServer252.cpandl.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. As Figure 7-14 shows, when you connect to multiple naming contexts, you have separate nodes for managing each context. For troubleshooting, you can connect to the same naming context on different servers in the same domain as well. By comparing the values associated with properties on one server with those on another, you can identify a replication problem.



**FIGURE 7-14** Navigate the naming contexts to examine related containers and properties.



## Examining Intersite Topology

The Inter-Site Topology Generator (ISTG) in 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 expand the site for the ISTG you want to locate in the console tree.
2. In the details pane, double-tap or double-click NTDS Site Settings. In the NTDS Site Settings dialog box, the current ISTG is listed in the Inter-Site Topology Generator panel.

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. Press and hold or right-click the server you want to designate as a preferred bridgehead, and then tap or 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 Inter-Site Data Transfer list and then tap or click Add. Repeat as necessary to specify both IP and SMTP. Tap or 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. Press and hold or right-click the server you want to stop using as a preferred bridgehead, and then tap or click Properties.
2. Select the transport protocol in the This Server Is A Preferred Bridgehead Server For The Following Transports list, and then tap or click Remove. Tap or click OK.

## Troubleshooting Active Directory

---

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, Domain Name System (DNS), Kerberos version 5 authentication, and Remote Procedure Call (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.

Active Directory tracks changes 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 called *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 using Repadmin. With Repadmin, most command-line parameters accept a list of the domain controllers you want to work with, called DCList. You can specify the values for DCList as follows:

- **\*** A wildcard 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** Includes 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. Table 7-2 shows some of these tasks.

**TABLE 7-2** Common Replication Tasks and Commands

TASK	COMMAND
Forcing the Knowledge Consistency Checker (KCC) to recalculate the intrasite replication topology for a specified domain controller.	<b>repadmin /kcc DCList [/async]</b>
Listing bridgehead servers that match the DCList.	<b>repadmin /bridgeheads [DCList] [/verbose]</b>
Listing calls made but not yet answered by the specified server to other servers.	<b>repadmin /showoutcalls [DCList]</b>

TASK	COMMAND
Listing domains trusted by a specified domain.	<b>repadmin /showtrust [DCList]</b>
Listing failed replication events that were detected by the KCC.	<b>repadmin /failcache [DCList]</b>
Listing connection objects for the specified domain controllers. Defaults to the local site.	<b>repadmin /showconn [DCList]</b>
Listing computers that have opened sessions with a specified domain controller.	<b>repadmin /showctx [DCList]</b>
Listing the name of the ISTG for a specified site.	<b>repadmin istg [DCList] [/verbose]</b>
Listing replication partners for each directory partition on the specified domain controller.	<b>repadmin /showrepl [DCList]</b>
Listing a summary of the replication state.	<b>repadmin /replsummary [DCList]</b>
Listing server certificates loaded on the specified domain controllers.	<b>repadmin /showcert [DCList]</b>
Listing tasks waiting in the replication queue.	<b>repadmin /queue [DCList]</b>
Listing the time between intersite replications using the ISTG Keep Alive time stamp.	<b>repadmin /latency [DCList] [/verbose]</b>

# Creating User and Group Accounts

- The Windows Server Security Model **296**
- Differences Between User and Group Accounts **300**
- Default User Accounts and Groups **306**
- Account Capabilities **309**
- Using Default Group Accounts **315**
- User Account Setup and Organization **318**
- Configuring Account Policies **322**
- Configuring User Rights Policies **327**
- Adding a User Account **330**
- Adding a Group Account **335**
- Handling Global Group Membership **338**
- Implementing Managed Accounts **341**

Managing accounts is one of your primary tasks as a Microsoft Windows administrator. Chapter 7, “Core Active Directory Administration,” discusses computer accounts. This chapter examines user and group 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 determine which actions users can perform as well as which computer systems and resources they can access.

Although you might be tempted to give users wide access, you need to balance a user’s need for job-related resources with your need to protect sensitive resources or 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 need that information have access to it.

In this chapter, you’ll learn how to manage domain accounts. Although local system accounts are discussed, they are not the primary focus. For further

discussion of configuring local system accounts, see Chapter 7, “Managing User Access and Security,” in *Microsoft Windows 8 Administration Pocket Consultant* (Microsoft Press, 2012). Keep in mind that Windows 8 adds a special type of local account called a *Microsoft account*. Microsoft accounts can be thought of as synchronized local accounts. Although Microsoft accounts are not available in domains, users can access the Windows Store using stored Windows credentials and they also can use *apps*. I use the term *apps* strictly to help distinguish between desktop apps and desktop programs. For an in-depth discussion on managing apps and managing access to the Windows Store, see Chapter 8 “Installing and Maintaining Programs” in *Microsoft Windows 8 Administration Pocket Consultant*.

## The Windows Server Security Model

---

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.

### Authentication Protocols

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 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 2012 supports many network authentication protocols. Active Directory uses Kerberos version 5 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 version 2 for authentication and session security if the server supports it. Active Directory can also use client certificates for authentication.

A key feature of the Windows Server authentication model is that it supports single sign-on, 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. Now 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, on the other hand, users must provide a user name and password every time they access a network resource.

Windows Server includes 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 as well as 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.

## 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.

## Claims-Based Access Controls

To the standard access controls, Windows Server 2012 adds Kerberos armoring, compound identities, and claims-based access controls. Kerberos with 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 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 using 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.

Once 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 it available for dynamic control. You then need to apply the policy across file servers 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 8 and Windows Server 2012 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 Windows Server 2012 domain controller, and file servers must run Windows Server 2012. By default, domain controllers are placed in the Domain Controllers 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.

## 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. Tap or click Tree View to display the tree view. Next, expand Dynamic Access Control in the left pane and then select Claim Types.
2. Use the Claim Types node to create and manage claim types. For example, right-click the Claim Types node, click New, and then select Claim Type to start creating a new claim type.
3. 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.

**NOTE** Resource properties are added as classification definition properties on file servers as well.

4. 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.
5. 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.

To complete the deployment, you need to edit the highest precedence GPO linked to the OU where you put file servers and enable central access policies. To do this, follow these steps:

1. In Group Policy Management, open the GPO for editing.
2. Navigate Computer Configuration policies to Windows Settings\Security Settings\File System.
3. Press and hold or right-click Central Access Policy, and then tap or click Manage Central Access Policies. This opens the Central Access Policies Configuration dialog box.

4. 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.

Once 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 elevated, administrator command prompt.

Servers that you want to apply dynamic controls to must have the File And Storage Services role with the File Server, Storage Services, and File Server Resource Manager role services at a minimum. 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 as well. You can speed this along by opening Windows PowerShell and entering **update-fsrmclassificationproperty-definition**. 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 using advanced security settings for the folder. On the Security tab, tap or click Advanced and then select 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 tap or 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. Tap or 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 using the advanced security settings for the folder.

## Differences Between User and Group Accounts

---

Windows Server 2012 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. You can use these objects as security principals just like user accounts. The *InetOrgPerson* object is fully enabled only when working in Windows Server 2003 or higher domain operations mode. In this mode, 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.

## 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.

## Logon Names, Passwords, and Public Certificates

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 cpandl.com domain, the full logon name is wrstaneK@cpandl.com. The pre-Windows 2000 logon name is CPANDL\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 the user `cpandl.com\users\wrstaneek`, `cpandl.com` is the DNS domain name, `users` is the container or organizational unit location, and `wrstaneek` 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 using a smart card and a smart card reader.

## Security Identifiers and User Accounts

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.

## 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 `cpandl\gmarketing` for the *Gmarketing* group in the `cpandl` 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 `cpandl.com\users\gmarketing`, `cpandl.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.

## Group Types

Windows Server supports three types of groups:

- **Local groups** 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** Groups that can have security descriptors associated with them. You define security groups in domains by using Active Directory Users And Computers.
- **Distribution groups** 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 distribution groups. Distribution groups are only for email distribution and are not for assigning or managing access.

## Group Scope

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** 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** 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** 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** 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. For more information, see "When to Use Domain Local, Global, and Universal Groups" later in this chapter.

When you work with groups, the group's scope restricts what you can and cannot do. Table 8-1 offers a quick summary of these items. For complete details on creating groups, see "Adding a Group Account" later in this chapter.

**TABLE 8-1** How Group Scope Affects Group Capabilities

GROUP CAPABILITY	DOMAIN LOCAL SCOPE	GLOBAL SCOPE	UNIVERSAL SCOPE
Members	Accounts, global groups, and universal groups from any domain; domain local groups from the same domain only	Accounts and global groups from the same domain only	Accounts from any domain, as well as global and universal groups from any domain
Member of	Can be put into other domain local groups and assigned permissions only in the same domain	Can be put into other groups and assigned permissions in any domain	Can be put into other groups and assigned permissions in any domain
Scope conversion	Can be converted to universal scope provided that it doesn't have as its member another group having domain local scope	Can be converted to universal scope provided that it's not a member of any other group having global scope	Can't be converted to any other group scope

### Security Identifiers and Group Accounts

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.

## When to Use Domain Local, Global, and Universal Groups

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** Groups with domain local scope have the smallest extent. Use groups with domain local scope to help you manage access to resources such as printers and shared folders.
- **Global groups** 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** Groups with universal scope have the largest extent. Use groups with universal scope to consolidate groups that span domains. Normally, 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. Say 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. 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 called `GMarketing` and add the members of the Seattle marketing department to it. In the Chicago domain, create a group called `GMarketing` and add the members of the Chicago marketing department to it. In the NY domain, create a group called `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. Call the printer group `LocalMarketingPrinter`. Call 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. Call 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.

## Default User Accounts and Groups

---

When you install Windows Server 2012, 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.



## 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** LocalSystem is a pseudoaccount for running system processes and handling system-level 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 have the privilege to interact with the desktop as well. Services that need alternative privileges or logon rights run under the LocalService or NetworkService account.
- **LocalService** LocalService is 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** NetworkService is 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 using the credentials of the computer account.

When you install add-ons or other applications on a server, other default accounts might be installed.

## 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.

### The Administrator Account

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 once 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.

## The Guest Account

The Guest account 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 to change the password regularly. As with the Administrator account, you might want to rename the account as an added security precaution.

## Built-in and Predefined Groups

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 and Special Identities

In Windows NT, implicit groups were assigned implicitly during logon and were based on how a user accessed a network resource. For example, if a user accessed a resource through an interactive logon, the user was automatically a member of the implicit group called Interactive. In Windows 2000 and later releases, the object-based approach to the directory structure has changed the original rules for implicit groups. Although you still 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 with other default groups, the availability of a specific implicit group depends on the current configuration. Implicit groups are discussed later in this chapter.

## Account Capabilities

---

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 capabilities to an account:

- **Privileges** A type of user right that grants permissions to perform specific administrative tasks. You can assign privileges to both user and group accounts. An example of a privilege is the ability to shut down the system.
- **Logon rights** A type of user right that 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.
- **Built-in capabilities** A type of user right that is 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** A type of user right that 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. Access permissions are discussed in Chapter 12, "Data Sharing, Security, and Auditing."

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.

## Privileges

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.

Table 8-2 provides a brief summary of each privilege you can assign to users and groups. To learn how to assign privileges, see "Configuring User Rights Policies" later in this chapter.

**TABLE 8-2** Windows Server 2012 Privileges for Users and Groups

PRIVILEGE	DESCRIPTION
Act As Part Of The Operating System	Allows a process to authenticate as any user and gain access to resources as any user. Processes that require this privilege should use the LocalSystem account, which already has this privilege.
Add Workstations To Domain	Allows users to add computers to the domain.
Adjust Memory Quotas For A Process	Allows users to adjust process-based memory usage quotas.
Back Up Files And Directories	Allows users to back up the system regardless of the permissions set on files and directories.
Bypass Traverse Checking	Allows users to pass through directories while navigating an object path regardless of permissions set on the directories. The privilege doesn't allow the user to list directory contents.
Change The System Time	Allows users to set the time for the system clock.
Change The Time Zone	Allows users to set the time zone for the system clock. All users have this privilege by default.
Create A Pagefile	Allows users to create and change the paging file size for virtual memory.

PRIVILEGE	DESCRIPTION
Create A Token Object	Allows processes to create token objects that can be used to gain access to local resources. Processes that require this privilege should use the LocalSystem account, which already has this privilege.
Create Global Objects	Allows processes to create global objects. LocalService and NetworkService have the privilege by default.
Create Permanent Shared Objects	Allows processes to create directory objects in the object manager. Most components already have this privilege; it's not necessary to specifically assign it.
Create Symbolic Links	Allows an application that a user is running to create symbolic links. Symbolic links make it appear as though a document or folder is in a specific location when it actually resides in another location. Use of symbolic links is restricted by default to enhance security.
Debug Programs	Allows users to perform debugging.
Enable Computer And User Accounts To Be Trusted For Delegation	Allows computers and users to change or apply the trusted-for-delegation setting, provided they have write access to the object.
Force Shutdown From A Remote System	Allows users to shut down a computer from a remote location on the network.
Generate Security Audits	Allows processes to make security log entries for auditing object access.
Impersonate A Client After Authentication	Allows Web applications to act as clients during the processing of requests. Services and users can also act as clients.
Increase A Process Working Set	Allows an application that a user is running to increase the memory that the related process working set uses. A <i>process working set</i> is the set of memory pages currently visible to a process in physical memory. Allowing for increases in memory pages reduces page faults and enhances performance.
Increase Scheduling Priority	Allows processes to increase the scheduling priority assigned to another process, provided that they have write access to the process.
Load And Unload Device Drivers	Allows users to install and uninstall Plug and Play device drivers. This doesn't affect device drivers that aren't Plug and Play, which can be installed only by administrators.

PRIVILEGE	DESCRIPTION
Lock Pages In Memory	Allows processes to keep data in physical memory, preventing the system from paging data to virtual memory on disk.
Manage Auditing And Security Log	Allows users to specify auditing options and access the security log. You must turn on auditing in the group policy first.
Modify An Object Label	Allows a user process to modify the integrity label of objects, such as files, registry keys, or processes owned by other users. This privilege can be used to lower the priority of other processes. Processes running under a user account can modify the label of any object the user owns without requiring this privilege.
Modify Firmware Environment Values	Allows users and processes to modify system environment variables.
Perform Volume Maintenance Tasks	Allows for the administration of removable storage, the disk defragmenter, and disk management.
Profile A Single Process	Allows users to monitor the performance of nonsystem processes.
Profile System Performance	Allows users to monitor the performance of system processes.
Remove Computer From Docking Station	Allows a laptop to be undocked and removed from the network.
Replace A Process Level Token	Allows processes to replace the default token for subprocesses.
Restore Files And Directories	Allows users to restore backed-up files and directories, regardless of the permissions set on files and directories.
Shut Down The System	Allows users to shut down the local computer.
Synchronize Directory Service Data	Allows users to synchronize directory service data on domain controllers.
Take Ownership Of Files Or Other Objects	Allows users to take ownership of files and any other Active Directory objects.

## Logon Rights

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.

Table 8-3 provides a brief summary of each logon right you can assign to users and groups. Assigning logon rights is covered in “Configuring User Rights Policies.”

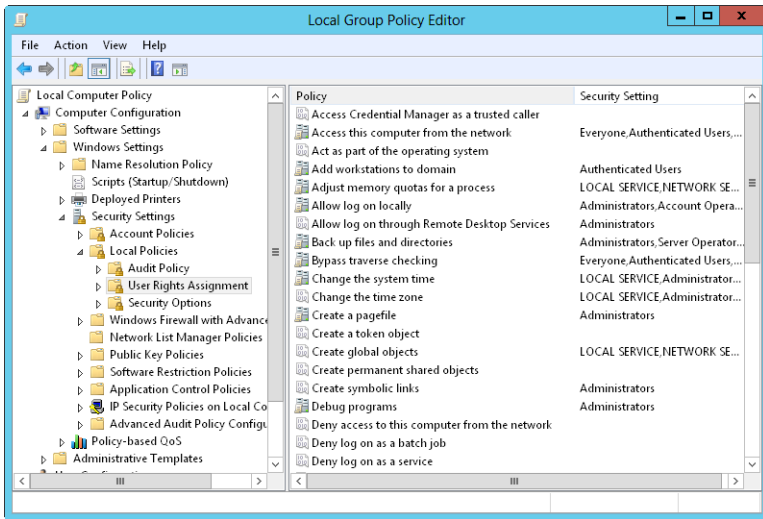
**TABLE 8-3** Windows Server 2012 Logon Rights for Users and Groups

LOGON RIGHT	DESCRIPTION
Access Credential Manager As A Trusted Caller	Grants permission to establish a trusted connection to Credential Manager. Credentials, such as a user name and password or smart card, provide identification and proof of identification.
Access This Computer From The Network	Grants remote access to the computer.
Allow Log On Locally	Grants permission to log on at the computer’s keyboard. On domain controllers, this right is restricted by default and only members of the following groups can log on locally: Administrators, Account Operators, Backup Operators, Print Operators, and Server Operators.
Allow Log On Through Remote Desktop Services	Grants access through Remote Desktop Services. This right is necessary for remote assistance and remote desktop use.
Deny Access To This Computer From The Network	Denies remote access to the computer through network services.
Deny Logon As Batch Job	Denies the right to log on through a batch job or script.
Deny Logon As Service	Denies the right to log on as a service.
Deny Logon Locally	Denies the right to log on by using the computer’s keyboard.
Deny Logon Through Remote Desktop Services	Denies the right to log on through Remote Desktop Services.
Log On As A Batch Job	Grants permission to log on as a batch job or script.
Log On As A Service	Grants permission to log on as a service. The LocalSystem account has this right. A service that runs under a separate account should be assigned this right.

## Built-in Capabilities for Groups in Active Directory

The built-in capabilities that are assigned to groups in Active Directory depend on a computer’s configuration. Using the Local Group Policy Editor, shown in Figure 8-1, you can view the capabilities that have been assigned to each group by expanding

Computer Configuration\Windows Settings\Security Settings\Local Policies and then selecting the User Rights Assignment node.



**FIGURE 8-1** View the built-in capabilities that are used with groups.

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.

Table 8-4 summarizes capabilities you can delegate to other users and groups. As you study the table, note that restricted accounts include the Administrator user account, the user accounts of administrators, and the group accounts for Administrators, Server Operators, Account Operators, Backup Operators, and Print Operators. Because these accounts are restricted, Account Operators can’t create or modify them.

**TABLE 8-4** Other Capabilities for Built-in and Local Groups

TASK	DESCRIPTION	GROUP NORMALLY ASSIGNED
Assign User Rights	Allows users to assign user rights to other users	Administrators
Create And Delete Groups	Allows users to create new groups and delete existing groups	Administrators, Account Operators
Create And Delete Printers	Allows users to create and delete printers	Administrators, Server Operators, Printer Operators



<b>TASK</b>	<b>DESCRIPTION</b>	<b>GROUP NORMALLY ASSIGNED</b>
Create, Delete, And Manage User Accounts	Allows users to administer domain user accounts	Administrators, Account Operators
Manage Group Policy Links	Allows users to apply existing group policies to sites, domains, and organizational units for which they have write access to the related objects	Administrators
Manage Network Configuration	Allows users to configure networking	Administrators, Network Configuration Operators
Manage Performance Logs	Allows users to configure performance logging	Administrators, Performance Log Users
Manage Printers	Allows users to modify printer settings and manage print queues	Administrators, Server Operators, Printer Operators
Modify The Membership Of A Group	Allows users to add and remove users from domain groups	Administrators, Account Operators
Monitor Performance Logs	Allows users to monitor performance logging	Administrators, Performance Monitor Users
Perform Cryptographic Operations	Allows users to manage cryptographic options	Administrators, Cryptographic Operators
Read All User Information	Allows users to view user account information	Administrators, Server Operators, Account Operators
Read Event Logs	Allows users to read event logs	Administrators, Event Log Readers
Reset Passwords On User Accounts	Allows users to reset passwords on user accounts	Administrators, Account Operators

## Using 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 2012 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.

## Groups Used by Administrators

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. Table 8-5 compares the administrator groups.

**TABLE 8-5** Administrator Groups Overview

ADMINISTRATOR GROUP TYPE	NETWORK ENVIRONMENT	GROUP SCOPE	MEMBERSHIP
Administrators	Active Directory domains	Domain local	Administrator, Domain Admins, Enterprise Admins
Administrators	Workgroups, computers not part of a domain	Local	Administrator
Domain Admins	Active Directory domains	Global	Administrator
Enterprise Admins	Active Directory domains	Global or Universal	Administrator
Schema Admins	Active Directory domains	Universal	Administrator

**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.

## Implicit Groups and Identities

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:

- **The Anonymous Logon identity** 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 the corporate presence servers.
- **The Authenticated Users identity** 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.
- **The Batch identity** 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.
- **The Creator Group identity** 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.
- **The Creator Owner identity** 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.

- **The Dial-Up identity** Any user accessing the system through a dial-up connection has the Dial-Up identity. This identity distinguishes dial-up users from other types of authenticated users.
- **The Enterprise Domain Controllers identity** Domain controllers with enterprisewide roles and responsibilities have the Enterprise Domain Controllers identity. This identity allows them to perform certain tasks in the enterprise using transitive trusts.
- **The Everyone identity** 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.
- **The Interactive identity** Any user logged on to the local system has the Interactive identity. This identity allows only local users to access a resource.
- **The Network identity** Any user accessing the system through a network has the Network identity. This identity allows only remote users to access a resource.
- **The Proxy identity** Users and computers accessing resources through a proxy have the Proxy identity. This identity is used when proxies are implemented on the network.
- **The Remote Desktop Services User identity** 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.
- **The Restricted identity** Users and computers with restricted capabilities have the Restricted identity.
- **The Self identity** The Self identity refers to the object itself and allows the object to modify itself.
- **The Service identity** Any service accessing the system has the Service identity. This identity grants access to processes being run by Windows Server services.
- **The System identity** 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.

## User Account Setup and Organization

---

A key part of your job as an administrator is to create accounts, and this chapter shows you how. User and group accounts allow Windows Server 2012 to track and manage information about users, including permissions and privileges. To create user accounts, you primarily use the following two account administration tools:

- Active Directory Users And Computers, which is designed to administer accounts throughout an Active Directory Domain Services 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.

## Account Naming Policies

A key policy you need to set is the naming scheme for accounts. 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. Logon names are discussed briefly in "Logon Names, Passwords, and Public Certificates" earlier in this chapter.

### Rules for Display Names

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.

### Rules for Logon Names

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 pre-Windows 2000 logon name is given to all accounts. By default, this logon name is set to the first 20 characters of the Windows logon name. The pre-Windows 2000 logon name must be unique throughout a domain.
- Users logging on to the domain using a computer that runs Windows 2000 or a later release can use their standard logon names or their pre-Windows 2000 logon names, regardless of the domain operations mode.
- 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.

**NOTE** Although Windows Server stores user names in the case that you enter, 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.

## Naming Schemes

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

**SECURITY ALERT** In environments with strict security, you can assign a numeric code for the logon name. This numeric code should be at least 20 characters. Combine this strict naming method with smart cards and smart card readers to allow users to quickly log on to the domain without having to type in all those characters. Don't worry, users can still have a display name that humans can read.

## 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 Windows NT 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.

But simply having a password isn't enough. 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 might also 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.

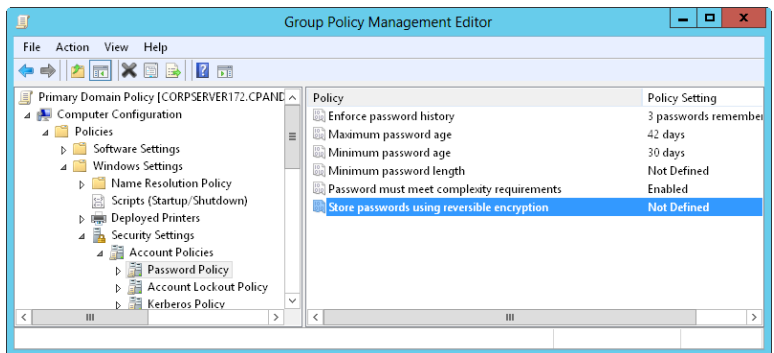
## Setting Account Policies

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 "Managing Local Group Policies" in Chapter 4, "Automating Administrative Tasks, Policies, and Procedures." You manage global group policies as explained in "Managing Site, Domain, and Organizational Unit Policies," also in Chapter 4.

Account policies should be configured in the highest precedence GPO linked to a domain. By default, the highest precedence GPO linked to a domain is the Default Domain Policy GPO. Once you access the Default Domain Policy GPO or other appropriate GPO, you can set account policies by following these steps:

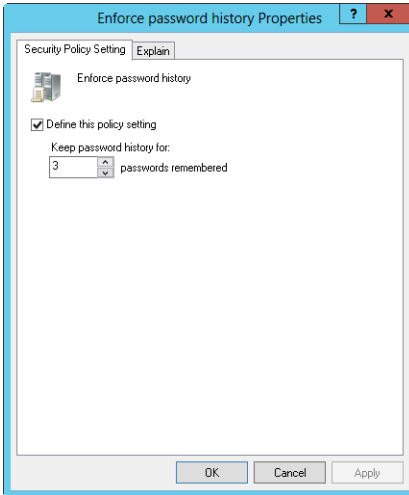
1. In the Group Policy Management Editor, shown in Figure 8-2, open the Account Policies node by expanding Computer Configuration, Windows Settings, and Security Settings. The console tree shows the name of the computer or domain you are configuring. Be sure that this is the appropriate network resource 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.



**FIGURE 8-2** Use the Account Policies node to set policies for passwords and general account use.

2. You can now manage account policies through the Password Policy, Account Lockout Policy, and Kerberos Policy nodes. To configure a policy, double-tap or double-click its entry, or press and hold or right-click it and then tap or click Properties. This opens a Properties dialog box for the policy, shown in Figure 8-3.



**FIGURE 8-3** Define and configure global group policies in the Properties dialog box.

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 standalone 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. Tapping or clicking Enabled turns on the policy restriction. Tapping or 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 following sections: "Configuring Password Policies," "Configuring Account Lockout Policies," and "Configuring Kerberos Policies."

## Configuring Account Policies

---

As you learned in the previous section, there are three types of account policies: password policies, account lockout policies, and Kerberos policies. The sections that follow show you how to configure each of these policies.



## Configuring Password Policies

Password policies, listed here, control security for passwords:

- Enforce Password History
- Maximum Password Age
- Minimum Password Age
- Minimum Password Length
- Passwords Must Meet Complexity Requirements
- Store Password Using Reversible Encryption For All Users In The Domain

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 using the Passwords Remembered box. Windows Server then tracks old passwords 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 to 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 less than 30 days away, users see a warning when they log on that they have to change their 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.

## Passwords 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 Password Using Reversible Encryption For All Users

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. Tap or 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 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. Normally, 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 has this right as well. 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 Group Policy Object (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.

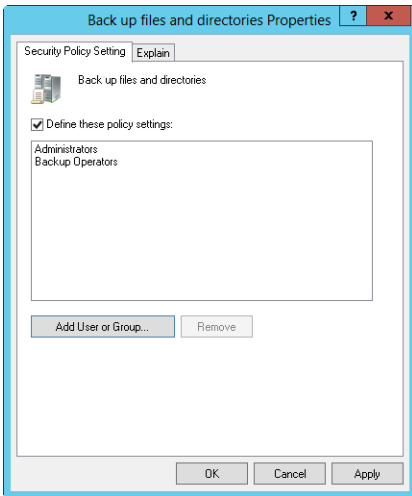
To administer user rights policies, follow these steps:

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 Computer Configuration, Windows Settings, Security Settings, and Local Policies.
2. Select User Rights Assignment to manage user rights. To configure a user rights assignment, double-tap or double-click a user right, or press and hold or right-click it and then tap or click Properties. This opens a Properties dialog box.
3. You can now configure the user rights. To configure local user rights, follow steps 1–3 in “Configuring Local User Rights.” To configure global user rights, follow steps 1–6 in the following section.

## Configuring Global User Rights

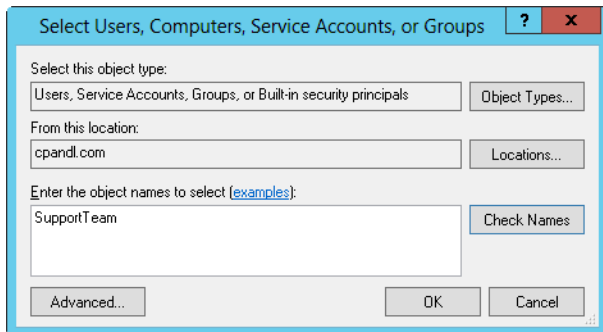
For a site, a domain, or an organizational unit, you configure individual user rights by following these steps:

1. Open the Properties dialog box for the user right, which is similar to the one shown in Figure 8-4. If the policy isn't defined, select Define These Policy Settings.



**FIGURE 8-4** In the Properties dialog box, define the user right, and then apply the right to users and groups.

2. To apply the right to a user or group, tap or click Add User Or Group. Then, in the Add User Or Group dialog box, tap or click Browse. This opens the Select Users, Computers, Service Accounts, Or Groups dialog box, shown in Figure 8-5.



**FIGURE 8-5** In the Select Users, Computers, Service Accounts, Or Groups dialog box, apply the user right to users and groups.

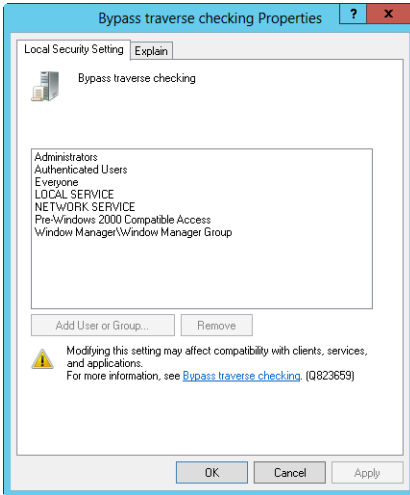
**SECURITY ALERT** Windows Firewall running on a domain controller might prevent you from using the Select Users, Computers, Service Accounts, Or Groups dialog box. This can occur when you aren't logged on locally to the domain controller and are working remotely. You might need to configure an exception on the domain controller for incoming TCP port 445. You can do this by expanding Computer Configuration\Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile. In the details pane, double-tap or double-click the Windows Firewall: Allow Inbound Remote Administration Exception policy, and then select Enabled. Alternatively, you can configure an exception by typing the following at a command prompt on the remote computer: **netsh firewall set portopening tcp 445 smb enable**. See Microsoft Knowledge Base Article 840634 ([support.microsoft.com/default.aspx?scid=kb;en-us;840634](http://support.microsoft.com/default.aspx?scid=kb;en-us;840634)) for more details.

3. Type the name of the user or group you want to use in the text box provided, and then tap or click Check Names. By default, the search is configured to find built-in security principals and user accounts. To add groups to the search, tap or click Object Types, select Groups in the list box, and then tap or click OK.
4. After you select the account names or groups to add, tap or click OK. The Add User Or Group dialog box should now show the selected accounts. Tap or 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 tapping or clicking Remove.
6. When you have finished granting the right to users and groups, tap or click OK.

# Configuring Local User Rights

For local computers, apply user rights by following these steps:

1. Open the Properties dialog box for the user right, which is similar to the one shown in Figure 8-6. Remember that site, domain, and organizational unit policies have precedence over local policies.



**FIGURE 8-6** In the Properties dialog box, define the user right and then apply the right to users and groups. If you can't edit local user rights, you might be working with a domain controller.

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 tap or click Remove.
3. You can apply the user right to additional users and groups by tapping or clicking Add User Or Group. This opens the Select Users, Computers, Service Accounts, Or Groups dialog box shown previously in Figure 8-5. You can now add users and groups.

## Adding a User Account

You need to create a user account for each user who wants to use your network resources. You create domain user accounts with Active Directory Users And Computers. You create local user accounts with Local Users And Groups.

## Creating Domain User Accounts

Generally, you can create new domain accounts in two ways:

- **Create a completely new user account** Press and hold or right-click the container in which you want to place the user account, tap or click New, and then tap or click User. This opens the New Object—User Wizard, shown in



Figure 8-7. When you create a new account, the default system settings are used.

The screenshot shows the 'New Object - User' wizard window. The window title is 'New Object - User' and it has a close button (X) in the top right corner. The window contains a user icon and the text 'Create in: cpandl.com/Users'. Below this, there are several text input fields and a dropdown menu. The 'First name:' field contains 'William', the 'Initials:' field contains 'R', and the 'Last name:' field contains 'Stanek'. The 'Full name:' field contains 'William R. Stanek'. The 'User logon name:' field contains 'wrstanek' and the domain dropdown is '@cpandl.com'. The 'User logon name (pre-Windows 2000):' field contains 'CPANDL\wrstanek'. At the bottom of the window are three buttons: '< Back', 'Next >', and 'Cancel'.

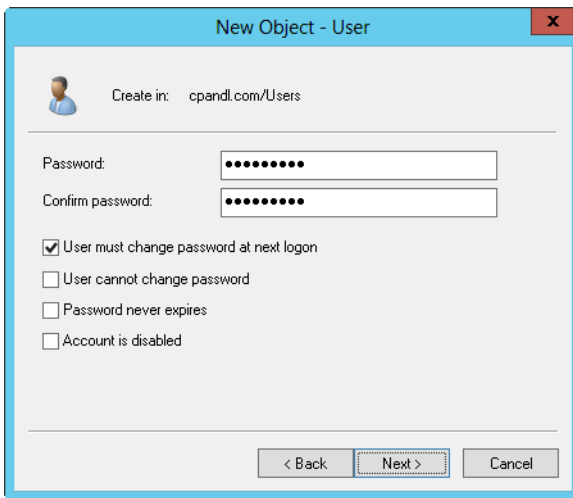
**FIGURE 8-7** Configure the user display and logon names.

- **Base the new account on an existing account** Press and hold or right-click the user account you want to copy in Active Directory Users And Computers, and then tap or click Copy. This starts the Copy Object—User Wizard, which is essentially the same as the New Object—User Wizard. However, when you create a copy of an account, the new account gets most of its environment settings from the existing account. For more information on copying accounts, see “Copying Domain User Accounts” in Chapter 9, “Managing User and Group Accounts.”

With either the New Object—User Wizard or the Copy Object—User Wizard, you can create an account by following these steps:

1. As shown in Figure 8-7, the first wizard page lets you configure the user display name and logon name. Type the user’s first name, middle initial, and last name in the text boxes provided. These boxes are used to create the full name, which is the user’s display name.
2. Make changes to the Full Name field as necessary. For example, you might want to type 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.
3. In the User Logon Name box, type the user’s logon name. Use the drop-down list to select the domain to associate the account with. This sets the fully qualified logon name.
4. The first 20 characters of the logon name are used to set the pre-Windows 2000 logon name. This logon name must be unique in the domain. If necessary, change the pre-Windows 2000 logon name.

5. Tap or click Next, and then configure the user's password on the page shown in Figure 8-8.



**FIGURE 8-8** Use the New Object—User Wizard to configure the user's password.

The options for this page are as follows:

- **Password** The password for the account. This password should follow the conventions of your password policy.
  - **Confirm Password** A text 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.
6. Tap or click Next, and then tap or click Finish to create the account. If you have problems creating the account, you'll see a warning, and you need to use the Back button to retype information in the user name and password pages as necessary.

After you create the account, you can set advanced properties for the account as discussed later in this chapter.

You also can create user accounts using Active Directory Administrative Center. To do this, follow these steps:

1. In the Active Directory Administrative Center console tree, press and hold or right-click the container in which you want to place the user account, tap or click New in the container pane, and then tap or click User. This opens the Create User dialog box shown in Figure 8-9.

**FIGURE 8-9** Creating a new user account in Active Directory Administrative Center

2. Type the user's first name, middle initial, and last name in the text boxes provided. These text boxes are used to create the full name, which is the user's display name.
3. Make changes to the Full Name box as necessary. The full name must be unique in the domain and must have 64 or fewer characters.
4. In the User UPN Logon box, type the user's logon name. Use the drop-down list to select the domain to associate the account with. This sets the fully qualified logon name.
5. The first 20 characters of the logon name are used to set the User SamAccountName Logon box. This is the user's pre-Windows 2000 logon name, which must be unique in the domain.
6. All other text boxes in the dialog box are optional. Set and confirm the user's password, if desired. Optionally, select Protect From Accidental Deletion to mark the account as protected in Active Directory. Protected accounts can be

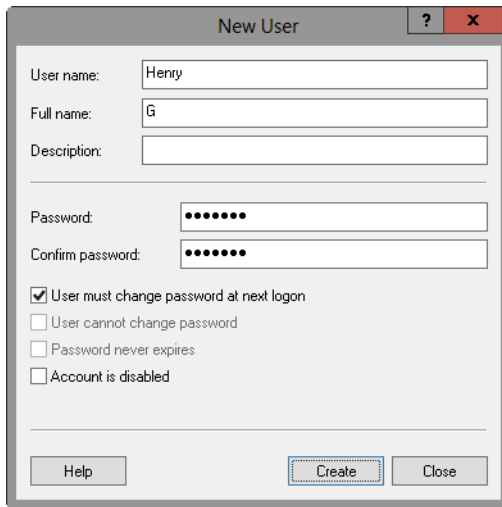
deleted only if you remove the Protect flag prior to attempting to delete the account.

7. Tap or click OK to create the user account.

## Creating Local User Accounts

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, tap or click Tools and then tap or click Computer Management. Alternatively, you can press Windows+X and then click Computer Management.
2. Press and hold or right-click the Computer Management entry in the console tree, and then tap or 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. Press and hold or right-click Users, and then tap or click New User. This opens the New User dialog box, shown in Figure 8-10. You use each of the text 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. Normally, you type the user's job title, such as Webmaster. You could also type 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.



**FIGURE 8-10** Configuring a local user account is different from configuring a domain user account.

5. Tap or click Create when you have finished configuring the new account.

## Adding a Group Account

---

You use group accounts to manage privileges for multiple users. You create global group accounts in Active Directory Users And Computers. You create local group accounts in Local Users And Groups.

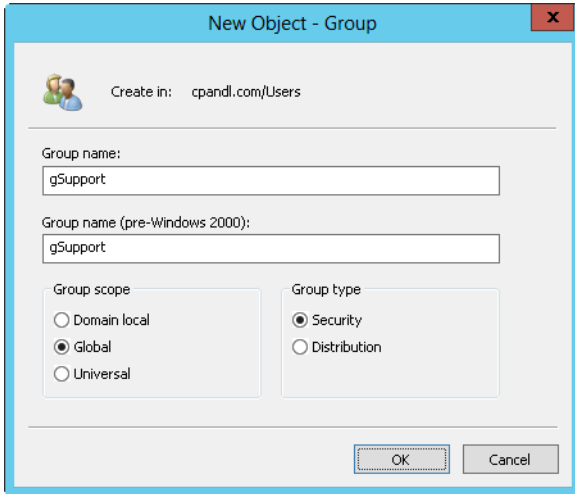
As you set out to create group accounts, remember that you create group accounts for similar types of users. The types of groups you might want to create include the following:

- **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. By creating groups based on roles within the organization, you can ensure that proper access is given to the users who need it.

# Creating a Global Group

To create a global group, follow these steps:

1. Start Active Directory Users And Computers. Press and hold or right-click the container in which you want to place the user account, tap or click New, and then tap or click Group. This opens the New Object—Group dialog box, shown in Figure 8-11.



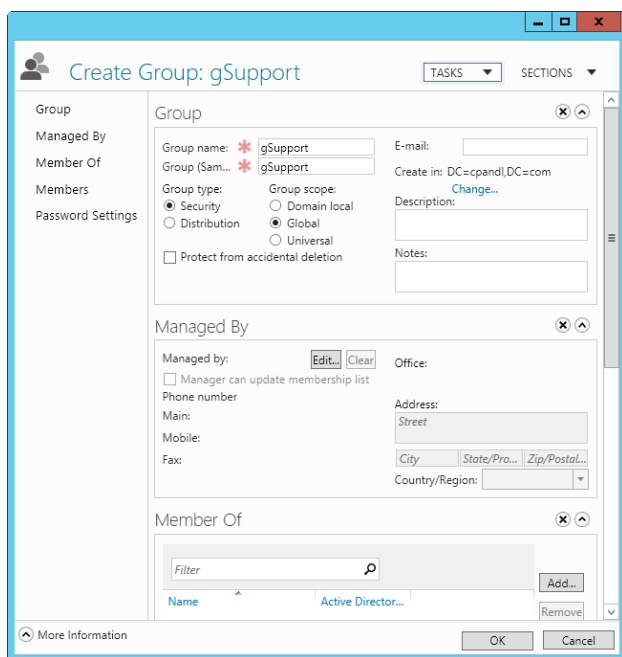
**FIGURE 8-11** The New Object—Group dialog box allows you to add a new global group to the domain.

2. Type a name for the group. Global group account names follow the same naming rules as display names for user accounts. They aren't case sensitive and can be up to 64 characters.
3. The first 20 characters of the group name are used to set the pre-Windows 2000 group name. This group name must be unique in the domain. If necessary, change the pre-Windows 2000 group name.
4. Select a group scope (Domain Local, Global, or Universal).
5. Select a group type (either Security or Distribution).
6. Tap or click OK to create the group. After you create the account, you can add members and set additional properties, as discussed later in this chapter.

You also can create groups using Active Directory Administrative Center. To do this, follow these steps:

1. In the Active Directory Administrative Center console tree, press and hold or right-click the container in which you want to place the group, tap or click New in the container pane, and then tap or click Group. This opens the Create Group dialog box shown in Figure 8-12.

2. Type a name for the group. Global group account names follow the same naming rules as display names for user accounts. They aren't case sensitive and can be up to 64 characters.
3. The first 20 characters of the group name are used to set the group SAM-AccountName group name. This is the pre-Windows 2000 group name, which must be unique in the domain.
4. Select a group type (either Security or Distribution).
5. Select a group scope (Domain Local, Global, or Universal).
6. All other selections in the dialog box are optional. 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. Tap or click OK to create the group.



**FIGURE 8-12** Creating a new group in Active Directory Administrative Center.


## Creating a Local Group and Assigning Members

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, tap or click Tools and then tap or click Computer Management. Press and hold or right-click the Computer Management entry in the console tree, and then tap or 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. Press and hold or right-click Groups, and then tap or click New Group. This opens the New Group dialog box, shown in Figure 8-13.



**FIGURE 8-13** In the New Group dialog box, you can add a new local group to a computer.

4. After you type a name and description of the group, tap or click the Add button to add names to the group. This opens the Select Users dialog box.
5. In the Select Users dialog box, type the name of a user you want to use in the Name box and then tap or click Check Names. If matches are found, select the account you want to use and then tap or click OK. If no matches are found, update the name you entered and try searching again. Repeat this step as necessary, and then tap or click OK.
6. The New Group dialog box is updated to reflect your selections. If you made a mistake, select a name and tap or click Remove.
7. Tap or click Create when you've finished adding or removing group members.

## Handling Global 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 membership
- Manage multiple memberships
- Set primary group membership for individual users and computers

## Managing Individual Membership

You can quickly add a user or a group to one or more groups by pressing and holding or right-clicking the account and 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-tap or 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 the user is currently a member of are listed. Tap or 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 tap or click Remove.
4. Tap or click OK.

If you're working exclusively with user accounts, you can add 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 tap or click the left mouse button on 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. Press and hold or right-click one of the selections, and then tap or 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. Tap or click OK.

## Managing Multiple Memberships in a Group

Another way to manage group membership is to use a group's Properties dialog box to add or remove multiple accounts. To do this, follow these steps:

1. Double-tap or 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, tap or 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 tap or click Remove.
4. Tap or click OK.

## Setting the Primary Group for Users and Computers

Users who access Windows Server through Services for Macintosh 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** Windows Server 2008 and later do not include Services for Macintosh. Services for Macintosh is only included with earlier releases of Windows Server. All user and computer accounts must have a primary group whether or not 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-tap or 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. Tap or 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 tap or click Set Primary Group.
2. In the Member Of list, tap or click the former primary group and then tap or click Remove. The group membership is now revoked.

## Implementing Managed Accounts

---

Microsoft Exchange Server, Internet Information Services, 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.

Windows 7 and all later releases of Windows support two additional types of accounts:

- Managed service accounts
- Managed virtual 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.

You can manage these accounts using the Active Directory module for Windows PowerShell. Because the Active Directory module is not imported into Windows PowerShell by default, you need to import the module before you can use the cmdlets it provides. Although not originally available with Windows 7 and Windows Server 2008 R2, Windows 8 and Windows Server 2012 support 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 Windows 8 and Windows Server 2012. 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, you must now set the *-RestrictToSingleComputer* option when creating the account. Don't forget that a single computer can have multiple managed service accounts as well.

In Active Directory schema, managed service accounts are represented by *msDS-ManagedServiceAccounts*. 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 fine-grained password policies. Instead, they use a randomly generated 240-byte (120-character) 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 using Active Directory Users And Computers or Add-ADGroupMember.

## Creating and Using Managed Service 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.

You create a managed service account 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 enabled by default as well, 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.

After you create a managed service account in Active Directory, you associate it with a target computer in Active Directory 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
```

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 Service Accounts

You can configure a service to run with the managed service account by following these steps:

1. In Server Manager, tap or click Tools and then tap or click Computer Management.
2. As necessary, connect to the computer you want to manage. In the left pane, press and hold or right-click the Computer Management node, and then tap or click Connect To Another Computer. Enter the host name, fully qualified domain name, or IP address of the remote server, and then tap or click OK.
3. In the left pane, expand the Services And Applications node, and then select the Services node.
4. Press and hold or right-click the name of the service you want to work with, and then tap or click Properties.
5. On the Log On tab, select This Account and then type the name of the managed service account in the format *DomainName\AccountName*, or tap or click Browse to search for the account.
6. Confirm that the password box is blank, and then tap or click OK.
7. Select the name of the service, and then tap or click Start to start the service, or tap 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 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.

## Removing Managed Service 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 snap-in 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 `Reset-ADServiceAccountPassword`. 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 Policy\Security Options`. Group Policy settings under `Account Policies>Password Policy` are not used to modify managed service account password-reset intervals, nor can the `NLTEST /SC_CHANGE_PWD` command be used to reset managed service account passwords.

## Moving Managed Service 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 `New-ADServiceAccount`.
2. Install the managed service account on the appropriate computer by using `Install-ADServiceAccount`, 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.

## 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, tap or click Tools and then tap or click Computer Management.
2. As necessary, connect to the computer you want to manage. In the left pane, press and hold or right-click the Computer Management node, and then tap or click Connect To Another Computer. Enter the host name, fully qualified domain name, or IP address of the remote server, and then tap or click OK.
3. In the left pane, expand the Services And Applications node, and then select the Services node.
4. Press and hold or right-click the name of the service you want to work with, and then tap or click Properties.
5. On the Log On tab, select This Account and then type the name of the service account in the format *SERVICE\ComputerName*.
6. Confirm that the password box is blank, and then tap or click OK.
7. Select the name of the service, and then tap or click Start to start the service, or tap 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.



# Managing User and Group Accounts

- Managing User Contact Information **347**
- Configuring the User's Environment Settings **351**
- Setting Account Options and Restrictions **355**
- Managing User Profiles **361**
- Updating User and Group Accounts **368**
- Managing Multiple User Accounts **375**
- Troubleshooting Logon Problems **378**
- Viewing and Setting Active Directory Permissions **380**

In a perfect world, you could create user and group accounts and never have to touch them again. Unfortunately, we live in the real world. After you create accounts, you'll spend a lot of time managing them. This chapter provides guidelines and tips to make that task easier.

## Managing User Contact Information

---

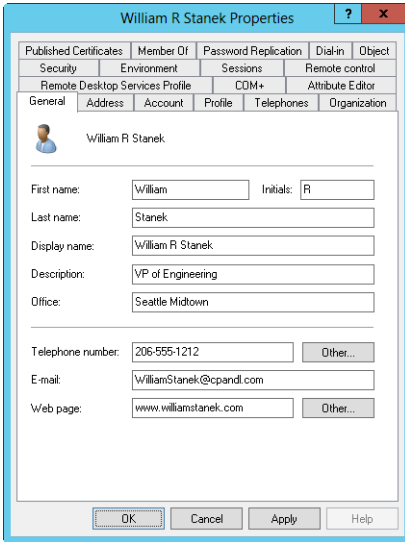
Active Directory is a directory service. When you create user accounts, those accounts can have detailed contact information associated with them. The contact information is then available for anyone in the domain tree or forest to use as criteria to search for users and to create address book entries.

### Setting Contact Information

You can set a user's contact information in Active Directory Users And Computers by following these steps:

1. Double-tap or double-click the user name in Active Directory Users And Computers. This opens the account's Properties dialog box.

2. Tap or click the General tab, shown in Figure 9-1. Set general contact information in the following text boxes:
  - **First Name, Initials, Last Name** Set the user's full name.
  - **Display Name** Sets the user's display name as seen in logon sessions and in Active Directory Domain Services.
  - **Description** Sets a description of the user.
  - **Office** Sets the user's office location.
  - **Telephone Number** Sets the user's primary business telephone number. If the user has other business telephone numbers you want to track, tap or click Other and then enter additional phone numbers in the Phone Number (Others) dialog box.
  - **E-Mail** Sets the user's business email address.
  - **Web Page** 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, tap or click Other and then enter additional webpage addresses in the Web Page Address (Others) dialog box.



**FIGURE 9-1** Configure general contact information for the user on the General tab.

**TIP** You must fill in the E-Mail and Web Page text boxes if you want to use the Send Mail and Open Home Page features of Active Directory Users And Computers. For more information, see "Updating User and Group Accounts" later in this chapter.

3. Tap or click the Address tab. Set the user's business or home address in the boxes provided. You'll usually want to enter the user's business address. You

can then track the business locations and mailing addresses of users at various offices.

**NOTE** You need to consider privacy issues before you enter users' home addresses. Discuss the matter with your human resources and legal departments. You might also want to get user consent before releasing home addresses.

4. Tap or click the Telephones tab. Enter the primary telephone numbers that should be used to contact the user, such as home, pager, mobile, fax, and IP phone.
5. You can configure other numbers for each type of telephone number. Tap or click the associated Other buttons, and then enter additional phone numbers in the dialog box provided.
6. Tap or click the Organization tab. As appropriate, enter the user's job title, department, and company.
7. To specify the user's manager, tap or click Change, and then select 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.
8. Tap or click Apply or OK to apply the changes.

You also can set contact information using Active Directory Administrative Center. Double-tap or double-click the user name. In the account's Properties dialog box, tap or click Organization to display the Organization panel. As Figure 9-2 shows, this panel provides a one-stop location for setting general, address, telephone, and organization details.

The screenshot shows the 'Organization' panel for a user. It is divided into several sections:

- General Information:** Display name (William R Stanek), Office (Seattle Midtown), E-mail (WilliamStanek@cpandl.com), Web page (www.williamstanek.com), Job title (VP of Engineering), Department (Engineering), Company (City Power & Light), and Manager (Rory R. Stanek).
- Phone Numbers:** Main (206-555-1212), Home, Mobile, Fax, Pager, and IP Phone. Includes a link for 'Other phone numbers...'.
- Direct Reports:** A list of direct reports (Dan Park, Yossi Banai, Thomas Axen, Hugo Garcia) with 'Add...' and 'Remove' buttons.
- Address:** Address (555 123rd St), City (Seattle), State (WA), ZIP (98501), and Country/Region (United States).
- Description:** A text box containing 'VP of Engineering'.

**FIGURE 9-2** The Organization panel provides a one-stop location for setting general, address, telephone, and organization details.

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, tap or 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. Tap or 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 text box. If no manager is set or you want to change the manager, tap or 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.

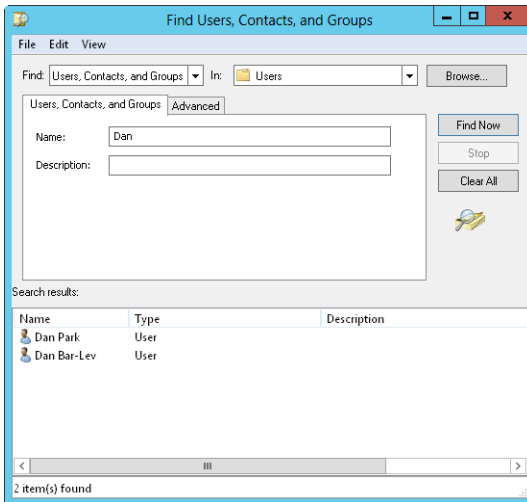
If the user has direct reports, she is listed under Direct Reports. You can add and remove direct reports using the related Add and Remove buttons. To add a direct report, tap or click Add, specify the direct report in the Select User Or Contact dialog box, and then tap or click OK. To remove a direct report, tap or click the name in the list and then tap or click Remove.

## Searching for Users and Groups in Active Directory

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, press and hold or right-click the domain or container, and then tap or click Find.
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 tap or click Browse to select a domain or container to search.
3. On the Users, Contacts, And Groups tab, type the name of the user, contact, or group you want to search for.
4. Tap or click Find Now to begin the search. If matches are found, the search results are displayed, as shown in Figure 9-3. Otherwise, type new search parameters and search again.
5. To manage an account, press and hold or right-click its entry. If you press and hold or right-click an account entry and then select Properties, you can open the account's Properties dialog box.

You can search for users and groups using the filter and global search features of Active Directory Administrative Center as well. For more information, see "Active Directory Administrative Center and Windows PowerShell" in Chapter 7, "Core Active Directory Administration."

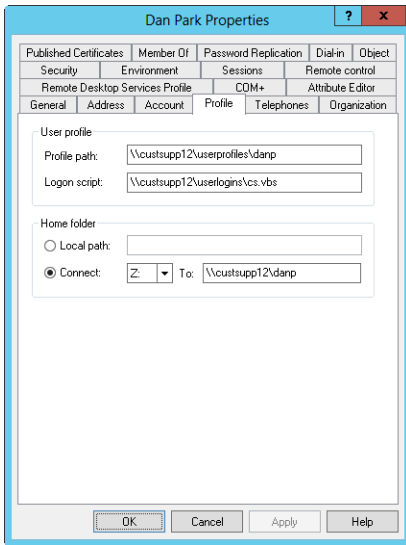


**FIGURE 9-3** Search for users in Active Directory, and then use the results to create address book entries.

## Configuring the User's Environment Settings


User accounts can also have profiles, logon scripts, and home directories associated with them. To configure these optional settings, double-tap or double-click a display name in Active Directory Users And Computers, and then tap or click the Profile tab, shown in Figure 9-4. 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. Setting the profile path is covered later in the chapter in "Managing User Profiles."
- **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. Chapter 4, "Automating Administrative Tasks, Policies, and Procedures," discusses logon scripts in detail.
- **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.



**FIGURE 9-4** The Profile tab allows you to create a user profile and thereby configure the network environment for a user.

In Active Directory Administrative Center, you configure a user's environment settings using the options on the Profile panel. To configure these settings, double-tap or double-click a display name in Active Directory Administrative Center and then tap or click Profile to display the Profile panel, shown in Figure 9-5.



**FIGURE 9-5** Configure the user's environment settings using the options on the Profile panel.


## System Environment Variables

System environment variables often come in handy 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 character, 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\Georgej. Use it with logon scripts.
- **%Processor\_Architecture%** The processor architecture of the user's computer, such as x86. Use it with logon scripts.

Figure 9-6 shows how you might use environment variables when creating user accounts. 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.



**FIGURE 9-6** When you use the Profile tab, environment variables can lessen the information you need to type, especially when you create an account based on another account.

## 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 Microsoft Windows commands. However, logon scripts can be any of the following:

- 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 tap or click the Profile tab.
2. Type the path to the logon script in the Logon Script text box. Be sure to set the full path to the logon script, such as **\\Zeta\User\_Logon\Eng.vbs**.

**NOTE** You can use other techniques to specify logon and logoff scripts. For complete details, see “User and Computer Script Management” in Chapter 4.

Creating logon scripts is easier than you might think, especially when you use the Windows command language. Just about any command you can type 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. You can set this information with the NET USE command. The following NET USE commands define a network printer and a network drive:

```
net use lpt1: \\zeta\techmain
net use G: \\gamma\corpfiles
```

If these commands were in the user’s logon script, the user would have a network printer on LPT1 and a network drive on G. You can create similar connections in a script. With 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 = "\\zeta\techmain"
wNetwork.AddWindowsPrinterConnection printerPath
wNetwork.SetDefaultPrinter printerPath

wNetwork.MapNetworkDrive "G:", "\\gamma\corpfiles"

Set wNetwork = vbEmpty
Set printerPath = vbEmpty
```

Here, you use the *AddWindowsPrinterConnection* method to add a connection to the TechMain printer on Zeta, 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 G.

## Assigning Home Directories

Windows Server 2012 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 disk drive or on a shared network drive. On a local drive, the directory is accessible only from a single workstation. On the other hand, shared network drives can be accessed from any computer on the network, which makes for a more versatile user environment.



**TIP** 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.

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.

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 tap or click the Profile tab.
2. Select Local Path in the Home Folder section, and then type the path to the home directory in the associated text 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 tap or 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. Type the complete path to the home directory using Universal Naming Convention (UNC) notation, such as **\\Gamma\User\_Dirs\%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** If you don't assign a home directory, Windows Server 2012 uses the default local home directory.

## Setting Account Options and Restrictions

---

Windows Server 2012 gives you many ways to control user accounts and their access to the network. You can define logon hours, permitted workstations for logon, dial-in privileges, and more.

### Managing Logon Hours

Windows Server 2012 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 normally 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

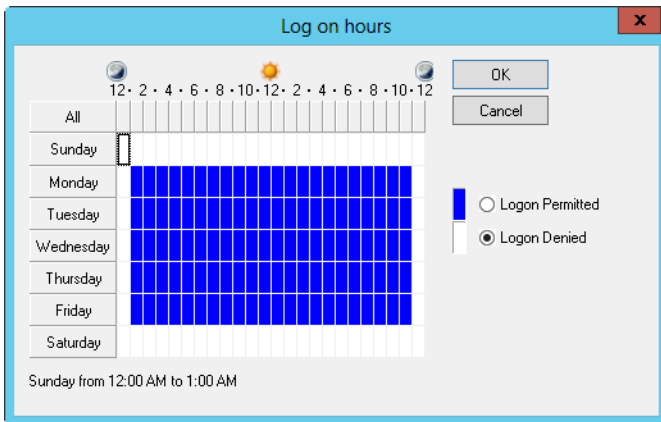
depends on the account policy you've set for them. Generally, one of two things happens to the user:

- **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.

## Configuring Logon Hours

To configure the logon hours, follow these steps:

1. Open the user's Properties dialog box. In Active Directory Users And Computers, tap or click the Account tab and then tap or click Logon Hours. In Active Directory Administrative Center, tap or click Log On Hours on the Account panel.
2. You can now set the valid and invalid logon hours using the Log On Hours dialog box, shown in Figure 9-7. 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.



**FIGURE 9-7** Configure logon hours for users.

3. To change the setting for an hour, tap or click it and then select either Logon Permitted or Logon Denied.

Table 9-1 lists Log On Hours dialog box options.

**TABLE 9-1** Log On Hours Dialog Box Options

FEATURE	FUNCTION
All	Allows you to select all the time periods
Days of the week buttons	Allow you to select all the hours in a particular day
Hourly buttons	Allow you to select a particular hour for all the days of the week
Logon Permitted	Sets the allowed logon hours
Logon Denied	Sets the disallowed logon hours

**TIP** 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 9–5 hours, 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.

## Enforcing Logon Hours

To forcibly disconnect users when their logon hours expire, follow these steps:

1. Access the Group Policy Object (GPO) you want to work with, as detailed in “Managing Site, Domain, and Organizational Unit Policies” in Chapter 4.
2. Open the Security Options node by working your way down through the console tree. Expand Computer Configuration, Windows Settings, and Security Settings. In Security Settings, expand Local Policies, and then select Security Options.
3. Double-tap or 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 tap or click Enabled. This turns on the policy restriction and enforces the logon hours. Tap or click OK.

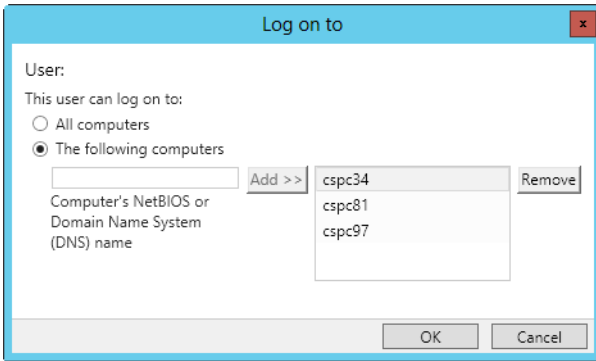
## Setting Permitted Logon Workstations

Windows Server 2012 has a formal policy that allows users to log on to systems locally. This policy controls whether a user can sit at the computer's keyboard and log on. By default, you can use any valid user account, including the Guest account, to log on locally to a workstation.

As you might imagine, allowing users to log on to any workstation is a security risk. 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. By defining a permitted workstation list, you close the opening in your domain and reduce the security risk. Now, not only must hackers find a user name and password, but they must also find the permitted workstations for the account.

For domain users, you define permitted logon workstations by following these steps:

1. Open the user's Properties dialog box. In Active Directory Users And Computers, tap or click the Account tab and then tap or click Log On To. In Active Directory Administrative Center, tap or click Log On To on the Account panel.
2. For the This User Can Log On To option, select The Following Computers, as shown in Figure 9-8.



**FIGURE 9-8** To restrict access to workstations, specify the permitted logon workstations.

3. Type the name of a permitted workstation, and then tap or click Add. Repeat this procedure to specify additional workstations.
4. If you make a mistake, select the erroneous entry and then tap or click Remove.

## Setting Dial-in and VPN Privileges

Windows Server 2012 lets you set remote access privileges for accounts on the Dial-In tab of the user's Properties dialog box. These settings control access for dial-in and virtual private networks (VPNs). Remote access privileges are controlled through Network Policy Server (NPS) Network Policy by default. This is the preferred method of controlling remote access. You can explicitly grant or deny dial-in privileges by selecting Allow Access or Deny Access. In any event, before users can remotely access the network, you need to follow these steps:

1. In Server Manager, add the role of Network Policy And Access Services.
2. To enable remote access connections, access the GPO for the site, domain, or organizational unit you want to work with, as specified in "Managing Site, Domain, and Organizational Unit Policies" in Chapter 4. In the policy editor, expand User Configuration, Administrative Templates, and then Network. Select Network Connections, and then configure the Network Connections policies as appropriate for the site, domain, or organizational unit.
3. Configure remote access 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 2012, you remove the payloads for roles and features that you are uninstalling using the `-Remove` parameter of the `Uninstall-WindowsFeature` cmdlet. Restore a removed payload using the `Install-WindowsFeature` cmdlet. By default, payloads are restored via Windows Update. Use the `-Source` parameter to restore a payload from a WIM mount point. In the following example, you restore the NPS and RRAS binaries via Windows Update:

```
install-windowsfeature -name npas-policy-server  
-includemanagementtools
```

```
install-windowsfeature -name remoteaccess  
-includeallsubfeature -includemanagementtools
```

After you grant a user permission to access the network remotely, follow these steps to configure additional dial-in parameters on the Dial-In tab of the user's Properties dialog box (as shown in Figure 9-9):

1. If the user must dial in from a specific phone number, select Verify Caller-ID and then type the telephone number from which this user is required to log on. Your telephone system must support Caller ID for this feature to work.

**NOTE** In Active Directory Administrative Center, the Dial-In tab is accessed from the Extensions panel. Tap or click Extensions and then tap or click Dial-In.

2. Define callback parameters using the following options:
  - **No Callback** Allows the user to dial in directly and remain connected. The user pays the long-distance telephone charges, if applicable.
  - **Set By Caller** Allows the user to dial in directly, and then the server prompts the user for a callback number. Once the number is entered, the user is disconnected, and the server dials the user back at the specified number to reestablish the connection. The company pays the long-distance telephone charges, if applicable.
  - **Always Callback To** Allows you to set a predefined callback number for security purposes. When a user dials in, the server calls back the preset number. The company pays the long-distance telephone charges, if applicable, and reduces the risk of an unauthorized person accessing the network.

**NOTE** You shouldn't assign callback numbers for users who dial in through a switchboard. The switchboard might not allow the user to properly connect to the network. You also shouldn't use preset callback numbers with multilinked lines. The multilinked lines won't function properly.

If necessary, you can also assign static IP addresses and static routes for dial-in connections by selecting Assign Static IP Addresses and Apply Static Routes, respectively.

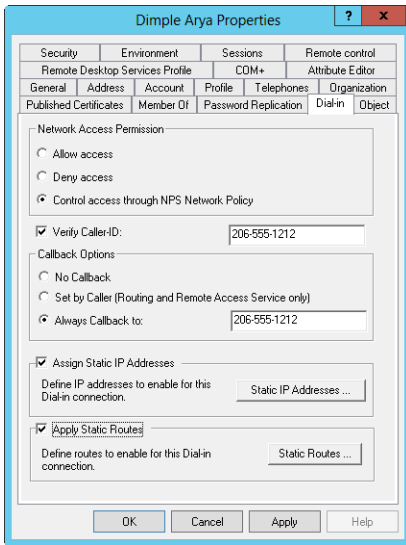


FIGURE 9-9 Dial-in privileges control remote access to the network.

## Setting Account 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 administrator accounts, you usually shouldn't use this option with normal user accounts.

- **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 using a smart card. The user can't log on to the workstation by typing 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 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 version 5 security procedure. The option to log on without it allows authentication from clients using a previous, or nonstandard, implementation of Kerberos.

**REAL WORLD** AES is one of several encryption standards. Another encryption standard is Data Encryption Standard (DES). Most computers running older versions of Windows support DES.

Computers running current releases of Windows support AES, which provides more secure encryption than DES. While U.S. versions support both 128-bit and 256-bit AES, versions exported for use outside the United States typically support only 128-bit encryption.

## Managing User Profiles

---

User profiles contain settings for the network environment, such as desktop configuration and menu options. 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 reboot the computer, go into Video Graphics Adapter (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.

Windows Server 2012 provides several ways to manage user profiles:

- You can assign profile paths in Active Directory Users And Computers or Active Directory Administrative Center.
- You can copy, delete, and change the type of an existing local profile with the System utility in Control Panel.
- You can set system policies that prevent users from manipulating certain aspects of their environment.

## Local, Roaming, and Mandatory Profiles

In Windows Server 2012, every user 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. 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.

### Working with Roaming and Mandatory Profiles

To reduce the confusion caused by multiple profiles, you can create a profile that other computers can access. This type of profile is called a *roaming profile*. By default, with a roaming profile, users can access the same profile no matter which computer they're using within the domain. Roaming profiles are server-based and can be stored on any Windows server. When a user with a roaming profile logs on, the profile is downloaded, which creates a local copy on the user's computer. When the user logs off, changes to the profile are updated both on the local copy and on the server.

**REAL WORLD** 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 or let users control 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.



**REAL WORLD** When a user has a mandatory profile, a temporary profile (meaning they are logged on as a guest), or a system profile, Windows 8 and Windows Server 2012 block 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 path.

## Restricting Roaming Profiles

Normally, users can access their roaming profile no matter which computer they're using within the domain. Windows 8 and Windows Server 2012 allow 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 by editing the advanced properties of a user or group in Active Directory and 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 the System\User Profiles path. 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 the System\Folder Redirection path.

The goal of these policies 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 normally 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 tap or 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. Tap or click *msDS-PrimaryComputer* and then tap or 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. Tap or click OK twice.

## Creating Local Profiles

User profiles are maintained either in a default directory or in the location set by the Profile Path text 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. If you don't change the default location, the user will have a local profile.

## Creating Roaming Profiles

Roaming profiles are stored on Windows servers. 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 directory by following these steps:

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. Type the path to the shared directory in the Profile Path text box. The path should have the form `\\server name\profile folder name\user name`. An example is `\\Zeta\User_Profiles\Georgej`, where *Zeta* is the server name, *User\_Profiles* is the shared directory, and *Georgej* is the user name.

The roaming profile is then stored in the `Ntuser.dat` file in the designated directory, such as `\\Zeta\User_Profiles\Georgej\Ntuser.dat`.

**NOTE** You don't usually need to create the profile directory. The directory is created automatically when the user logs on, and NTFS permissions are set so that only the user has access. 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 tapping or clicking the user names. Then when you right-click one of the selected users and then tap or click Properties, you can edit properties for all the selected users. Be sure to use `%UserName%` in the profile path, such as `\\Zeta\User_Profiles\%UserName%`.

3. 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 use the default local profile. Any changes the user makes to this profile are saved when the user logs off. The next time the user logs on, he has a personal profile.

## Creating Mandatory Profiles

Mandatory profiles are stored on servers running Windows Server. If you want a user to have a mandatory profile, you define the profile as follows:

1. Follow steps 1 and 2 in the previous section, "Creating Roaming Profiles."
2. Create a mandatory profile by renaming the `Ntuser.dat` file as `%UserName%\Ntuser.man`. The next time the user logs on, she will have a mandatory profile.

**NOTE** `Ntuser.dat` contains the registry settings for the user. When you change the extension for the file to `Ntuser.man`, you tell Windows Server to create a mandatory profile.

## Using the System Utility to Manage Local Profiles

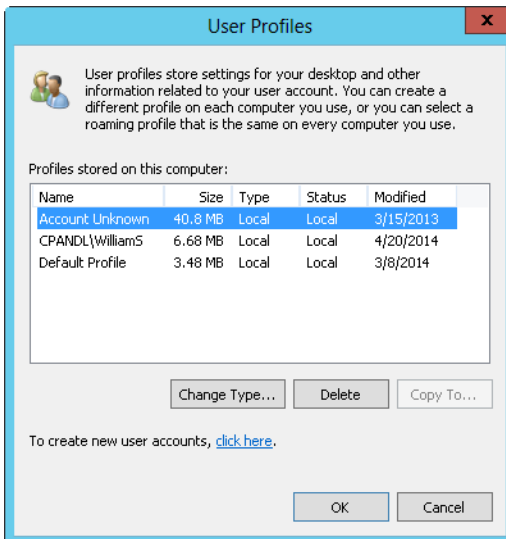
To manage local profiles, you need to log on to the user's computer. Then you can use the System utility in Control Panel to manage local profiles. To view current profile information, tap or click System And Security in Control Panel and then tap or click System. On the System page in Control Panel, tap or click Advanced System Settings. In the System Properties dialog box, under User Profiles, tap or click Settings.

As shown in Figure 9-10, 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 ADATUM\WrstaneK tells you that the original profile is from the domain adatum and the user account is wrstaneK.

**NOTE** If you delete an account but don't delete the associated profile, you might also see an entry that says Account Deleted or Account Unknown. 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.



**FIGURE 9-10** The User Profiles dialog box lets you manage existing local profiles.

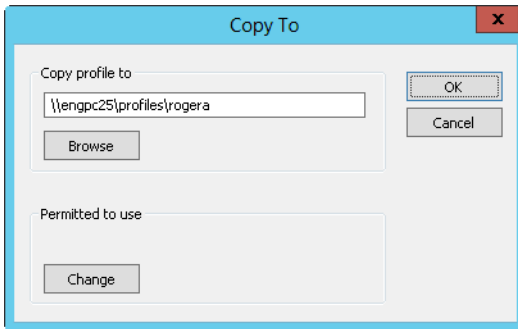
## Creating a Profile by Hand

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.

## Copying an Existing Profile to a New User Account

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. To do this, follow these steps to use the System Control Panel utility:

1. Start the System Control Panel utility. On the System page in Control Panel, tap or click Advanced System Settings. In the System Properties dialog box, under User Profiles, tap or click Settings.
2. Select the profile you want to copy from the Profiles Stored On This Computer list. (See Figure 9-10.)
3. Copy the profile to the new user's account by tapping or clicking Copy To. In the Copy Profile To box, shown in Figure 9-11, type the path to the new user's profile directory. For example, if you were creating the profile for georgej, you'd type `\\Zeta\User_Profiles\Georgej`.



**FIGURE 9-11** In the Copy To dialog box, enter the location of the profile directory and assign access permissions to the user.

4. Now you need to give the user permission to access the profile. In the Permitted To Use area, tap or click Change, and then use the Select User Or Group dialog box to grant access to the new user account.
5. Tap or click OK to close the Copy To dialog box. Windows then copies the profile to the new location.

**TIP** If you know the name of the user or group you want to use, you can save time by typing it directly into the Name box.

## Copying or Restoring a Profile

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 maintain 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, and start the System Control Panel utility. On the System page in Control Panel, tap or click Advanced System Settings. In the System Properties dialog box, under User Profiles, tap or click Settings.
2. In the Profiles Stored On This Computer list, select the profile you want to copy.
3. Copy the profile to the new location by tapping or clicking Copy To, and then type 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 type `\\Gamma\User_Profiles\Janew`.
4. To give the user permission to access the profile, tap or 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.
5. When you have finished, tap or click OK to close the Copy To dialog box. Windows then copies the profile to the new location.

## Deleting a Local Profile and Assigning a New One

Profiles are accessed when a user logs on to a computer. 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. For example, if a user's local profile becomes corrupt, you can delete the profile and assign a new one. 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 using an account with administrator privileges, and then start the System utility.
2. Tap or click Advanced System Settings. In the System Properties dialog box, under User Profiles, tap or click Settings.
3. Select the profile you want to delete, and then tap or click Delete. When asked to confirm that you want to delete the profile, tap or 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 reboot the computer.

The next time the user logs on, Windows Server does one of two things. Either the operating system gives the user the default local profile for that system, or it retrieves the user's roaming profile stored on another computer. To prevent the use of either of these profiles, you need to assign the user a new profile. To do this, you can do one of the following:

- Copy an existing profile to the user's profile directory. Copying profiles is covered in "Copying or Restoring a Profile" earlier in the chapter.
- Update the profile settings for the user in Active Directory Users And Computers. Setting the profile path is covered in "Creating Roaming Profiles."

### 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 tap or 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 (that was 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.

## Updating User and Group Accounts

---

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, press and hold or right-click the domain name and then tap or click Find.

2. In the Find list, select Custom Search. This updates the Find dialog box to display a Custom Search tab.
3. In the In list, select the area you want to search. To search the enterprise, select Entire Directory.
4. On the Custom Search tab, tap or 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 tap or click Add. If prompted to confirm, tap or click Yes.
6. Tap or 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 once is to hold down the Shift key, click the first object, and then click the last object.
8. Press and hold or 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, tap or click Field, choose Computer, and then select Computer Name (Pre-Windows 2000). With groups, do a custom search, tap or click Field, choose Group, and then select Group Name (Pre-Windows 2000).

The sections that follow examine other techniques you can use to update (re-name, copy, delete, and enable) accounts as well as to change and reset passwords. You'll also learn how to troubleshoot account logon problems.

## Renaming User and Group Accounts

When you rename a user account, you give the account a new label. As discussed in Chapter 8, "Creating User and Group 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 text 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. Press and hold or right-click the user account, and then tap or 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, shown in Figure 9-12.

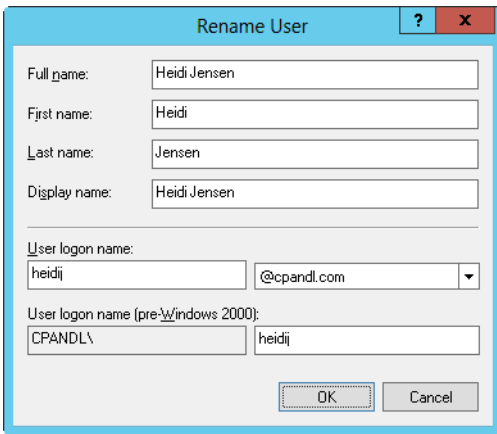


FIGURE 9-12 Fully rename an account.

3. Make the necessary changes to the user's name information, and then tap or click OK. If the user is logged on, you'll see a warning prompt telling you that the user should log off and then log back on 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.

## 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. Press and hold or right-click the account you want to copy, and then tap or 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.

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. Instead, Active Directory Users And Computers tries to copy only the information you need and to discard the information that you need to update. The following properties are retained:

- City, state, ZIP code, and country values set on the Address tab
- Department and company set on the Organization tab
- Account options set 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

**NOTE** If you used environment variables to specify the profile settings in the original account, the environment variables are used for the copy of the account as well. 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 2012 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 text file as the import source. You can run CSVDE 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=cpan1,DC=com",user,williams,William,Stanek,
williams@cpan1.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 run CSVDE using the general parameters listed previously as well 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=cpan1,DC=com" -r
"(objectClass=user)"
```

## Deleting User and Group Accounts

Deleting an account permanently removes the account. Once you delete an account, you can't create an account with the same name to get 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 2012 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 pressing and holding or right-clicking and selecting Delete. When prompted, tap or 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 tapping or 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 tapping or clicking the last account in the range.

**NOTE** When you delete a user account, Windows Server 2012 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.

## 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. Press and hold or right-click the account name, and then tap or click Reset Password or Set Password.
3. Type 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. Tap or 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. The following sections describe what to do when an account is disabled, locked out, or expired.

### Account Disabled

Active Directory Users And Computers and Active Directory Administrative Center depict disabled accounts with a down 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. Press and hold or 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, type `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 using the Enable Account option or disable them by selecting Disable Account. In Active Directory Administrative Center, enable the accounts using the Enable All option or disable them using Disable All.

### Account Locked Out

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-tap or 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.

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 on setting account policy, see "Configuring Account Policies" in Chapter 8.

## Account Expired

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-tap or double-click the user's account name. Open the Account tab or panel.
3. Under Account Expires, select End Of, and then tap or click the down 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.

## Managing Multiple User Accounts

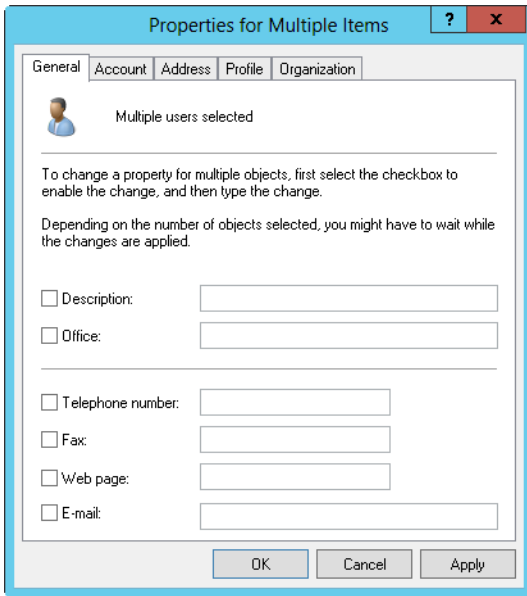
---

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 press and hold or 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. You'll see Add To Group, Disable All, Enable All, Unlock All, Move, Delete, and Properties.

The Properties option is the one we'll look at in the sections that follow. As shown in Figure 9-13, the Properties For Multiple Items dialog box has a different interface than the Properties dialog box for standard users.



**FIGURE 9-13** The Properties dialog box has a different interface when you work with multiple accounts.

**NOTE** The examples shown here and in the sections that follow are for Active Directory Users And Computers. The management techniques are similar for Active Directory Administrative Center.

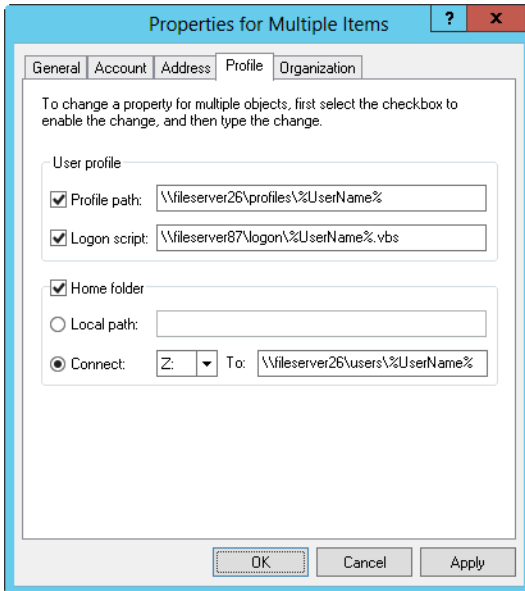
You should note the following differences:

- 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 properties' check boxes. After you do this, the value you enter in the text box is applied to all the selected accounts.

## Setting Profiles for Multiple Accounts

You set the profile information for multiple accounts with the options on the Profile tab. One of the best reasons to work with multiple accounts in Active Directory Users And Computers is that you can set all their environment profiles 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. Thus, the users named bobs, janew, and ericl would all be assigned the following unique logon scripts: Bobs.cmd, Janew.cmd, and Ericl.cmd.

Figure 9-14 shows an example of setting environment profile information for multiple accounts. Note that the %UserName% variable is used to assign the user profile path, the user logon script name, and the home folder.



**FIGURE 9-14** Use the %UserName% environment variable to assign paths and file names based on individual user names.

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.

## Setting Logon Hours for Multiple Accounts

When you select multiple user accounts in Active Directory Users And Computers, you can manage their logon hours collectively. To do this, follow these steps:

1. Select the accounts you want to work with in Active Directory Users And Computers.
2. Press and hold or right-click the highlighted accounts, and then tap or click Properties. In the Properties dialog box, tap or click the Account tab.
3. Select the Logon Hours check box, and then tap or click Logon Hours. You can then set the logon hours as discussed in "Configuring Logon Hours" earlier in the chapter.

**NOTE** 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.

## Setting Permitted Logon Workstations for Multiple Accounts

You set the permitted logon workstations for multiple accounts with the Logon Workstations dialog box. To open this dialog box, follow these steps:

1. Select the accounts you want to work with in Active Directory Users And Computers.
2. Press and hold or right-click the highlighted accounts, and then tap or click Properties. In the Properties dialog box, tap or click the Account tab.
3. Select the Computer Restrictions check box, and then tap or click Log On To.
4. If you want to allow the users to log on to any workstation, select All Computers. If you want to specify which workstations users are permitted to use, tap or click The Following Computers button, and then enter the names of the workstations. When you tap or click OK, these settings are applied to all the selected user accounts.

## Setting Logon, Password, and Expiration Properties for Multiple Accounts

User accounts have many options that control logon, passwords, and account expiration. You set these values on the Account tab of the Properties dialog box. 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 tap or 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 tap or 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.

## 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, press and hold or right-click the account in Active Directory Users And Computers and then tap or click Enable Account.

You can also search the entire domain for users with disabled accounts by typing **dsquery user –disabled** at a command prompt. To enable a disabled account from the command line, type **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 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 2012 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 Computer Configuration\Windows Settings\Security Settings\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 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.

When a domain is operating in Windows 2000 native or Windows Server 2003 mode, 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 "Configuring User Rights Policies" 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 “Local, Roaming, and Mandatory Profiles” earlier in the 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 “Setting Permitted Logon Workstations for Multiple Accounts” earlier in the chapter.

## Viewing and Setting Active Directory Permissions

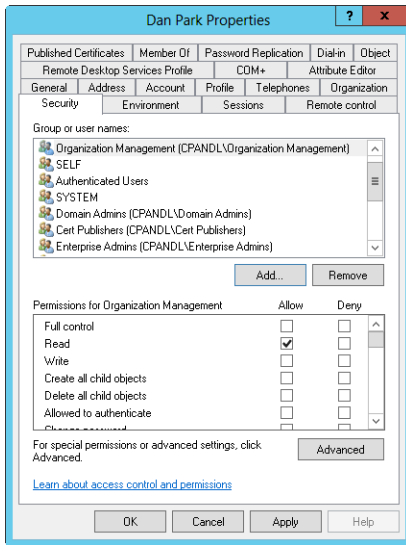
---

As you know from previous discussions, user, group, and computer accounts are represented in Active Directory as objects. Active Directory objects 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, press and hold or right-click the user, group, or computer account you want to work with, and then tap or click Properties.
2. In the Properties dialog box, tap or click the Security tab. As shown in Figure 9-15, you should now see a list of groups and users that have been assigned permissions on the object you previously selected. If the permissions are dimmed, 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, tap or 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. Tap or click Check Names, and then tap or click OK. In the check boxes in the Permissions area, allow or deny permissions. Repeat this step for other users, computers, or groups.



**FIGURE 9-15** View and configure object permissions on the Security tab.

6. Tap or 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, press and hold or right-click the user, group, or computer account you want to work with, and then tap or click Properties.
2. In the Properties dialog box, tap or click the Security tab and then tap or click Advanced. You should now see a list of individual permission entries for the previously selected object. 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 tap or 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. Tap or click OK twice when you have finished.



**PART III**

# Windows Server 2012 Data Administration

**CHAPTER 10** Managing File Systems and Drives **385**

**CHAPTER 11** Configuring Volumes and RAID Arrays **419**

**CHAPTER 12** Data Sharing, Security, and Auditing **457**

**CHAPTER 13** Data Backup and Recovery **519**



# Managing File Systems and Drives

- Managing the File Services Role **385**
- Adding Hard Disk Drives **389**
- Working with Basic, Dynamic, and Virtual Disks **401**
- Using Basic Disks and Partitions **406**
- Compressing Drives and Data **411**
- Encrypting Drives and Data **413**

A hard disk drive is the most common storage device used on network workstations and servers. Users depend on hard disk drives to store their word-processing 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 available on most workstations and servers is an example of a local file system. You access the C drive 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 using the Map Network Drive feature of File Explorer.

Wherever disk resources are located, your job as a system administrator is to manage them. The tools and techniques you use to manage file systems and drives are discussed in this chapter. Chapter 11, "Configuring Volumes and RAID Arrays," looks at partition management, volume sets, and fault tolerance.

## Managing the File Services Role

---

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. In earlier releases of the Microsoft Windows Server operating system, all servers were installed with basic file services.

With Windows Server 2012, you must specifically configure a server to be a file server by adding the File Services role and configuring this role to use the appropriate role services.

Table 10-1 provides an overview of the role services associated with the File Services role. When you install the File Services role, you might also want to install the following optional features, available through the Add Features Wizard:

- **Windows Server Backup** The standard backup utility included with Windows Server 2012.
- **Enhanced Storage** Supports additional functions made available by devices that support hardware encryption and enhanced storage. Enhanced storage devices support Institute of Electrical and Electronic 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.

If the binaries for the tools have been removed, you need to install the tools by specifying a source, as discussed in “Server Manager Essentials and Binaries” in Chapter 2, “Managing Servers Running Windows Server 2012.”

**TABLE 10-1** Role Services for File Servers

ROLE SERVICE	DESCRIPTION
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 128-KB 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	Allows 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.



ROLE SERVICE	DESCRIPTION
DFS Replication	Allows you to synchronize folders on multiple servers across local or wide area network connections 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	Allows 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. Using FSRM, administrators can generate storage reports, configure quotas, and define file-screening policies.
File Server VSS Agent Service	Allows 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 non-Windows iSCSI initiators and 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 Services for Network File System (NFS), users can transfer files between Windows Server and UNIX operating systems by using the NFS protocol.
Storage Services	Allows 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.

You can add the File Services role to a server by following these steps:

1. In Server Manager, tap or click Manage and then tap or 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 tap or click Next.

2. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Tap or 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, tap or click Browse and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, tap or click Next.

**NOTE** Only servers running Windows Server 2012 and that have been added for management in Server Manager are listed.

4. On the Server Roles page, select File And Storage Services. Expand the related node, and select the additional role services to install. If additional features are required to install a role, you'll see an additional dialog box. Tap or click Add Features to close the dialog box and add the required features to the server installation. When you are ready to continue, tap or click Next.

**NOTE** A summary of each role service is provided in Table 10-1. To allow for interoperability with UNIX, be sure to add Server For NFS.

5. On the Features page, select any features you want to install. If additional features are required to install a feature you selected, you'll see an additional dialog box. Tap or click Add Features to close the dialog box and add the required features to the server installation. When you are ready to continue, tap or click Next.
6. On the Confirm page, tap or click the Export Configuration Settings link to generate an installation report that can be displayed in Internet Explorer.
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.

**REAL WORLD** You also can specify an alternate path for the required source files. To do this, click the Specify An Alternate Source Path link, type that alternate path in the box provided, and then tap or 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.

8. After you review the installation options and save them as necessary, tap or click Install to begin the installation process. The Installation Progress page tracks the progress of the installation. If you close the wizard, tap or click the Notifications icon in Server Manager and then tap or click the link provided to reopen the wizard.
9. 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.

If the File Services role is installed already on a server and you want to install additional services for a file server, you can add role services to the server by using a similar process.

## Adding Hard Disk Drives

---

Before you make a hard disk drive available to users, you need to configure it and consider how it will be used. With Windows Server 2012, you can configure hard disk drives in a variety of ways. The technique you choose depends primarily on the type of data you're working with 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 hard disk 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 2012 supports drive sets and arrays using redundant array of independent disks (RAID) technology, which is built into the operating system.

### 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 2012 supports both Standard Format and Advanced Format hard drives. Standard Format drives have 512 bytes per physical sector and are also referred to as *512b drives*. Advanced Format drives have 4096 bytes per physical sector and are also referred to as *512e drives*. 512e represents a significant shift for the hard 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; 512e disks work with data 4096 bytes at a time. At an elevated, administrator prompt, you can use `Fsutil` to determine bytes per physical sector by typing the following:

```
Fsutil fsinfo ntfsinfo DriveDesignator
```

where *DriveDesignator* is the designator of the drive to check, such as:

```
Fsutil fsinfo sectorinfo c:
```

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, hard disks 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 2012 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 hard disk drives. This 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), though popular, is giving way to 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 low-cost 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. eSATA (also known as external SATA) is meant for externally connected drives.

**NOTE** Windows Server 2012 features enhancements to provide improved support for SATA drives. These enhancements reduce metadata inconsistencies and allow drives to cache data more efficiently. Improved disk caching helps to protect cached data in the event of an unexpected power loss.

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 input/output (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 Ultra320 SCSI drives with 15,000 RPM rotational speed from different vendors, you will probably find similar transfer rates and MTTF. For example, the Maxtor Atlas 15K II has a maximum sustained data-transfer rate of up to 98 megabytes per second (MBps). The Seagate Cheetah 15K.4 has a maximum sustained data-transfer rate of up to 96 MBps. Both have an MTTF of 1.4 million hours. Transfer rates can also be expressed in 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 see 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. The Seagate Barracuda 7200 SATA II drive has a rotational speed of 7200 RPM and an average sustained transfer rate of 58 MBps. With an average seek time of 8.5 msec and an MTTF of 1 million hours, the drive performs comparably to other 7200 RPM SATA II drives. However, most Ultra320 SCSI drives perform better and are better at multiuser read/write operations, too.

**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. With SATA, the maximum data transfer rate is usually around 150 MBps or 300 MBps. With PATA/IDE, the maximum data transfer rate is usually around 100 MBps.

Temperature is another important factor to consider when you're selecting a drive, but it's a factor few administrators take into account. Typically, the faster a drive rotates, 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 run hot, and you must be sure to carefully regulate temperature. Both the Maxtor Atlas 15K II and the Seagate Cheetah 15K.4 can become nonoperational at temperatures of 70 degrees Centigrade or higher (as would most other drives).

Windows Server 2012 adds support for disk drives with hardware encryption (referred to as encrypted hard drives). Encrypted hard drives have built-in processors that shift the encryption-decryption activities from the operating system to hardware, freeing up operating system resources. Windows Server 2012 will use hardware encryption with BitLocker when available. Other security features available in Windows Server 2012 include Secured Boot and Network Unlock. Secured 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. For more information on TPM, BitLocker, Secured Boot, Network Unlock, and encrypted hard drives, see “Using TPM and BitLocker Drive Encryption” in Chapter 11 of *Windows 8 Administration Pocket Consultant* (Microsoft Press, 2012).

## 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.

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 hard disk 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, Itanium-based computers. GPT is recommended for disks larger than 2 TB on x86 and x64 systems or any disks used on Itanium-based computers. 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 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.

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 Basic, Dynamic, and Virtual 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 2012 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 you are working with 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 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). exFAT defines its file allocation tables using 64 bits. This allows exFAT to overcome the 4-GB file-size limit and the 32-GB volume-size 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 to store information about files. This database is called the master file table (MFT).

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.

Resilient File System 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, 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 due to 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, there is no need for a Check Disk feature 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 supports 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 "Standards-Based Storage Management" in Chapter 11.

## Using Disk Management

You use the Disk Management snap-in for the Microsoft Management Console (MMC) to configure drives. Disk Management makes it easy to work with the internal and external drives on a local or remote system. 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.

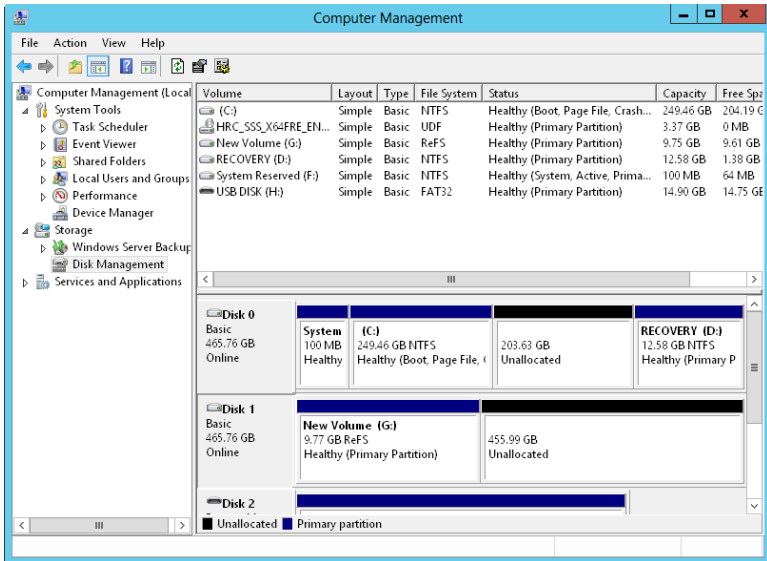
Disk Management has three views: Disk List, Graphical View, and Volume List. With remote systems, you're limited in the tasks you can perform with Disk Management. Remote management tasks you can perform include viewing drive details, changing drive letters and paths, and converting disk types. With 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** Before you work with Disk Management, you should know several things. If you create a partition but don't format it, the partition is labeled as Free Space. If you haven't assigned a portion of the disk to a partition, this section of the disk is labeled Unallocated.

In Figure 10-1, the Volume List view is in the upper-right corner and Graphical View is used 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.



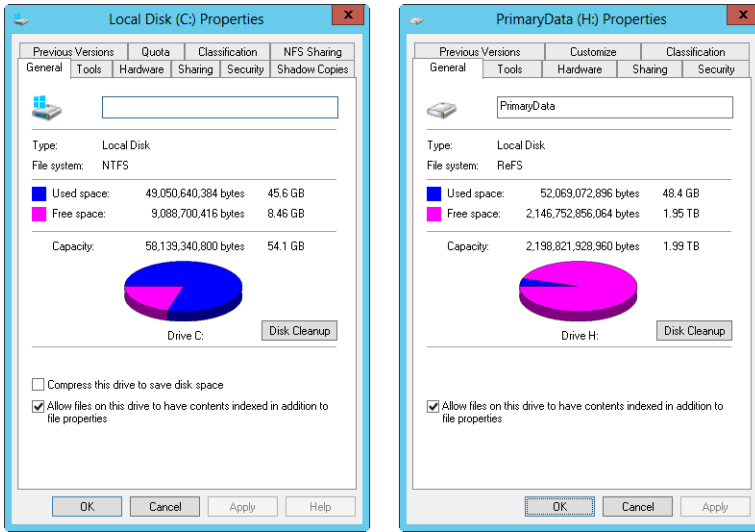


**FIGURE 10-1** In Disk Management, the upper view provides a detailed summary of all the drives on the computer and the lower view provides an overview of the same drives by default.

Windows Server 2012 supports four types of disk configurations:

- **Basic** The standard fixed disk type used in previous versions of Windows. Basic disks are divided into partitions and are the original disk type for early Windows operating systems.
- **Dynamic** An enhanced fixed disk type for Windows Server 2012 that you can update without having to restart the 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.

From the Disk Management window, you can get more detailed information on a drive section by pressing and holding or right-clicking it and then selecting Properties. When you do this, you see a dialog box. Figure 10-2 shows the dialog boxes for two fixed disks. The one on the left uses NTFS. The one on the right uses ReFs. Both disks have additional tabs based on the server configuration.



**FIGURE 10-2** The General tab of the Properties dialog box provides detailed information about a drive.

If you've configured remote management through Server Manager and MMCs, as discussed in Chapter 2, 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.

## Removable Storage Devices

Removable storage devices can be formatted with 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 universal serial bus (USB) or a FireWire interface. When working with USB and FireWire, the transfer

speed and overall performance of the device from a user's perspective depends primarily on the version supported. Currently, several versions of USB and FireWire are used.

USB 2.0 is the industry standard, while the world transitions to USB 3.0. USB 2.0 devices can be rated as either full speed (up to 12 Mbps) or high speed (up to 480 Mbps). Although high-speed USB 2.0 supports data transfers at a maximum rate of 480 Mbps, sustained data-transfer rates usually are 10–30 Mbps. The actual sustainable transfer rate depends on many factors, including the type of device, the data you are transferring, and the speed of a computer. Each USB controller on a computer has a fixed amount of bandwidth, which all devices attached to the controller must share. The data transfer rates are significantly slower if a computer's USB port is an earlier version than the device you are using. For example, if you connect a USB 2.0 device to a USB 1.0 port or vice versa, the device operates at the significantly reduced USB 1.0 transfer speed.

USB 1.0, 1.1, and 2.0 ports all look alike. However, most USB 3.0 ports I've seen have a special coloring to differentiate them. Still, the best way to determine which type of USB ports a computer has is to refer to the documentation that comes with the computer. Newer monitors have USB 2.0 ports to which you can connect devices as well. When you have USB devices connected to a monitor, the monitor acts like a USB hub device. As with any USB hub device, all devices attached to the hub share the same bandwidth, and the total available bandwidth is determined by the speed of the USB input to which the hub is connected on a computer.

FireWire (IEEE 1394) is a high-performance connection standard that uses a peer-to-peer architecture in which peripherals negotiate bus conflicts to determine which device can best control a data transfer. Like USB, several versions of FireWire are currently used. FireWire 400 (IEEE 1394a) has maximum sustained transfer rates of up to 400 Mbps. IEEE 1394b allows 400 Mbps (S400), 800 Mbps (S800), and 1600 Mbps (S1600). As with USB devices, if you connect an IEEE 1394b device to a IEEE 1394a port or vice versa, the device operates at the significantly reduced FireWire 400 transfer speed.

As with USB ports, the sustained transfer rate for IEEE 1394a and IEEE 1394b ports will be considerably less than the maximum rate possible. IEEE 1394a and IEEE 1394b ports and cables have different shapes, making it easier to tell the difference between them—if you know what you're looking for. FireWire 400 cables without bus power have four pins and four connectors. FireWire 400 cables with bus power have six pins and six connectors. FireWire 800 and FireWire 1600 cables always have bus power and have nine pins and nine connectors.

Another option is external SATA (eSATA), which is available on newer computers and is an ultra-high-performance connection for data transfer to and from external mass storage devices. eSATA operates at speeds up to 3 Gbps. You can add support for eSATA devices by installing an eSATA controller card.

When you are purchasing an external device for a computer, you'll also want to consider what interfaces it supports. In some cases, you might be able to get a device with more than one interface, such as one that supports USB 3.0 and eSATA. A device with multiple interfaces gives you more options.

Working with removable disks is similar to working with fixed disks. You can do the following:

- Press and hold or right-click a removable disk and select Open or Explore to examine the disk's contents in File Explorer.
- Press and hold or right-click a removable disk and 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.
- Press and hold or right-click a removable disk and select Properties to view or set properties. On the General tab of the Properties dialog box, you can set the volume label as discussed in "Changing or Deleting the Volume Label" in Chapter 11.

When you work with removable disks, you can customize disk and folder views. To do this, press and hold or right-click the disk or folder, select Properties, and then tap or 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.

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.

Removable disks differ from standard NTFS sharing in that they don't necessarily have an underlying security architecture. 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.

## 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 do this, 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, reboot.

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 initialize the disks by following these steps:

1. Each disk you install needs to be initialized. 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. Tap or 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 disk. In the Disk List view, the 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. You can then press and hold or right-click the disk's icon and select Online. Press and hold or right-click the disk's icon again, and select Initialize Disk. You can then initialize the disk as discussed previously.

## 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. Table 10-2 summarizes the most common status values.

**TABLE 10-2** Common Drive Status Values

STATUS	DESCRIPTION	RESOLUTION
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 pressing and holding or 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).

STATUS	DESCRIPTION	RESOLUTION
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.	Press and hold or right-click the disk, and then tap or 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 see this status if the card is unformatted or improperly formatted. You might also see 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. Press and hold or right-click the disk, and then tap or click Rescan Disk (on the Action menu) to try to correct the problem. You might also want to reboot the system.
Unrecognized	The disk is of an unknown type and can't be used on the system. A drive from a non-Windows system 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.
Not Initialized	The disk doesn't have a valid signature. A drive from a non-Windows system 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 2012, press and hold or right-click the disk, and then tap or 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 Basic, Dynamic, and Virtual Disks

---

Windows Server 2012 supports basic, dynamic, and virtual disk configurations. This section discusses techniques for working with each disk configuration type.

**NOTE** You can't use dynamic disks on portable computers or with removable media.

### Using Basic and Dynamic Disks

Normally, Windows Server 2012 disk partitions are initialized as basic disks. You can't create new fault-tolerant drive sets 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, 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 you are working with 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

## 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.

**NOTE** You can mark a partition as active using Disk Management. In Disk Management, press and hold or right-click the primary partition you want to mark as active, and then tap or 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

Basic disks are designed to be used with previous versions of Windows. Dynamic disks are designed to let you take advantage of the latest Windows features. Only computers running Windows 2000 or later releases of Windows can use dynamic disks. However, you can use dynamic disks with other operating systems, such as UNIX. To do this, you need to create a separate volume for the non-Windows operating system. You can't use dynamic disks on portable computers.

Windows Server 2012 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 other versions of Windows. Only computers running Windows 2000 and later releases of Windows can use 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 using Windows Server 2012.

To convert a basic disk to a dynamic disk, follow these steps:

1. In Disk Management, press and hold or 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 tap or click Convert To Dynamic Disk.
2. In the Convert To Dynamic Disk dialog box, select the check boxes for the disks you want to convert. If you're converting a spanned, striped, mirrored, or RAID-5 volume, be sure to select all the basic disks in this set. You must convert the set together. Tap or click OK to continue. This displays the Disks To Convert dialog box.

The Disks To Convert dialog box 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.
3. To begin the conversion, tap or click Convert. Disk Management warns you that after the conversion is complete, you won't be able to boot previous versions of Windows from volumes on the selected disks. Tap or click Yes to continue.
  4. 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, press and hold or right-click the disk and select Convert To Basic Disk. This changes 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, press and hold or right-click the dynamic disk you want to reactivate, and then tap or 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 them from one computer to another. For example, if after setting up a computer you decide that you don't really need an additional hard disk, you can move it to another computer where it can be better used.

Windows Server 2012 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 drives are currently installed. Check the status of the drives, and ensure that they're marked as Healthy. If the status isn't Healthy, you should repair partitions and volumes before you move the disk drives.

**NOTE** Drives with BitLocker Drive Encryption cannot be moved 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 hard disk subsystems on the original computer and the computer to which you want to transfer the disk. Both computers should have identical hard disk subsystems. If they don't, the Plug and Play ID on the system disk 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. Press and hold or right-click each disk you want to move, and then tap or click Uninstall. If you are unsure which disks to uninstall, press and hold or right-click each disk and tap or click Properties. In the Properties dialog box, tap or click the Volumes tab and then select Populate. This shows you 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, press and hold or right-click each disk, and then tap or 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, press and hold or right-click any disk marked Foreign, and then tap or 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

Using Disk Management, you can create, attach, and detach virtual hard disks. You can create a virtual hard disk by choosing Create VHD from the Action menu. In the Create And Attach Virtual Hard Disk dialog box, tap or click Browse. Use the Browse

Virtual Disk Files dialog box to select the location where you want to create the .vhd file for the virtual hard disk, and then tap or click Save.

In the Virtual Hard Disk Size list, enter the size of the disk in MB, GB, or TB. Specify whether the size of the VHD dynamically expands to its fixed maximum size as data is saved to it or uses a fixed amount of space regardless of the amount of data stored on it. When you tap or click OK, Disk Management creates the virtual hard disk.

The VHD is attached automatically and added as a new disk. To initialize the disk for use, press and hold or right-click the disk entry in Graphical View and then tap or 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 tap or click OK.

After initializing the disk, press and hold or right-click the unpartitioned space on the disk and create a volume of the appropriate type. After you create the volume, the VHD is available for use.

Once 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 pressing and holding or 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 pressing and holding or right-clicking the disk entry in Graphical View, selecting Detach VHD, and then tapping or 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, tap or click the Attach VHD option on the Action menu.
2. In the Attach Virtual Hard Disk dialog box, tap or click Browse. Use the Browse Virtual Disk Files dialog box to select the .vhd file for the virtual hard disk, and then tap or click Open.
3. If you want to attach the VHD in read-only mode, select Read-Only. Tap or click OK to attach the VHD.

## 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 using Disk Management.

### Partitioning Basics

In Windows Server 2012, 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 to assign drive letters. 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 2012. 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 2012 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.

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.

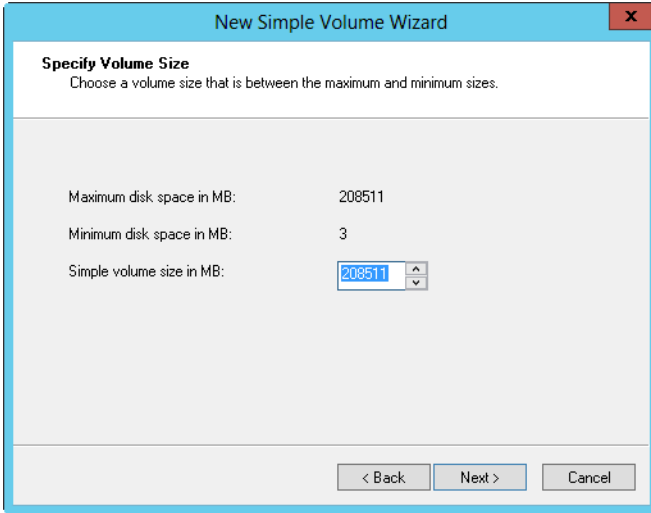
To help you differentiate between primary partitions and extended partitions with logical drives, Disk Management color codes the partitions. For example, primary partitions might be color coded with a dark-blue band and logical drives in extended partitions might be color coded with a light-blue band. 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 2012 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:

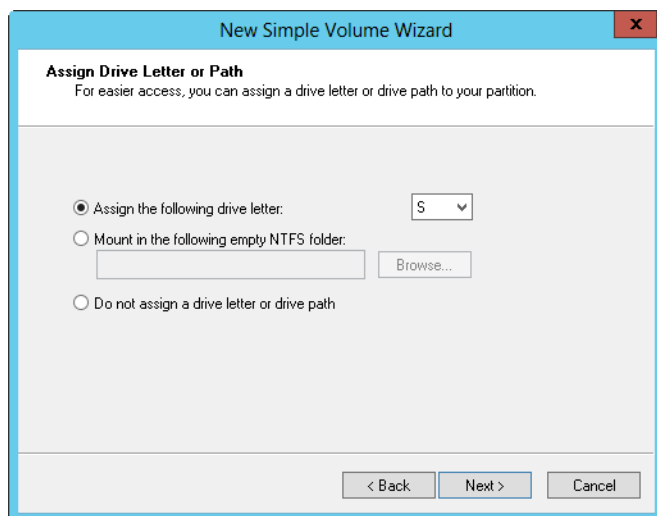
1. In Disk Management's Graphical View, press and hold or right-click an unallocated or free area, and then tap or click New Simple Volume. This starts the New Simple Volume Wizard. Read the Welcome page, and then tap or click Next.
2. The Specify Volume Size page, shown in Figure 10-3, specifies the 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 tap or click Next.



**FIGURE 10-3** Set the size of the volume on the Specify Volume Size page.

3. On the Assign Drive Letter Or Path page, shown in Figure 10-4, specify whether you want to assign a drive letter or path and then tap or 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 2012 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 tap 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.

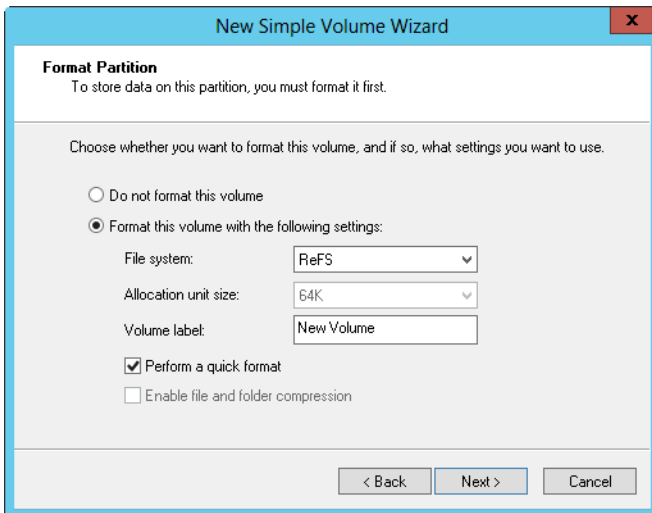
**NOTE** You don't have to assign volumes a drive letter or a path. A volume with no designators is considered to be unmounted and is for the most part unusable. An unmounted volume can be mounted by assigning a drive letter or a path at a later date. See "Assigning Drive Letters and Paths" in Chapter 11.



**FIGURE 10-4** On the Assign Drive Letter Or Path page, assign the drive designator or choose to wait until later.

4. On the Format Partition page, shown in Figure 10-5, 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 pressing and holding or right-clicking the volume in File Explorer, tapping or clicking Properties, and typing a new value in the Label box provided on the General tab.

- **Perform A Quick Format** Tells Windows Server 2012 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.



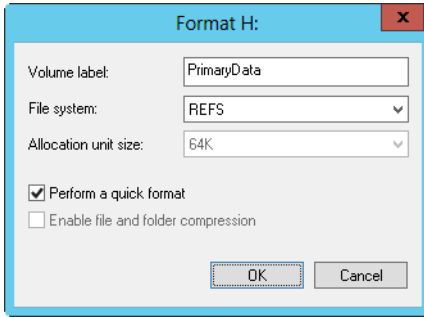
**FIGURE 10-5** Set the formatting options for the partition on the Format Partition page.

5. Tap or click Next, confirm your options, and then tap or click Finish.

## 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, press and hold or right-click the partition and then tap or click Format. This opens the Format dialog box, shown in Figure 10-6.





**FIGURE 10-6** Format a partition in the Format dialog box by specifying its file system type and volume label.

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 2012 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, tap or click OK. Because formatting a partition destroys any existing data, Disk Management gives you one last chance to cancel the procedure. Tap or 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.

## Compressing Drives and Data

When you format a drive for NTFS, Windows Server 2012 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. Note that File Explorer shows the names of compressed resources in blue.

**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 2012 automatically decrypts the data and then compresses it. Likewise, if you try to encrypt compressed data, Windows Server 2012 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, press and hold or right-click the drive you want to compress and then tap or click Properties.
2. On the General tab, select Compress Drive To Save Disk Space and then tap or click OK.
3. In the Confirm Attribute Changes dialog box, select whether to apply the changes to subfolders and files and then tap or click OK.

## Compressing Directories and Files

If you decide not to compress a drive, Windows Server 2012 lets you selectively compress directories and files. To compress a file or directory, follow these steps:

1. In File Explorer, press and hold or right-click the file or directory you want to compress and then tap or click Properties.
2. On the General tab of the Properties dialog box, tap or click Advanced. In the Advanced Attributes dialog box, select the Compress Contents To Save Disk Space check box. Tap or 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 tap or click OK. Once 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, press and hold or right-click the drive that contains the data you want to expand, and then tap or click Properties.
2. Clear the Compress Drive To Save Disk Space check box, and then tap or click OK.
3. In the Confirm Attribute Changes dialog box, select whether to apply the change to subfolders and files and then tap or 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 you want to expand a compressed file or directory, follow these steps:

1. Press and hold or right-click the file or directory in File Explorer, and then tap or click Properties.
2. On the General tab of the Properties dialog box, tap or click Advanced. Clear the Compress Contents To Save Disk Space check box. Tap or 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 tap or 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 of the major advantages is the capability to automatically encrypt and decrypt data 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.

## 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, 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 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 discussed in "Backing Up and Restoring the System State" in Chapter 13, "Data Backup and Recovery." 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 stand-alone 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, press and hold or right-click the file or directory you want to encrypt and then tap or click Properties.
2. On the General tab of the Properties dialog box, tap or click Advanced and then select the Encrypt Contents To Secure Data check box. Tap or 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 tap or 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 pressing and holding or right-clicking the file or folder in File Explorer and then selecting Properties. On the General tab of the Properties dialog box, tap or click Advanced. In the Advanced Attributes dialog box, tap or 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, tap or click Add. If a user certificate is available for the user, select the user's name in the list provided and then tap or click OK. Otherwise, tap or 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. Press and hold or right-click the file, and then select Properties. On the General tab of the Properties dialog box, tap or click Advanced. The Encrypt Contents To Secure Data option should be selected.

## Configuring Recovery Policy

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 stand-alone workstation.

Through the Group Policy console, you can view, assign, and delete recovery agents. To do that, follow these steps:

1. Open the Group Policy console for the local computer, site, domain, or organizational unit you want to work with. For details on working with Group Policy, see “Understanding Group Policies” in Chapter 4, “Automating Administrative Tasks, Policies, and Procedures.”
2. Open the Encrypted Data Recovery Agents node in Group Policy. To do this, expand Computer Configuration, Windows Settings, Security Settings, and 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, press and hold or right-click Encrypting File System and then tap or 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. Tap or 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, tap or click Browse Directory and then—in the Find Users, Contacts, And Groups dialog box—select the user you want to work with. You’ll then be able to use the published certificate of that user. If you want to use a certificate file, tap or 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, tap or 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, press and hold or right-click the file or directory and then tap or click Properties.
2. On the General tab of the Properties dialog box, tap or click Advanced. Clear the Encrypt Contents To Secure Data check box. Tap or 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 tap or click OK.

**TIP** Windows Server also provides a command-line utility called Cipher (Cipher.exe) for encrypting and decrypting your data. Typing **cipher** at a command prompt without additional parameters shows you the encryption status of all folders in the current directory.



# Configuring Volumes and RAID Arrays

- Using Volumes and Volume Sets **420**
- Improving Performance and Fault Tolerance with RAID **426**
- Implementing RAID on Windows Server 2012 **427**
- Managing RAID and Recovering from Failures **432**
- Standards-Based Storage Management **435**
- Managing Existing Partitions and Drives **443**

Storage management has changed substantially over the past few years, as have the ways Microsoft Windows Server works with disks. 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.

Generally, 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).

When I need to differentiate between the two storage-management approaches, I refer to the former technique as *traditional* and the latter technique as *standards-based*. In this chapter, I look at traditional techniques for creating volume sets and arrays first, and then I look at standards-based techniques for creating volumes. Whether a volume is created using the traditional approach or the standards-based approach, you manage it using similar techniques. For this

reason, in the final section of this chapter, I discuss techniques for working with existing volumes and drives.

**REAL WORLD** Standards-based approaches to storage management can be used with a server's internal disks as well. When internal disks are used in this way, however, the internal disks—such as virtual disks on attached storage—are resources to be allocated using standards-based 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 also can enable data deduplication on your virtual disks. You can't, however, use the operating system's volume set or RAID array features. The reason for this is that standards-based, storage-management approaches rely on the storage subsystem to manage the physical disk architecture.

## Using Volumes and Volume Sets

---

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 using the built-in features of the operating system (a software approach) or by using hardware. Windows Server 2012 supports three levels of software RAID: 0, 1, and 5. RAID arrays are implemented as mirrored, striped, and striped with parity volumes.

You create volume sets and RAID arrays on dynamic drives, which are accessible only by Microsoft Windows 2000 and later releases. However, computers running earlier versions of Windows can access the drives over the network, just as they can any other network drive.

You create and manage volumes in much the same way you create and manage partitions. A *volume* is a drive section you can use to store data directly.

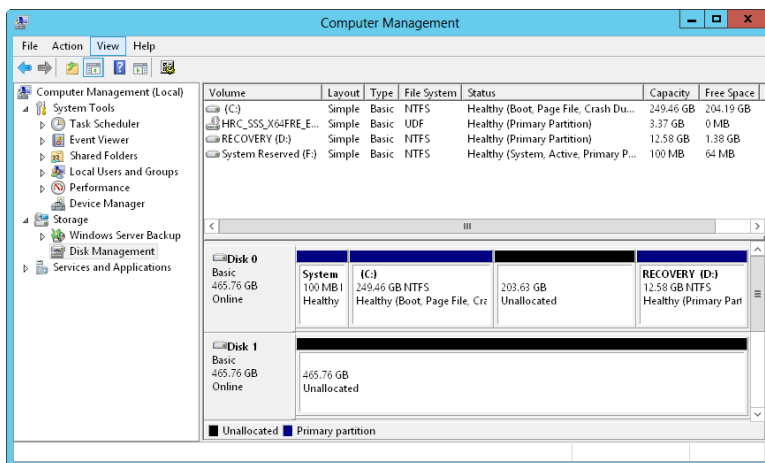
**NOTE** With spanned and striped volumes on basic disks, you can delete a volume but you can't create or extend volumes. With mirrored volumes on basic disks, you can delete, repair, and resync the mirror. You can also break the mirror. For striped with parity volumes (RAID-5) on basic disks, you can 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. As Figure 11-1 shows, 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*.

- **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.
- **Capacity** The total storage size of the drive.
- **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.



**FIGURE 11-1** Disk Management displays volumes much like it does partitions.

An important advantage of dynamic volumes over basic volumes is that they let you 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 2012. You can install other operating systems and dual boot a Windows Server 2012 system. To do this, you should create a separate volume for the other operating system. For example, you could install Windows Server 2012 on volume C and Windows 8 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 this chapter
- 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 10, “Managing File Systems and Drives”

## Understanding Volume Sets

With volume sets, you can create volumes that span several drives. To do this, you use 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. Table 11-1 summarizes status values for dynamic volumes.

**TABLE 11-1** Understanding and Resolving Volume Status Issues

STATUS	DESCRIPTION	RESOLUTION
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 volume, 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, press and hold or right-click the volume and then tap or 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.

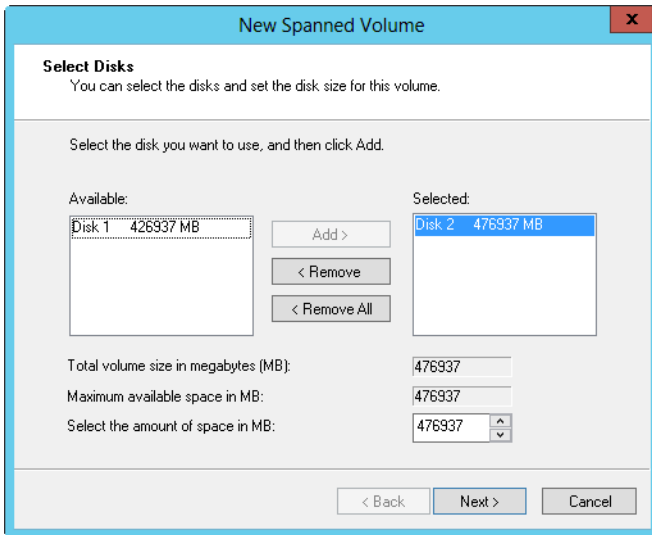
<b>STATUS</b>	<b>DESCRIPTION</b>	<b>RESOLUTION</b>
Formatting	A temporary status that indicates the volume is being formatted.	The progress of the formatting is indicated as the percent complete unless you choose the Perform A Quick Format option.
Healthy	The normal volume status.	The volume doesn't 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.	Press and hold or right-click the volume, and then tap or 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. NTFS formatting allows 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. You do this by selecting an area of free space and adding it to the volume. You can extend a simple volume within the same disk. 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, press and hold or right-click an unallocated area and then tap or click New Spanned Volume or New Striped Volume as appropriate. Read the Welcome page, and then tap or click Next.
2. You should see the Select Disks page, shown in Figure 11-2. Select disks that you want to be part of the volume, and size the volume segments on those disks.



**FIGURE 11-2** On the Select Disks page, select disks to be a part of the volume, and then size the volume on each disk.

3. Available disks are shown in the Available list. If necessary, select a disk in this list and then tap or 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 tapping or clicking Remove.

**CAUTION** The disk wizards in Windows Server 2012 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 tapping or 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. Tap or 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 tap or click Next. You use these 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 tap 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. The NTFS file system is the only option within Disk Management.
  - **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. 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.
  - **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 10.
7. Tap or click Next, and then tap or 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 2012.

To delete volumes, follow these steps:

1. In Disk Management, press and hold or right-click any volume in the set and then tap or click Delete Volume. You can't delete a portion of a spanned volume without deleting the entire volume.
2. Tap or 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 "Managing Existing Partitions and Drives" later in this chapter.

## Improving Performance and Fault Tolerance with RAID

---

You'll often want to give important data increased protection from drive failures. To do this, you can use RAID technology to add fault tolerance to your file systems. 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. These implementations are described in terms of levels. Currently, RAID levels 0 to 5 are defined. Each RAID level offers different features. Windows Server 2012 supports RAID levels 0, 1, and 5. You can use RAID-0 to improve the performance of your drives. You use RAID-1 and RAID-5 to provide fault tolerance for data.

Table 11-2 provides a brief overview of the supported RAID levels. This support is completely software-based.

The most common RAID levels in use on servers running Windows Server 2012 are level 1 (disk mirroring), and level 5 (disk striping with parity). With respect to up-front 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.



**TABLE 11-2** Windows Server 2012 Support for RAID

RAID LEVEL	RAID TYPE	DESCRIPTION	MAJOR ADVANTAGES
0	Disk striping	Two or more volumes, each on a separate drive, are configured as a striped set. Data is broken into blocks, called <i>stripes</i> , and then written sequentially to all drives in the striped set.	Speed and performance.
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.)	Redundancy. Better write performance than disk striping with parity.
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.	Fault tolerance with less overhead than mirroring. Better read performance than disk mirroring.

**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 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 on Windows Server 2012

Windows Server 2012 supports disk mirroring, disk striping, and disk striping with parity. Implementing these RAID techniques is discussed in the sections that follow.

**CAUTION** Some operating systems, such as MS-DOS, don't support RAID. If you dual boot your system to one of these noncompliant operating systems, your RAID-configured drives will be unavailable.

### 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 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. Data backup and recovery is discussed in Chapter 13, “Data Backup and Recovery.”

**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 the smallest volume is 20 GB and you want to create a three-volume striped set, the maximum size for the striped set is 60 GB.

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, press and hold or right-click an area marked Unallocated on a dynamic disk, and then tap or click New Striped Volume. This starts the New Striped Volume Wizard. Read the Welcome page, and then tap or click Next.
2. Create the volume as described in “Creating Volumes and Volume Sets” earlier in this chapter. The key difference is that you need at least two dynamic disks to create a striped volume.

After you create a striped volume, you can use the volume as you would any other volume. You can't extend a striped set once it's created. Therefore, you should carefully consider the setup before you implement it.

## Implementing RAID-1: Disk Mirroring

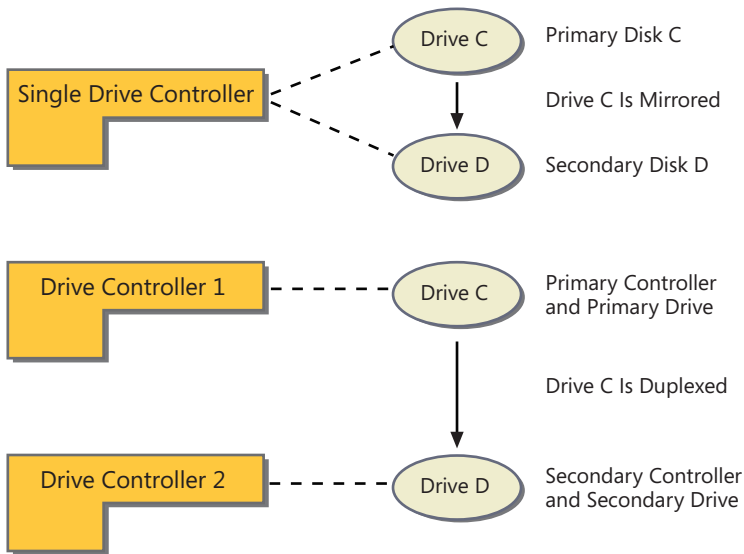
RAID level 1 is disk mirroring. With disk mirroring, you use identically sized volumes on two different drives to create a redundant data set. The drives are written with identical sets of information, and if one of the drives fails, you can still obtain the data from the other drive.

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 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. This provides increased protection against failure of the disk controller. If one of the disk controllers fails, the disk on 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*. Figure 11-3 shows the 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.



**FIGURE 11-3** Although disk mirroring typically uses a single drive controller to create a redundant data set, disk duplexing uses two drive controllers.

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, press and hold or right-click an area marked Unallocated on a dynamic disk, and then tap or click New Mirrored Volume. This starts the New Mirrored Volume Wizard. Read the Welcome page, and then tap or click Next.
2. Create the volume as described in “Creating Volumes and Volume Sets” earlier in this chapter. The key difference 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 you want to work with.

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.

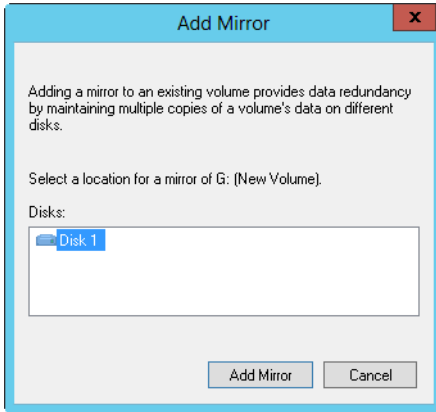
**NOTE** The status of a normal mirror is Healthy. During the creation of a mirror, you’ll see a status of Resyncing, which tells you that Disk Management is creating the mirror.

## 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 dynamic drive of equal or larger size than the existing volume.

In Disk Management, you mirror an existing volume by following these steps:

1. Press and hold or right-click the simple volume you want to mirror, and then tap or click Add Mirror. This displays the Add Mirror dialog box.
2. In the Disks list, shown in Figure 11-4, select a location for the mirror, and then tap or click Add Mirror. Windows Server 2012 begins the mirror creation process. In Disk Management, you’ll see a status of Resyncing on both volumes. The disk on which the mirrored volume is being created has a warning icon.



**FIGURE 11-4** Select the location for the mirror.

## 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 is essentially an enhanced version of RAID-1, with the key addition of fault tolerance. 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, press and hold or right-click an area marked Unallocated on a dynamic disk, and then tap or click New RAID-5 Volume. This starts the New RAID-5 Volume Wizard. Read the Welcome page, and then tap or click Next.
2. Create the volume as described previously in "Creating Volumes and Volume Sets." The key difference 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. This ensures that if you have problems, you can recover your data.

In Disk Management, you can break a mirrored set by following these steps:

1. Press and hold or right-click one of the volumes in the mirrored set, and then tap or click Break Mirrored Volume.
2. Confirm that you want to break the mirror by tapping or clicking Yes. If the volume is in use, you'll see another warning dialog box. Confirm that it's okay to continue by tapping or clicking Yes.

Windows Server 2012 breaks the mirror, creating two independent volumes.

## Resynchronizing and Repairing a Mirrored Set

Windows Server 2012 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 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, press and hold or right-click the failed volume, and tap or click Reactivate Volume. The drive status should change to Regenerating and then to Healthy. If the volume doesn't return to the Healthy status, press and hold or right-click the volume, and then tap or click Resynchronize Mirror.

- If the status is Online (Errors), press and hold or right-click the failed volume, and then tap or click Reactivate Volume. The drive status should change to Regenerating and then to Healthy. If the volume doesn't return to the Healthy status, press and hold or right-click the volume, and then tap or 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, press and hold or right-click the failed volume, and then tap or click Remove Mirror. Next, press and hold or right-click the remaining volume in the original mirror, and then tap or 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 booting. 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 2012, 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 boot 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 boot 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 boot to the primary system volume, restart the system and select the Windows Server 2012—Secondary Plex option for the operating system you want to start. The system should start up normally. After you successfully boot the system to the secondary drive, you can schedule the maintenance necessary to rebuild the mirror. You need to follow these 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. Press and hold or right-click the remaining volume that was part of the original mirror, and then tap or 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. Press and hold or right-click the original system volume, and then tap or click Add Mirror. Now re-create 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, press and hold or right-click one of the volumes in the mirrored set, and then tap or click Remove Mirror. This displays 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 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, press and hold or 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, press and hold or right-click the volume and select Regenerate Parity.



- If the status is Online (Errors), press and hold or 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, press and hold or 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 reboot 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. Press and hold or 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.

**BEST PRACTICES** If possible, you should back up the data before you perform this procedure. This ensures that if you have problems, you can recover your data.

## Standards-Based Storage Management

---

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 Standards-Based Storage

With standards-based management, the physical layout of disks (spindles) is abstracted. Here, 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 *storage 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** Trust me when I say this all sounds more complicated than it is. When you throw storage subsystems into the mix, it's really a three-layered 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 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.

## Working with Standards-Based Storage

To use standards-based storage management, 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:

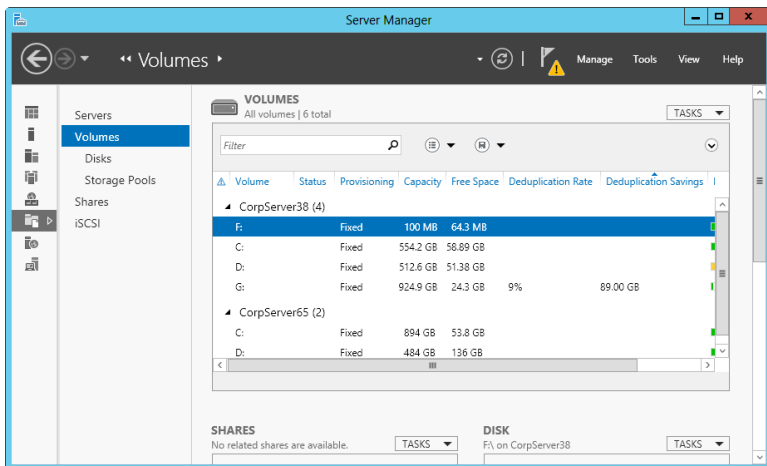
- Add the Data Deduplication role service if you want to enable data deduplication.
- Add the iSCSI Target Server and iSCSI Target Storage Provider role services if you want the server to host iSCSI virtual disks.

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 there will be additional options as well. The Servers subnode lists file servers that have been configured for standards-based management.

As Figure 11-5 shows, the Volumes subnode lists allocated storage on each server according to how volumes are provisioned and how much free space each volume has. Volumes are listed regardless of whether the underlying disks are physical or virtual. Press and hold or 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.

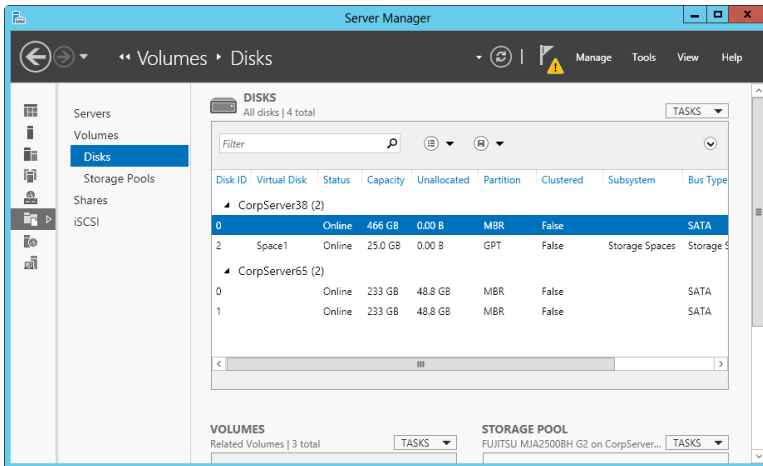
- **New Share** Allows you to create new Server Message Block (SMB) or Network File System (NFS) shares on the volume.
- **Properties** Displays information about the volume type, file system, health, capacity, used space, and free space. You also can use this option to set the volume label.
- **Repair 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.



**FIGURE 11-5** Note how volumes are provisioned.

As Figure 11-6 shows, 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. Press and hold or right-click a disk to display management options, including the following:

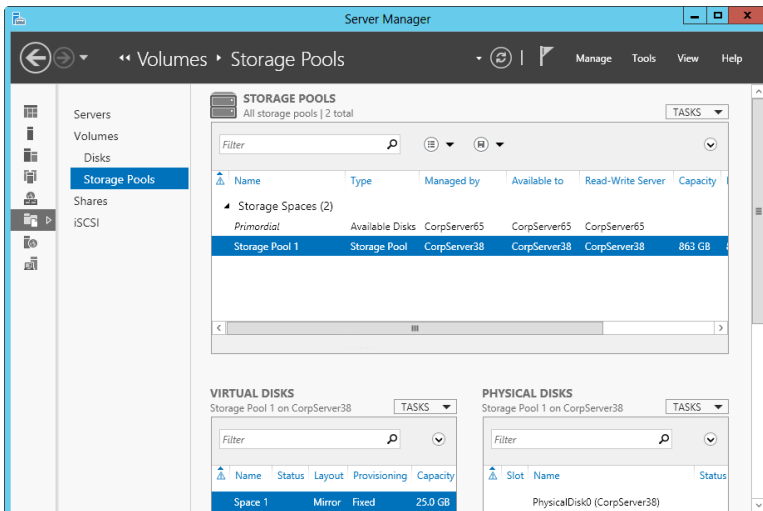
- **Bring Online** Allows you to take an offline disk and make it available for use.
- **Take Offline** Allows you to take a disk offline so that it can no longer be used.
- **Reset Disk** Allows you to completely reset the disk, which deletes all volumes on the disk and removes all available data on the disk.
- **New Volume** Allows you to create a new volume on the disk.



**FIGURE 11-6** Note the disks available and how much unallocated space is available.

## Creating Storage Pools and Allocating Space

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. As Figure 11-7 shows, 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.



**FIGURE 11-7** Create and manage storage pools.

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 sections that follow examine procedures related to each of these steps.

## Creating a Storage Space

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.

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

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. Tap or click Tasks in the Storage Pools panel, and then tap or click 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 tap or click Next.
3. On the Specify A Storage Pool Name And Subsystem page, type a name and description of the storage pool. Then select the primordial pool you want to work with. A primordial pool is simply a group of disks managed by and available to a specific server via a storage subsystem. Tap or click 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 CorpServer 38.

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 Data Store to allocate the disk to the pool and make it available for use.
  - 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.
5. When you are ready to continue, tap or click Next. After you confirm your selections, tap or click 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 tap or click Close.
- If any portion of the configuration failed, note the reason for the failure and take corrective actions as appropriate before repeating this procedure.

## Creating a Virtual Disk in a Storage Space

After you create a storage pool, you can allocate space from the pool to virtual disks that are available to your servers. 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 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. Here, 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 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. Tap or click Tasks in the Virtual Disks panel, and then tap or click New Virtual Disk. This starts the New Virtual Disk Wizard.

3. On the Storage Pool page, tap or click the storage pool in which you want to create the virtual disk and then tap or click 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 CorpServer 38.

4. On the Specify The Virtual Disk Name page, type a name and description for the virtual disk. Tap or click Next.
5. 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. Tap or click Next.
6. 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. Tap or click Next.
7. On the Specify The Size Of The Virtual Disk page, use the options provided to set the size of the virtual disk. By selecting the Create The Largest Virtual Disk Possible check box, you ensure the disk is created and sized within the available space. For example, if you are trying to create a 2-TB fixed disk with a simple layout and only 1.5 TB of space is available, a 1.5-TB fixed disk will be created. Keep in mind that if a disk is mirrored or striped, it will use more free space than you specify.
8. When you are ready to continue, tap or click Next. After you confirm your selections, tap or click 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.
9. When you tap or click Close, the New Volume Wizard should start automatically. Use the wizard to create a volume on the disk as discussed in "Creating a Standard Volume."

## 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 you want to work with, 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 you want to work with, 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 tap or click 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 tap or click 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. Tap or click 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 tap or click 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 type the path to an existing folder on an NTFS drive, or tap or click 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 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.
6. 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, tap or click Next.



7. After you confirm your selections, tap or click 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. Tap or click Close.

## Managing Existing Partitions and Drives

---

Disk Management provides many ways to manage existing partitions and drives. Use these features to assign drive letters, delete partitions, set the active partition, and more. In addition, Windows Server 2012 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.

**NOTE** Windows Vista as well as all later releases of Windows support hot-pluggable media that use NTFS volumes. This new feature allows you to format USB flash devices and other similar media with NTFS. There are also enhancements to prevent data loss when ejecting NTFS-formatted removable media.

### 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.

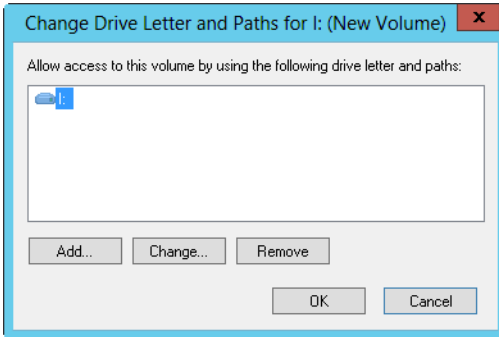
Windows cannot modify the drive letter of system, boot, or page-file volumes. To change the drive letter of a system or boot volume, you need to edit the registry as described in Microsoft Knowledge Base article 223188 ([support.microsoft.com/kb/223188/](http://support.microsoft.com/kb/223188/)). Before you can change the drive letter of a page-file volume, you might need to move the page file to a different volume.

To manage drive letters and paths, press and hold or right-click the drive you want to configure in Disk Management, and then tap or click Change Drive Letter And Paths. This opens the dialog box shown in Figure 11-8. You can now do the following:

- **Add a drive path** Tap or click Add, select Mount In The Following Empty NTFS Folder, and then type the path to an existing folder, or tap or click Browse to search for or create a folder.
- **Remove a drive path** Select the drive path to remove, tap or click Remove, and then tap or click Yes.
- **Assign a drive letter** Tap or 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 tap or 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, tap or click Remove, and then tap or click Yes.

**NOTE** If you try to change the letter of a drive that's in use, Windows Server 2012 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 tapping or clicking Yes when prompted.



**FIGURE 11-8** You can change the drive letter and path assignment in the Change Drive Letter And Paths dialog box.

## 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, the volume label can be up to 32 characters. Additionally, although FAT doesn't allow you to use some special characters—including \* / \ [ ] ; | = , . + " ? < >—NTFS does allow you to use these special characters.

Because the volume label is displayed when the drive is accessed in various Windows Server 2012 utilities, including File Explorer, it can provide information about a drive's contents. You can change or delete a volume label using Disk Management or File Explorer.

Using Disk Management, you can change or delete a label by following these steps:

1. Press and hold or right-click the partition, and then tap or click Properties.
2. On the General tab of the Properties dialog box, type a new label for the volume in the Label text box or delete the existing label. Tap or click OK.

Using File Explorer, you can change or delete a label by following these steps:

1. Press and hold or right-click the drive icon, and then tap or click Properties.
2. On the General tab of the Properties dialog box, type a new label for the volume in the Label text box or delete the existing label. Tap or click OK.

## 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 2012 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, press and hold or right-click the partition, volume, or drive you want to delete, and then tap or click Explore. Using File Explorer, move all the data to another volume or verify an existing backup to ensure the data was properly saved.
2. In Disk Management, press and hold or right-click the partition, volume, or drive again, and then tap or click Delete Partition, Delete Volume, or Delete Logical Drive as appropriate.
3. Confirm that you want to delete the selected item by tapping or 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:

1. Delete all the logical drives on the partition following the steps listed in the previous procedure.
2. Select the extended partition area itself and delete it.

## Converting a Volume to NTFS

Windows Server 2012 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 using this tool, the file and directory structure is preserved and no data is lost. Keep in mind, however, that Windows Server 2012 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 to re-create the partition as a FAT volume.

## 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. You are not prompted for a volume label 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:

<code>volume</code>	Sets the volume to work with
<code>/FS:NTFS</code>	Converts to NTFS
<code>/V</code>	Sets verbose mode
<code>/X</code>	Forces the volume to dismount before the conversion (if necessary)
<code>/CvtArea: filename</code>	Sets the name of a contiguous file in the root directory to be a placeholder for NTFS system files
<code>/NoSecurity</code>	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 2012 displays a prompt asking if you want to schedule the drive to be converted the next time the system starts. If you tap or 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.

## Resizing Partitions and Volumes

Windows Server 2012 doesn't use `Ntldr` and `Boot.ini` to load the operating system. Instead, Windows Server 2012 has a preboot environment in which Windows Boot Manager is used to control startup and load the boot application you selected.

Windows Boot Manager also finally frees the Windows operating system from its reliance on MS-DOS so that you can use drives in new ways. With Windows Server 2012, you can extend and shrink both basic and dynamic disks. You can use either Disk Management or DiskPart to extend and shrink volumes. You cannot shrink or extend striped, mirrored, or striped-with-parity volumes.

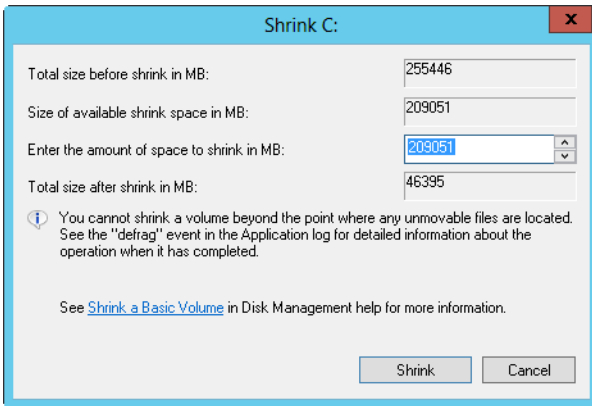
In extending a volume, you convert areas of unallocated space and add them to the existing volume. For spanned volumes on dynamic disks, the space can come from any available dynamic disk, not only from those on which the volume was

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. First, you can extend simple and spanned volumes only if they are formatted and the file system is NTFS. You can't extend striped volumes. You can't extend volumes that aren't formatted or 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, press and hold or right-click the volume you want to shrink and then tap or 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, shown in Figure 11-9, enter the amount of space to shrink.



**FIGURE 11-9** Specify the amount of space to shrink from the volume.

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. Tap or click Shrink to shrink the volume.  
You can extend a simple volume or a spanned volume by following these steps:
    1. In Disk Management, press and hold or right-click the volume you want to extend and then tap or 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 tap or 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.
    4. With dynamic disks, you can specify the additional space you want to use on other disks by performing the following tasks:
      - Tap or click the disk, and then tap or 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. Tap or click Next, confirm your options, and then tap or click Finish.

## Repairing Disk Errors and Inconsistencies Automatically

Windows Server 2012 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. This 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 self-healing 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 2012 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 2012 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. 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 2012 performs a proactive scan of NTFS volumes. As with other automated maintenance, Windows scans disks using Check Disk at 3:00 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 3:00 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.

## Checking Disks Manually

With Windows Server 2012, 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 allows Windows to rapidly repair volumes. It also allows you to 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 2012 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 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 typing 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. Unless you specify further options, Check Disk won't repair problems, however. 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 see a prompt asking whether you would like to force a dismount of the volume for the repair. In this case, you can type **Y** to proceed or **N** to cancel the dismount. If you cancel the dismount, you'll see 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 type **Y** to schedule the repair or **N** to cancel the repair.
- For system volumes, you'll see 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 type **Y** to schedule the repair or **N** to cancel the repair.

You can't run Check Disk with both the `/scan` and `/spotfix` options. The reason for this is that the scan and repair tasks are now independent of each other.

The complete syntax for Check Disk is shown here:

```
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:

<code>volume</code>	Sets the volume to work with.
<code>[path]filename</code>	FAT only. It specifies files to check for fragmentation.
<code>/B</code>	Reevaluates bad clusters on the volume (NTFS only; implies <code>/R</code> ).
<code>/C</code>	NTFS only. It skips the checking of cycles within the folder structure.
<code>/F</code>	Fixes errors on the disk using the offline (legacy) scan and fix behavior.
<code>/I</code>	NTFS only. It performs a minimum check of index entries.
<code>/L:size</code>	NTFS only. It changes the log file size.
<code>/R</code>	Locates bad sectors, and recovers readable information (implies <code>/F</code> ).
<code>/V</code>	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.
<code>/X</code>	Forces the volume to dismount first if necessary (implies <code>/F</code> ).

For NTFS volumes, Check Disk supports these enhanced options:


<code>/forceofflinefix</code>	Must be used with <code>/scan</code> . It bypasses all online repair and queues errors for offline repair.
<code>/offlinescanandfix</code>	Performs an offline scan and fix of the volume.
<code>/perf</code>	Performs the scan as fast as possible using more system resources.
<code>/scan</code>	Performs an online scan of the volume (the default). Errors detected during the scan are added to the <code>\$corrupt</code> system file.

<code>/sdcleanup</code>	Cleans up unneeded security descriptor data. It implies <code>/F</code> (with legacy scan and repair).
<code>/spotfix</code>	Allows certain types of errors to be repaired online (the default).

## Running Check Disk Interactively

You can run Check Disk interactively by using File Explorer or Disk Management. Follow these steps:

1. Press and hold or right-click the drive, and then tap or click Properties.
2. On the Tools tab of the Properties dialog box, tap or click Check. This displays the Error Checking dialog box, shown in Figure 11-10.



**FIGURE 11-10** Use Check Disk to check a disk for errors and repair any that are found.

3. Click Scan Drive to start the scan. If no errors are found, Windows confirms this. 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 tap or 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

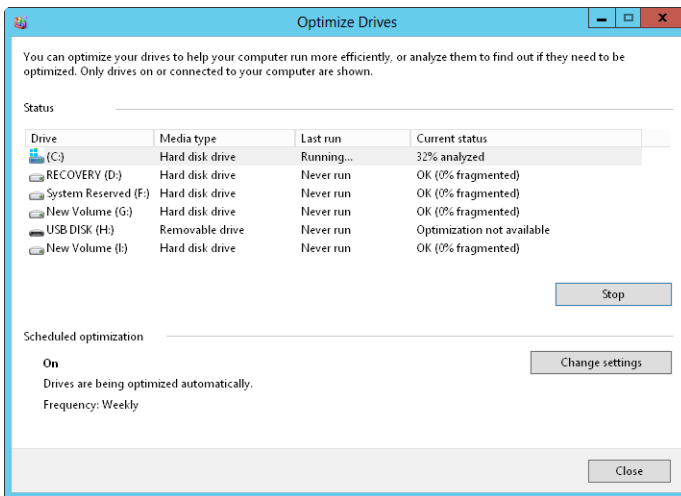
Any time you add files to or remove files from a drive, the data on the drive can become fragmented. When a drive is fragmented, large files can't be written to a single continuous area on the disk. As a result, the operating system must write the file to several smaller areas on the disk, which means more time is spent reading the file from the disk. To reduce fragmentation, Windows Server 2012 can manually or automatically analyze and optimize disks 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 Computer Management, select the Storage node and then the Disk Management node. Press and hold or right-click a drive, and then tap or click Properties.
2. On the Tools tab, tap or click Optimize. In the Optimize Drives dialog box, select a disk and then tap or click Analyze. Optimize Drives then analyzes the disk, as shown in Figure 11-11, 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 tap or click Optimize.

**NOTE** Depending on the size of the disk, defragmentation can take several hours. You can tap or click Stop Operation at any time to stop defragmentation.



**FIGURE 11-11** Optimize Drives analyzes and defragments disks efficiently.

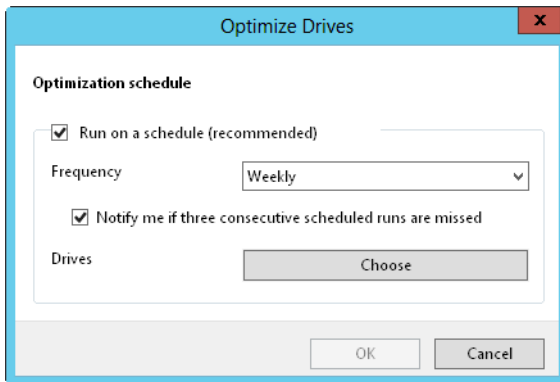
Automatic analysis and optimization of disks can occur while the disks are online, so long as the computer is on AC power and the operating system is running but otherwise idle. By default, disk optimization is a weekly task rather than a daily task—and there’s a good reason for this. Normally, you need 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 2012 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. Press and hold or right-click a drive, and then tap or click Properties.
2. On the Tools tab, tap or click Optimize. This displays the Optimize Drives dialog box.
3. If you want to change how optimization works, tap or click Change Settings. This displays the dialog box shown in Figure 11-12. To cancel automated defragmentation, clear the Run On A Schedule check box. To enable automated defragmentation, select Run On A Schedule.



**FIGURE 11-12** Set the run schedule for automated defragmentation.

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, tap or 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. Tap or click OK to save your settings.
6. Tap or click OK, and then tap or click Close.



# Data Sharing, Security, and Auditing

- Using and Enabling File Sharing **458**
- Configuring Standard File Sharing **461**
- Managing Share Permissions **469**
- Managing Existing Shares **474**
- Configuring NFS Sharing **479**
- Using Shadow Copies **481**
- Connecting to Network Drives **484**
- Object Management, Ownership, and Inheritance **485**
- File and Folder Permissions **488**
- Auditing System Resources **498**
- Using, Configuring, and Managing NTFS Disk Quotas **503**
- Using, Configuring, and Managing Resource Manager Disk Quotas **513**

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. Windows 8 and Windows Server 2012 support SMB version 3.0 and include an SMB 3.0-compatible client.

SMB 3.0 brings many enhancements for performance, especially when you use clustered file servers. A key enhancement that doesn't rely on a special configuration is end-to-end encryption of SMB data, which eliminates the need to use Internet Protocol security (IPsec), specialized hardware, or wide area network (WAN) accelerators to protect data from eavesdropping. SMB encryption can be enabled on a per-share basis.

With SMB, Windows Server 2012 supports two file-sharing models: *standard file sharing* and *public folder sharing*. Standard file sharing allows 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 allows 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.

With public folder sharing, you share files simply by copying or moving files to the computer's Public folder. Public files are available to anyone who logs on to a computer locally regardless of whether that person has a standard user account or an administrator user account on the computer. You can also grant network access to the Public folder. If you do this, however, there are no access restrictions. The Public folder and its contents are open to everyone who can access the computer over the local network.

## Using and Enabling File Sharing

---

The sharing settings on a computer determine the way files can be shared. The two file-sharing models that Windows Server 2012 supports have the following differences:

- **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.
- **Public folder sharing** Allows local users and (optionally) remote users to access any files placed in the computer's %SystemDrive%\Users\Public folder. Access permissions on the Public folder determine which users and groups have access to publicly shared files as well as the level of access those users and groups have. When you copy or move files to the Public folder, access permissions are changed to match those of the Public folder. Some additional permissions are added as well. When a computer is part of a workgroup, you can add password protection to the Public folder. Separate password protection isn't needed in a domain. In a domain, only domain users can access Public folder data.

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. With public folder sharing, on the other hand, files copied or moved to the Public folder are available to anyone who logs on locally. You can grant network access to the Public folder as well. Doing so, however, makes



the Public folder and its contents open to everyone who can access the computer over the network.

Windows Server 2012 adds new layers of security through compound identities, claims-based access controls, and central access policies. With both Windows 8 and Windows Server 2012, you can assign claims-based access controls to file and folder resources on NTFS and ReFS volumes. With Windows Server 2012, 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.

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. Both Windows 8 and Windows Server 2012 support SMB 3.0. (They have an SMB 3.0 client.)

**REAL WORLD** Although ReFS provides a highly reliable file system, keep in mind that ReFS does not support shadow copies. Therefore, if you create shares on ReFS volumes, users won't be able to go back to previous versions of files and folders stored in shares.

Public folder sharing is designed to allow 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. In File Explorer, double-tap or double-click the system drive and then access 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).

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:

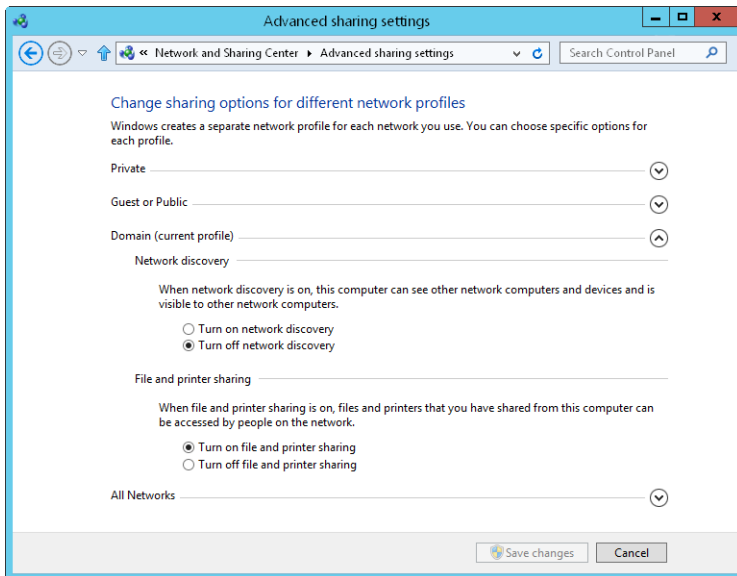
- Allow users logged on to the computer to view and manage public files but restrict network users from accessing public files. When you configure this

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 2012 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.

As shown in Figure 12-1, 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.



**FIGURE 12-1** Network And Sharing Center shows the current sharing configuration.

You can manage a computer's sharing configuration by following these steps:

1. In Control Panel, tap or click View Network Status And Tasks under the Network And Internet heading. This opens Network And Sharing Center.

2. In Network And Sharing Center, tap or 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 Public Folder Sharing panel by tapping or clicking the related expand button. 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. Tap or 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 using Distributed File System (DFS). Server Manager displays information about standard SMB folder shares, SMB folders shared using DFS, and folders shared using 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. You're connected to the local computer by default. If you want to connect to a remote computer, press and hold or right-click the Computer Management node and then tap or click Connect To Another Computer. Choose Another Computer, type the name or IP address of the computer you want to connect to, and then tap or 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, as shown in Figure 12-2.

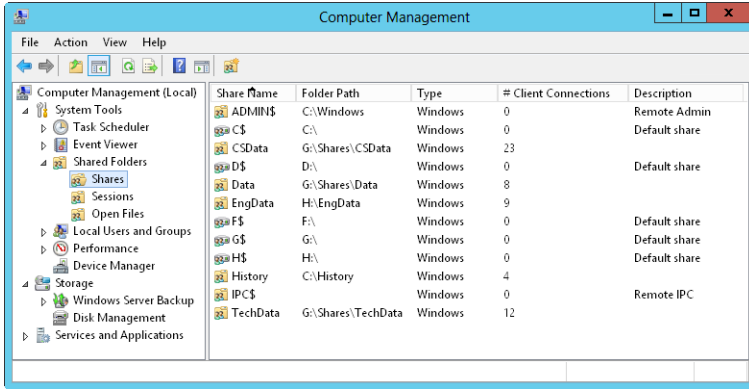


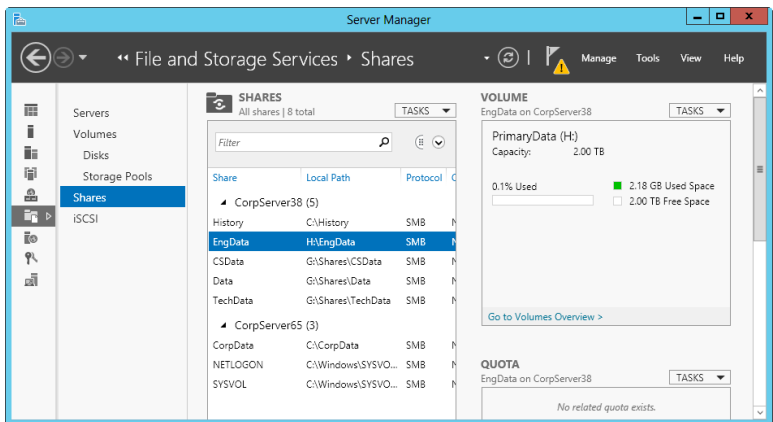
FIGURE 12-2 Available shares are listed in the Shared Folders node.

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:

1. Select the File And Storage Services node, and then select the related Shares subnode.
2. As Figure 12-3 shows, 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.



**FIGURE 12-3** Tap or click Shares in the main pane (on the left) to view the available shares.

3. When you tap or click a share in the Shares pane, information about the related volume is displayed in the Volume pane.

**REAL WORLD** Network File System (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 Shared Folders in Computer Management

Windows Server 2012 provides several ways to share folders. You can share local folders using File Explorer, and you can share local and remote folders 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

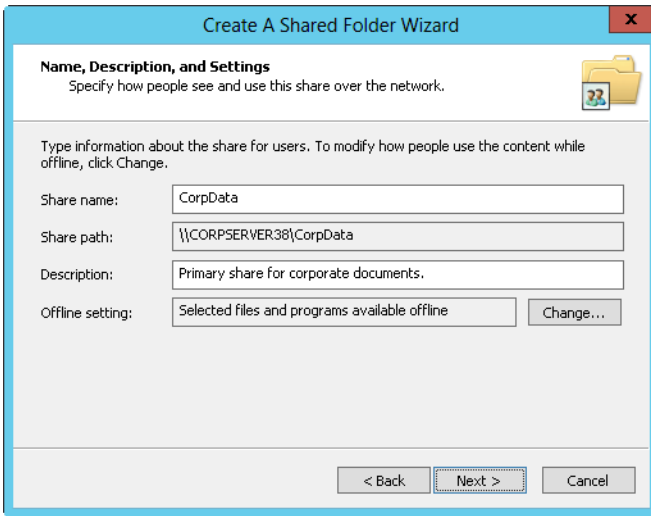
In Computer Management, you share a folder by following these steps:

1. If necessary, connect to a remote computer. In the console tree, expand System Tools, expand Shared Folders, and then select Shares. The current shares on the system are displayed.
2. Press and hold or right-click Shares, and then tap or click New Share. This starts the Create A Shared Folder Wizard. Tap or click Next.
3. In the Folder Path text box, type the local file path to the folder you want to share. The file path must be exact, such as **C:\EntData\Documents**. If you don't know the full path, tap or click Browse, use the Browse For Folder

dialog box to find the folder you want to share, and then tap or click OK. Tap or click Next.

**TIP** If the file path you specified doesn't exist, the wizard can create it for you. Tap or click Yes when prompted to create the necessary folder or folders.

4. In the Share Name text box, type a name for the share, as shown in Figure 12-4. This is the name of the folder to which users will connect. Share names must be unique for each system.



**FIGURE 12-4** Use the Create A Shared Folder Wizard to configure the essential share properties, including name, description, and offline resource usage.

**TIP** If you want to hide a share from users (which means that they won't be able to see the shared resource when they try to browse to it in File Explorer or at the command line), type 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.

5. If you want to, type 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.
6. 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 allows users to take advantage of the new Always Offline feature. If you want to use different offline file settings, tap or click Change, select the appropriate options in the Offline Settings dialog

box, and then tap or 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 Share 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.
7. Tap or 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 only view files and read data. 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."
8. When you tap or 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. Tap or 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 4, when you name the share, make sure that you use a different name.

- In step 5, 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 Shared Folders 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.
2. In the Shares pane, tap or click Tasks and then tap or click New Share. This starts the New Share Wizard. Choose one of the available file share profiles and then tap or click Next. The New Share Wizard has several file share profiles:
  - **SMB Share—Quick** A basic profile for creating SMB file shares that allows you to configure their settings and permissions.
  - **SMB Share—Advanced** An advanced profile for creating SMB file shares that allows you to configure their settings, permissions, management properties, and NTFS quota profile (if applicable).
  - **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.

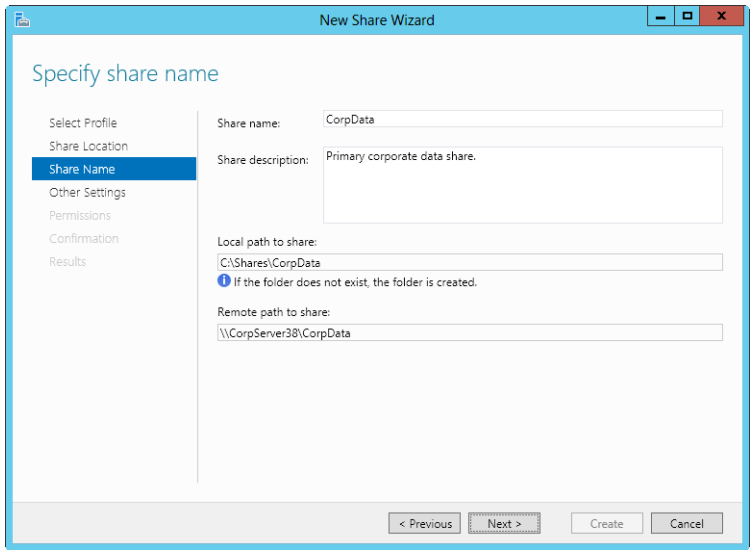
**NOTE** If you are using the Server For NFS role service, options are available for creating NFS shares as well.

**REAL WORLD** SMB 3.0 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 3.0, packets use large Maximum Transmission Units (MTUs) as well, which enhances 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 and volume on which you want the share to be created. Only file servers you've added for management are available. When you are ready to continue, tap or click Next.

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 type the desired share path, such as C:\Data, or click Browse to use the Select Folder dialog box to select the share path.
4. On the Specify Share Name page, type a name for the share, as shown in Figure 12-5. This is the name of the folder to which users will connect. Share names must be unique for each system.





**FIGURE 12-5** Set the name and description for the share.

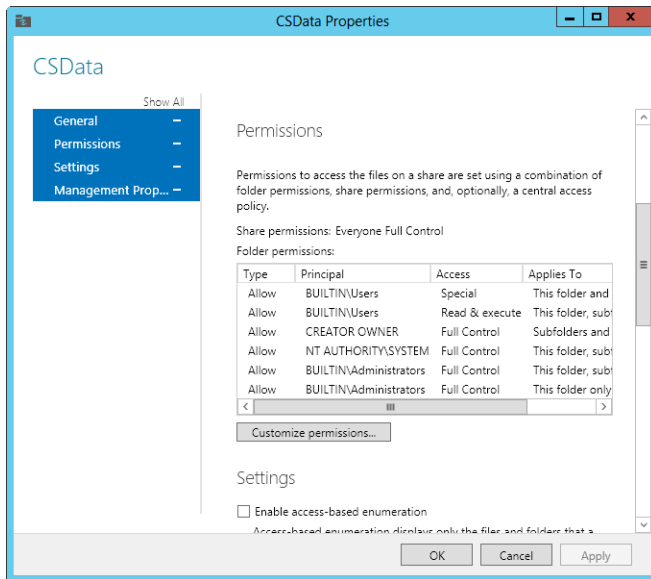
5. If you want to, type a description of the share in the Description text box. When you view shares on a particular computer, the description is displayed in Computer Management.
6. Note the local and remote paths to the share. These paths are set based on the share location and share name you specified. When you are ready to continue, tap or click Next.
7. 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.
8. 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, tap or click Customize Permissions and then use the Advanced Security Settings dialog box to configure the desired permissions. Setting share permissions is discussed fully in “Managing Share Permissions.” Setting folder permissions is discussed fully in “Understanding File and Folder Permissions” later in the chapter.
- 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.
9. If you are using the advanced profile, optionally set the folder management properties and then tap or 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.
  10. If you are using the advanced profile, optionally apply a quota based on a template to the folder and then tap or click Next. You can select only quota templates that have already been created. For more information, see “Managing Disk Quota Templates” later in this chapter.
  11. On the Confirm Selections page, review your selections. When you tap or 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. Tap or click Close.

## Changing Shared Folder 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.
2. Press and hold or right-click the share you want to work with, and then tap or click Properties.
3. In the Properties dialog box, shown in Figure 12-6, you have several options panels that can be accessed using controls in the left pane. You can expand the panels one by one or tap or click Show All to expand all the panels at the same time.
4. Use the options provided to modify the settings as necessary, and then tap or click OK. The options available are the same whether you use the basic, advanced, or applications profile to create the shared folder.



**FIGURE 12-6** Modify share settings using the options provided.

**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, press and hold or right-click the share in Computer Management, and then tap or 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 tap or 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, as well as share permissions, to further constrain actions within the share. With FAT volumes, share permissions control only access.

### 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.

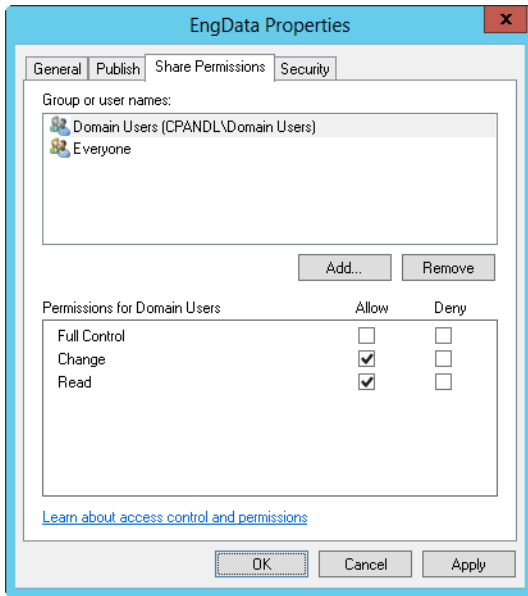
- **Read** With this permission, users can do the following:
  - View file and subfolder names
  - Access the subfolders in the share
  - Read file data and attributes
  - Run program files
- **Change** Users have Read permission and the ability to do the following:
  - Create files and subfolders
  - Modify files
  - Change attributes on files and subfolders
  - Delete files and subfolders
- **Full Control** Users have Read and Change permissions, as well as the following additional 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 “Implicit Groups and Special Identities” in Chapter 8, “Creating User and Group Accounts.”

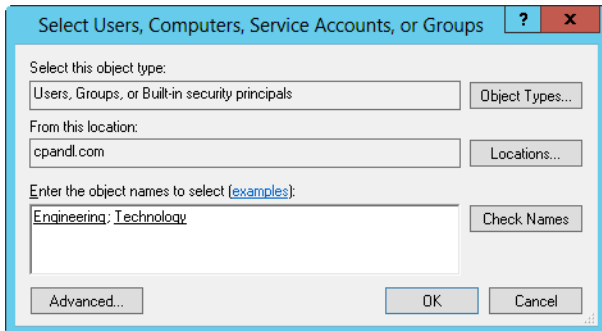
## Viewing and Configuring Share Permissions

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. Press and hold or right-click the share you want to work with, and then tap or click Properties.
3. In the Properties dialog box, tap or click the Share Permissions tab, shown in Figure 12-7. 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 tapping or clicking Remove. You can change permissions for these users and groups by doing the following:
  - a. Select the user or group you want to change.
  - b. Allow or deny access permissions in the Permissions list box.
5. To add permissions for another user or group, tap or click Add. This opens the Select Users, Computers, Service Accounts, Or Groups dialog box, shown in Figure 12-8.



**FIGURE 12-7** The Share Permissions tab shows which users and groups have access to the share and what type of access they have.



**FIGURE 12-8** Add users and groups to the share.

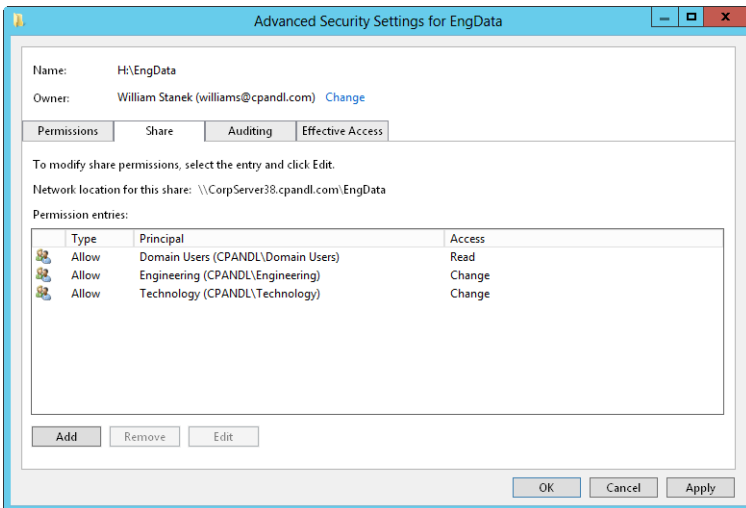
6. Type the name of a user, computer, or group in the current domain, and then tap or click Check Names. This produces one of the following results:
  - If a single match is found, the dialog box is automatically updated and the entry is underlined.
  - If no matches are found, you either entered an incorrect name part or you're working with an incorrect location. Modify the name and try again, or tap or click Locations to select a new location.
  - If multiple matches are found, select the name or names you want to use, and then tap or click OK. To assign permissions to other users, computers, or groups, type a semicolon (;) and then repeat this step.

**NOTE** The Locations button allows you to access account names in other domains. Tap or click Locations to see a list of the current domain, 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. Tap or 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. Tap or click OK. To assign additional security permissions for NTFS, see "File and Folder Permissions" later in this chapter.

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. Press and hold or right-click the share you want to work with, and then tap or click Properties.
3. In the Properties dialog box, tap or 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, tap or click Customize Permissions. Next, select the Share tab in the Advanced Security Settings dialog box, as shown in Figure 12-9.

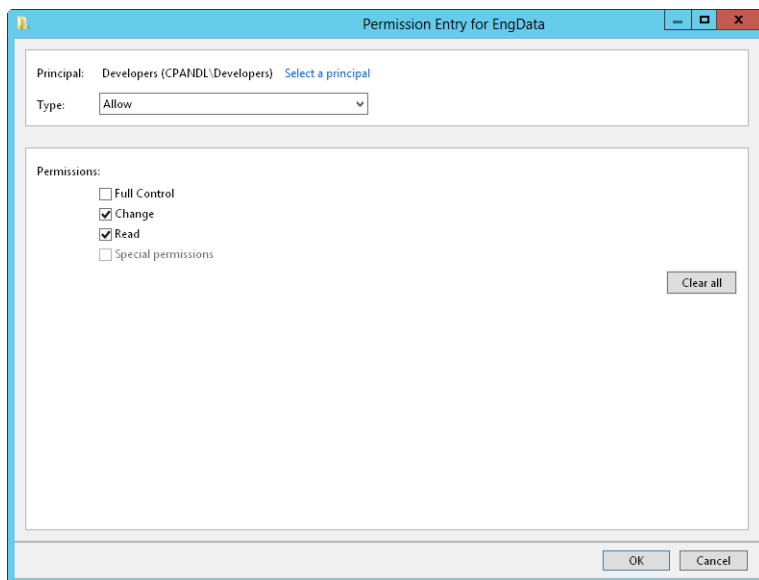


**FIGURE 12-9** The Share tab shows which users and groups have access to the share and what type of access they have.

5. Users or groups that already have access to the share are listed in the Permission Entries list. You can remove permissions for these users and groups by

selecting the user or group you want to remove and then tapping or clicking Remove. You can change permissions for these users and groups by doing the following:

- a. Select the user or group you want to change, and then select Edit.
  - b. Allow or deny access permissions in the Permission Entries list, and then tap or click OK.
6. To add permissions for another user or group, tap or click Add. This opens the Permission Entry dialog box, shown in Figure 12-10.



**FIGURE 12-10** Add permission entries for a particular user or group.

7. Tap or click Select A Principal to display the Select User, Computer, Service Account Or Group dialog box. Type 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.
8. Tap or 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 see 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 tap 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 tap or click OK.
9. Tap or click OK. The user and group is added as the Principal, and the Permission Entry dialog box is updated to show this.
10. 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.

11. Tap or click OK to return to the Advanced Security Settings dialog box. To assign additional security permissions for NTFS, see “File and Folder Permissions” later in this chapter.

## 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*. These shares 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\AutoShareServer
- HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\parameters\AutoShareWks

Which special shares are available depends on your system configuration. Table 12-1 lists special shares you might see and how they're used.

**TABLE 12-1** Special Shares Used by Windows Server 2012

SHARE NAME	DESCRIPTION	USAGE
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.

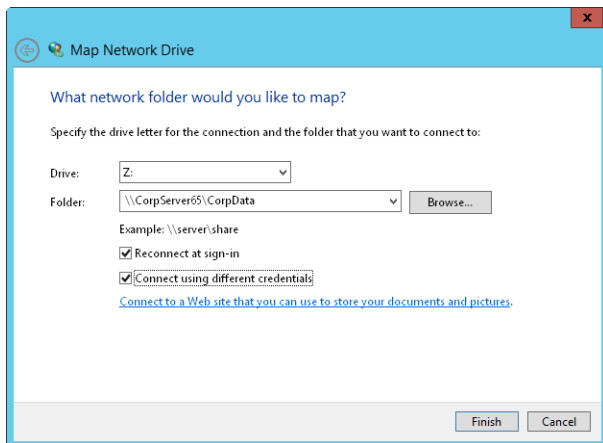


SHARE NAME	DESCRIPTION	USAGE
PRINT\$	Supports shared printer resources by providing access to printer drivers.	Used by shared printers. Everyone has Read access. Administrators, server operators, and printer operators have Full Control.
SYSVOL	Supports Active Directory.	Used to store data and objects for Active Directory.
<i>Driveletter\$</i>	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

Special shares end with the \$ symbol. Although these shares aren't displayed in File Explorer, administrators and certain operators can connect to them. To connect to a special share, follow these steps:

1. Open File Explorer, tap or click the leftmost option button in the address list, and then tap or click Computer.
2. Next, tap or click the Map Network Drive button on the Computer panel and then tap or click Map Network Drive. This displays the Map Network Drive dialog box, shown in Figure 12-11.



**FIGURE 12-11** Connect to special shares by mapping them with 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, type 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 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. Tap or 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 **Cpandl\Williams**. Before tapping or 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 2012 system. Whenever a user or computer connects to a shared resource, Windows Server 2012 lists a connection in the Sessions node.

To view connections to shared resources, type **net session** at a command prompt or 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 the user is actively working with. 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.

## 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, type **net session \\computername /delete** at a command prompt or 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 select Sessions.
3. Press and hold or right-click the user sessions you want to end, and then tap or click Close Session.
4. Tap or 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 press and hold or right-click Sessions.
3. Tap or click Disconnect All Sessions, and then tap or 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 once 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.

## 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.

## 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 you want to work with.
2. In the console tree, expand System Tools, expand Shared Folders, and then select Open Files.
3. Press and hold or right-click the open file you want to close, and then tap or click Close Open File.
4. Tap or click Yes to confirm the action.

## 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 press and hold or right-click Open Files.
3. Tap or click Disconnect All Open Files, and then tap or click Yes to confirm the action.

## 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. Press and hold or right-click the share you want to remove, tap or click Stop Sharing, and then tap or click Yes to confirm the action.

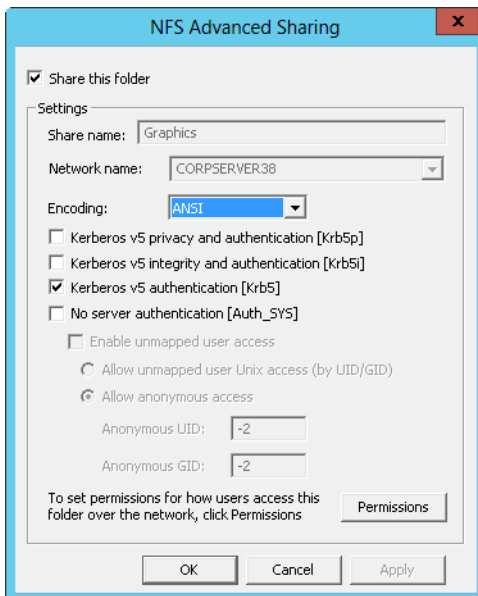
**CAUTION** You should never delete a folder containing shares without first stopping the shares. If you fail to stop the shares, Windows Server 2012 attempts to reestablish the shares the next time the computer is started, and the resulting error is logged in the system event log.

## Configuring NFS Sharing

As discussed in Chapter 10, “Managing File Systems and Drives,” 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 2012, OS X, and UNIX operating systems using the Network File System (NFS) protocol.

You can configure NFS sharing for local folders on NTFS volumes using File Explorer. You can also configure NFS sharing of local and remote folders on NTFS volumes by using Server Manager. In File Explorer, follow these steps to enable and configure NFS sharing:

1. Press and hold or right-click the share you want to manage, and then tap or click Properties. This displays a Properties dialog box for the share.
2. On the NFS Sharing tab, tap or click Manage NFS Sharing.
3. In the NFS Advanced Sharing dialog box, select the Share This Folder check box, as shown in Figure 12-12.



**FIGURE 12-12** You can use NFS sharing to share resources between Windows and UNIX computers.

4. In the Share Name text box, type 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 as well as 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, tap or click Permissions, set the permissions you want to use in the NFS Share Permissions dialog box, and then tap or click OK. You can configure no access, read-only access, or read/write access by client computer name and client computer groups.
9. Tap or 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. Press and hold or right-click the share you want to manage, and then tap or click Properties. This displays a Properties dialog box for the share.
2. On the NFS Sharing tab, tap or click Manage NFS Sharing.
3. In the NFS Advanced Sharing dialog box, clear the Share This Folder check box and then tap or 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, tap or click Tasks and then tap or click New Share. This starts the New Share Wizard. Choose NFS Share—Quick or NFS Share—Advanced as the share profile, and then tap or 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 tap or 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. Tap or click Close.

## Using Shadow Copies

---

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. These 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 2012 includes a feature enhancement that allows you to revert an entire (nonsystem) volume to a previous shadow copy state.

## Understanding 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 100 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.

You configure and view current Shadow Copy settings on the Shadow Copies tab of the disk's Properties dialog box. In File Explorer or Computer Management, press and hold or right-click the icon for the disk you want to work with, tap or click Properties, and then tap or 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.

## Creating Shadow Copies

To create a shadow copy on an NTFS volume with shared folders, follow these steps:

1. Open Computer Management. If necessary, connect to a remote computer.
2. In the console tree, expand Storage, and then select Disk Management. The volumes configured on the selected computer are displayed in the details pane.
3. Press and hold or right-click Disk Management, point to All Tasks, and then tap or click Configure Shadow Copies.
4. On the Shadow Copies tab, select the volume you want to work with in the Select A Volume list.
5. Tap or click Settings to configure the maximum size of all shadow copies for this volume and to change the default schedule. Tap or click OK.
6. After you configure the volume for shadow copying, tap or click Enable if necessary. When prompted to confirm this action, tap or click Yes. Enabling shadow copying creates the first shadow copy and sets the schedule for later shadow copies.

**NOTE** If you create a run schedule when configuring the shadow copy settings, shadow copying is enabled automatically for the volume when you tap or 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 tap or 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, press and hold or right-click the share for which you want to access previous file versions, tap or click Properties, and then tap or click the Previous Versions tab.
2. On the Previous Versions tab, select the folder version you want to work with. Each folder has a date and time stamp. Tap or click the button corresponding to the action you want to perform:
  - Tap or click Open to open the shadow copy in File Explorer.
  - Tap or click Copy to display the Copy Items dialog box, which lets you copy the snapshot image of the folder to the location you specify.



- Tap or click Restore to roll back the shared folder to its state at the time of the snapshot image you selected.

## Reverting an Entire Volume to a Previous Shadow Copy

Windows Server 2012 features a shadow copy enhancement that allows 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. Open Computer Management. If necessary, connect to a remote computer.
2. In the console tree, expand Storage. Press and hold or right-click Disk Management, point to All Tasks, and then tap or click Configure Shadow Copies.
3. On the Shadow Copies tab, select the volume you want to work with in the Select A Volume list.
4. 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 tap or click Revert.
5. To confirm this action, select the Check Here If You Want To Revert This Volume check box and then tap or click Revert Now. Tap or 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. This recovers the disk space used by the shadow copies.

To delete a shadow copy, follow these steps:

1. Open Computer Management. If necessary, connect to a remote computer.
2. In the console tree, expand Storage. Press and hold or right-click Disk Management, point to All Tasks, and then tap or click Configure Shadow Copies.
3. On the Shadow Copies tab, select the volume you want to work with in the Select A Volume list.
4. 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 tap or click Delete Now. Tap or 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. Open Computer Management. If necessary, connect to a remote computer.
2. In the console tree, expand Storage. Press and hold or right-click Disk Management, point to All Tasks, and then tap or click Configure Shadow Copies.
3. On the Shadow Copies tab, select the volume you want to work with in the Select A Volume list and then tap or click Disable.
4. When prompted, confirm the action by tapping or clicking Yes. Tap or click OK to close the Shadow Copies dialog box.

## 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 2012 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 2012, you connect to a network drive by mapping to it using NET USE and 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: \\ROME0\DOCS
```

or

```
net use * \\ROME0\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.

If the client computer is running Windows 8, you can map network drives by completing the following steps:

1. In File Explorer, tap or click the leftmost option button in the address list, and then tap or click Computer.
2. Next, tap or click the Map Network Drive button in the Computer panel and then tap or click Map Network Drive.

3. Use the Drive list to select a free drive letter to use, and then tap or 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 see a list of shared folders. Select the shared folder you want to work with, and then tap or 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. Tap or click Finish. If the currently logged-on user doesn't have appropriate access permissions for the share, select Connect Using Different Credentials and then tap or click Finish. After you tap or 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 **Cpandl\Williams**. Before tapping or clicking OK, select Remember My Credentials if you want the credentials to be saved. Otherwise, you need to provide credentials in the future.

## Disconnecting a Network Drive

To disconnect a network drive, follow these steps:

1. In File Explorer, tap or click the leftmost option button in the address list, and then tap or click Computer.
2. Under Network Location, press and hold or right-click the network drive icon, and then tap or click Disconnect.

## Object Management, Ownership, and Inheritance

---

Windows Server 2012 takes an object-based approach to describing resources and managing permissions. Objects that describe resources are defined on NTFS volumes and in Active Directory. 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 12-2.

**TABLE 12-2** Windows Server 2012 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 2012, 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. This ensures that authorized administrators 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, provided that 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, provided that 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. Press and hold or right-click the object you want to take ownership of, and then tap or click Properties. In the Properties dialog box, tap or click the Security tab.
3. On the Security tab, tap or click Advanced to display the Advanced Security Settings dialog box where the current owner is listed under the file or folder name.

4. Tap or click Change. Use the options in the Select User, Computer, Service Account, Or Group dialog box to select the new owner.
5. Tap or 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. Here, you would take ownership of all child objects.

## Object Inheritance

Objects are defined using a parent-child structure. A parent object is a top-level object. 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 2012 objects are created with inheritance enabled by default. This means that child objects automatically inherit the permissions of the parent. Because of this, the parent object permissions control access to the child object. If you want to change permissions on a child object, you must do the following:

1. Edit the permissions of the parent object.
2. Stop inheriting permissions from the parent object, and then assign permissions to the child object.
3. 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. Press and hold or right-click the object you want to work with, and then tap or click Properties. In the Properties dialog box, tap or click the Security tab.
3. Tap or click Advanced to display the Advanced Security Settings dialog box.
4. On the Permissions tab, tap or 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. Tap or 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. Press and hold or right-click the object you want to work with, and then tap or click Properties. In the Properties dialog box, tap or click the Security tab.
3. Tap or click Advanced to display the Advanced Security Settings dialog box.
4. On the Permissions tab, tap or click Enable Inheritance and then tap or 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. These permissions grant or deny access to the files and folders. Because Windows Server 2012 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, press and hold or right-click the file or folder you want to work with and then tap or click Properties. In the Properties dialog box, tap or 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.

As discussed earlier in the chapter, 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. Press and hold or right-click the folder you want to work with, and then tap or click Properties. This displays a Properties dialog box.
3. When you tap or click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
4. To get more information, tap or click Customize Permissions to open the Advanced Security Settings dialog box.

On file servers running Windows Server 2012, you also can 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 Table 12-3. 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.

**TABLE 12-3** File and Folder Permissions Used by Windows Server 2012

PERMISSION	MEANING FOR FOLDERS	MEANING FOR FILES
Read	Permits viewing and listing files and subfolders	Permits viewing or accessing a file's contents
Write	Permits adding files and subfolders	Permits writing to a file
Read & Execute	Permits viewing and listing files and subfolders as well as executing files; inherited by files and folders	Permits viewing and accessing a file's contents as well as executing a file
List Folder Contents	Permits viewing and listing files and subfolders as well as executing files; inherited by folders only	N/A
Modify	Permits reading and writing of files and subfolders; allows deletion of the folder	Permits reading and writing of a file; allows deletion of a file
Full Control	Permits reading, writing, changing, and deleting files and subfolders	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. A user can still delete the 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. Table 12-4 shows special permissions used to create the basic permissions for files. 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.
- Actions that users can perform are based on the sum of all the permissions assigned to the user and to all the groups the user is a member of. 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.

**TABLE 12-4** Special Permissions for Files

SPECIAL PERMISSIONS	BASIC PERMISSIONS				
	FULL CONTROL	MODIFY	READ & EXECUTE	READ	WRITE
Traverse Folder/ Execute File	Yes	Yes	Yes		
List Folder/Read Data	Yes	Yes	Yes	Yes	
Read Attributes	Yes	Yes	Yes	Yes	
Read Extended Attributes	Yes	Yes	Yes	Yes	
Create Files/Write Data	Yes	Yes			Yes
Create Folders/ Append Data	Yes	Yes			Yes
Write Attributes	Yes	Yes			Yes
Write Extended Attributes	Yes	Yes			Yes
Delete Subfolders And Files	Yes				
Delete	Yes	Yes			
Read Permissions	Yes	Yes	Yes	Yes	Yes
Change Permissions	Yes				
Take Ownership	Yes				

Table 12-5 shows special permissions used to create the basic permissions for folders. As you study the 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.



**TABLE 12-5** Special Permissions for Folders

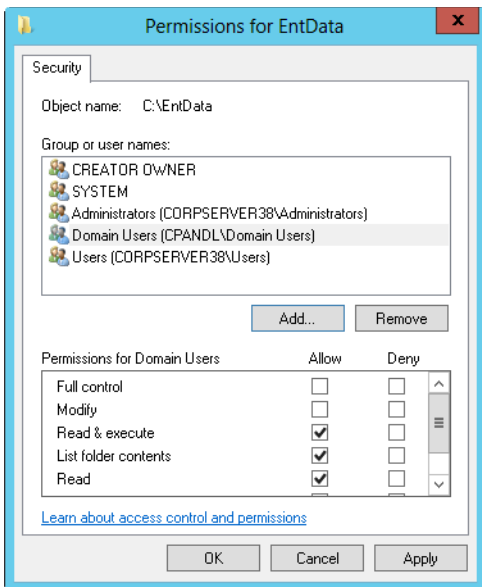
SPECIAL PERMISSIONS	BASIC PERMISSIONS					
	FULL CONTROL	MODIFY	READ & EXECUTE	LIST FOLDER CONTENTS	READ	WRITE
Traverse Folder/Execute File	Yes	Yes	Yes	Yes		
List Folder/Read Data	Yes	Yes	Yes	Yes	Yes	
Read Attributes	Yes	Yes	Yes	Yes	Yes	
Read Extended Attributes	Yes	Yes	Yes	Yes	Yes	
Create Files/Write Data	Yes	Yes				Yes
Create Folders/Append Data	Yes	Yes				Yes
Write Attributes	Yes	Yes				Yes
Write Extended Attributes	Yes	Yes				Yes
Delete Subfolders And Files	Yes					
Delete	Yes	Yes				
Read Permissions	Yes	Yes	Yes	Yes	Yes	Yes
Change Permissions	Yes					
Take Ownership	Yes					

## Setting Basic File and Folder Permissions

To set basic NTFS permissions for files and folders, follow these steps:

1. In File Explorer, press and hold or right-click the file or folder you want to work with and then tap or click Properties. In the Properties dialog box, tap or click the Security tab.
2. Tap or click Edit to display an editable version of the Security tab, as shown in Figure 12-13.
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 doing the following:
  - a. Select the user or group you want to change.
  - b. Grant or deny 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.



**FIGURE 12-13** Configure basic permissions for the file or folder on the Security tab.

4. To set access permissions for additional users, computers, or groups, tap or click Add. This displays the Select Users, Computers, Service Accounts, Or Groups dialog box.
5. Type the name of a user, computer, or group in the current domain, and then tap or 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 tap or click Locations to select a new location.
  - If multiple matches are found, select the name or names you want to use and then tap or click OK. To add more users, computers, or groups, type a semicolon (;) and then repeat this step.

**NOTE** The Locations button allows you to access account names in other domains. Tap or click Locations to see a list of the current domain, trusted domains, and other resources you can access. Because of the transitive trusts in Windows Server 2012, 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. Tap or click OK.

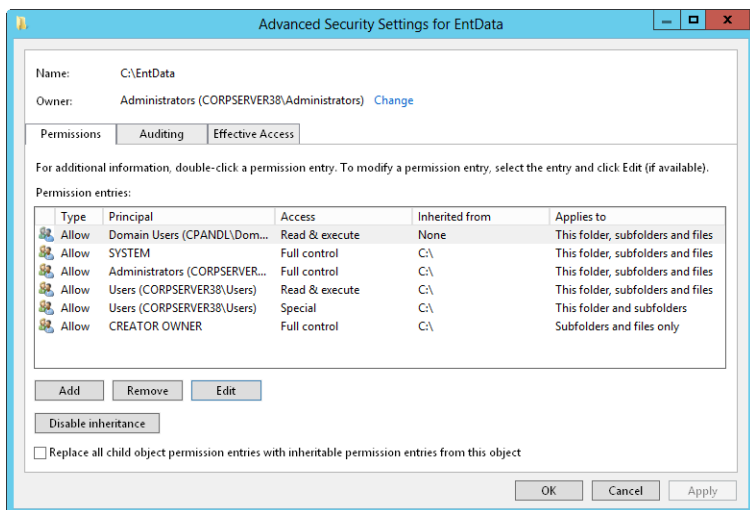
Because shared folders also have NTFS permissions, you might want to set basic NTFS permissions using Server Manager. To do this, follow these steps:

1. In Server Manager, press and hold or right-click the folder you want to work with and then tap or click Properties. This displays a Properties dialog box.
2. When you tap or click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
3. Tap or 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.

## Setting Special Permissions on Files and Folders

To set special NTFS permissions for files and folders, follow these steps:

1. In File Explorer, press and hold or right-click the file or folder you want to work with and then tap or click Properties.
2. In the Properties dialog box, select the Security tab and then tap or click Advanced to display the Advanced Security Settings dialog box. Before you can modify permissions, you must click Change Permissions. As shown in Figure 12-14, 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.

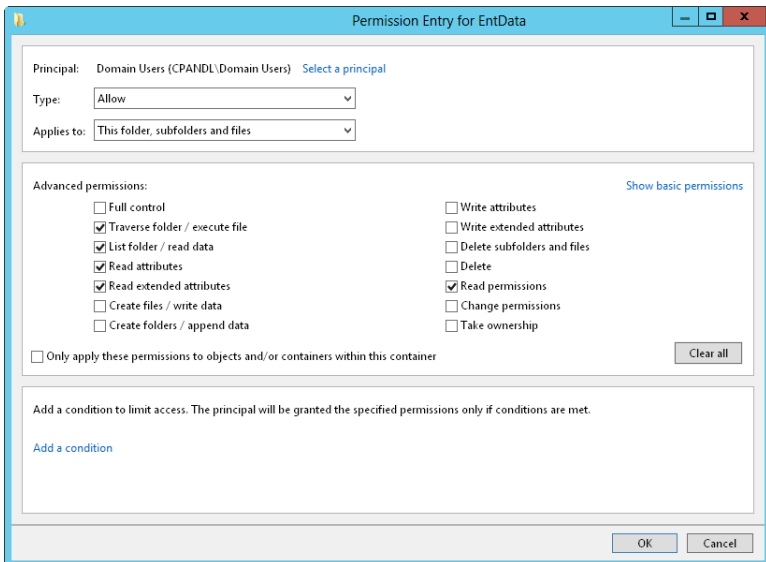


**FIGURE 12-14** Configure special permissions on files and folders.

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, tap or click Add to display the Permission Entry dialog box. Tap or click Select A Principal to display the Select User, Computer, Service Account, Or Group dialog box.
5. Type 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. Tap or 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 see 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 tap 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 tap or click OK.
7. Tap or 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. Tap or click Show Advanced Permissions to display the special permissions, as shown in Figure 12-15.



**FIGURE 12-15** Configure the special permissions that should be allowed or denied.

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, tap or click OK.

Because shared folders also have NTFS permissions, you might want to set special NTFS permissions using Server Manager. To do this, follow these steps:

1. In Server Manager, select File And Storage Services and then select Shares. Next, press and hold or right-click the folder you want to work with and then tap or click Properties. This displays a Properties dialog box.
2. When you tap or click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
3. Tap or 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 to 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 a user is a member of and the groups the user's computer is a member of 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. You do this 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. Using these claim types, you could then fine-tune access to ensure only users, devices, or both that belong to specific business categories and have certain country codes are granted access to a resource. You also could 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, as well as the servers where 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 non-domain-joined computers, 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 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 8 and

Windows Server 2012 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 controllers organizational unit (OU).

Once you've enabled and configured the claims-based policy, you can define claim conditions by completing these steps:

- 1.** In File Explorer, press and hold or right-click the file or folder you want to work with and then tap or click Properties. In the Properties dialog box, select the Security tab and then tap or 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, tap or click the user you want to work with, tap or click Edit, and then skip steps 3–6.
- 3.** Tap or click Add to display the Permission Entry dialog box. Tap or click Select A Principal to display the Select User, Computer, Service Account, Or Group dialog box.
- 4.** Type 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.** Tap or 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 see 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 tap 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 tap or click OK.
- 6.** Tap or click OK. The user and group are added as the Principal. Tap or 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, tap or click OK.

Because shared folders also have NTFS permissions, you might want to set claims-based permissions using Server Manager. To do this, follow these steps:

1. In Server Manager, press and hold or right-click the folder you want to work with and then tap or click Properties. This displays a Properties dialog box.
2. When you tap or click Permissions in the left pane, the current share permissions and NTFS permissions are shown in the main pane.
3. Tap or 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, 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 2012 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.

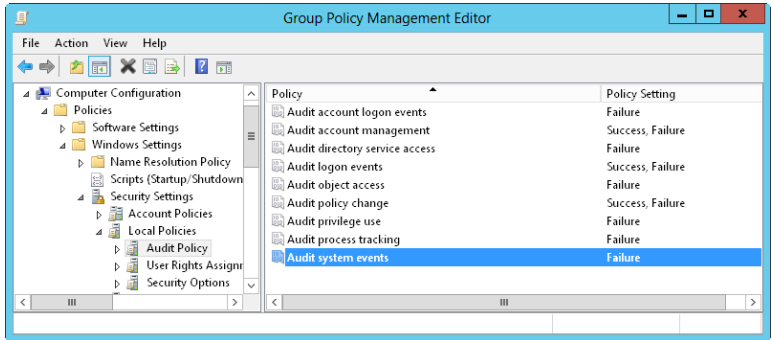
## Setting Auditing Policies

Auditing policies are essential to 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 you want to work with, you can set auditing policies by following these steps:

1. In the Group Policy Management Editor, shown in Figure 12-16, 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.



**FIGURE 12-16** Set auditing policies using the Audit Policy node in Group 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, as well as actions that affect system security or the security log.

3. To configure an auditing policy, double-tap or double-click its entry, or press and hold or right-click the entry, and then tap or 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. Tap or click OK.

When auditing is enabled, the security event log will reflect the following:

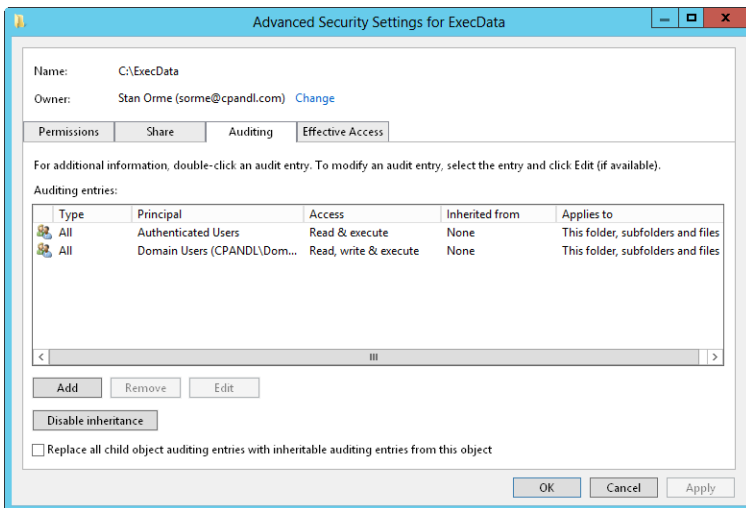
- Event IDs of 560 and 562 detailing user audits
- Event IDs of 592 and 593 detailing process audits

## 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 allows 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, press and hold or right-click the file or folder to be audited and then tap or click Properties.
2. Tap or click the Security tab, and then tap or click Advanced. This displays the Advanced Security Settings dialog box.
3. On the Auditing tab, tap or click Continue. You can now view and manage auditing settings by using the options shown in Figure 12-17.



**FIGURE 12-17** After you audit object access, you can set auditing policies on individual files and folders on the Auditing tab.

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 tap or click Remove.
5. To configure auditing for additional users, computers, or groups, tap or click Add. This displays the Select Users, Computers, Service Accounts, Or Groups dialog box.
6. Type the name of a user, computer, or group in the current domain, and then tap or click Check Names. If a single match is found, the dialog box is automatically updated and the entry is underlined. Otherwise, you'll see 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 tap 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 tap or click OK.
7. Tap or 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, tap or 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 listed in Tables 12-4 and 12-5, 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, tap or 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 allows 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, type **regedit**, or type **regedit** in the Apps 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, tap or click Advanced. In the Advanced Security Settings dialog box, tap or click the Auditing tab.
4. Tap or click Add to display the Auditing Entry dialog box. Tap or 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, type **Everyone**, tap or click Check Names, and then tap or click OK.
6. In the Auditing Entry dialog box, only basic permissions are listed by default. Tap or 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. Tap or 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. This allows you to 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-tap or double-click the object to be audited. This opens the related Properties dialog box.
3. Tap or click the Security tab, and then tap or click Advanced.
4. In the Advanced Settings dialog box, tap or 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 tap or click Remove.
5. To add specific accounts, tap or click Add to display the Auditing Entry dialog box. Tap or click Select A Principal to display the Select User, Computer, Service Account, Or Group dialog box.
6. Type the name of a user, computer, or group in the current domain, and then tap or click Check Names. If a single match is found, the dialog box is automatically updated and the entry is underlined. Otherwise, you'll see 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 tap 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 tap or click OK.
7. Tap or 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. Tap or click OK. Repeat this process to audit other users, groups, or computers.

## Using, Configuring, and Managing NTFS Disk Quotas

---

Windows Server 2012 supports two mutually exclusive types of disk quotas:

- **NTFS disk quotas** NTFS disk quotas are supported on all editions of Windows Server 2012 and allow you to manage disk space usage by users. You configure quotas on a per-volume basis. Although users who exceed limits see 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 2012 and allow you to manage disk space usage by folder 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.

The sections that follow discuss NTFS disk quotas.

**NOTE** 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.

**REAL WORLD** 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 be unable to apply changes to Group Policy because the Group Policy client runs within a Local-System 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. You would be unable to disable or modify the quota settings through Group Policy, for example.

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 allows 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.

## Understanding NTFS Disk Quotas and How NTFS Quotas Are Used

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 do this. Sometimes you want to track disk space usage on a per-user 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. NTFS disk quotas don't apply to administrators. Administrators can't be denied disk space even if they exceed enforced disk quota limits.

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 to 20 to 100 GB. For a user data share, you might want to set the level much lower, such as 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 to 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. 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. You can do this 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 2012 manages disk quotas according to security identifiers (SIDs). Because SIDs track 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. That's because Windows Server 2012 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 2012 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.

## Setting NTFS Disk Quota Policies

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 Computer Configuration\Administrative Templates\System\Disk Quotas. Table 12-6 summarizes the available policies.

**TABLE 12-6** Policies for Setting NTFS Disk Quotas

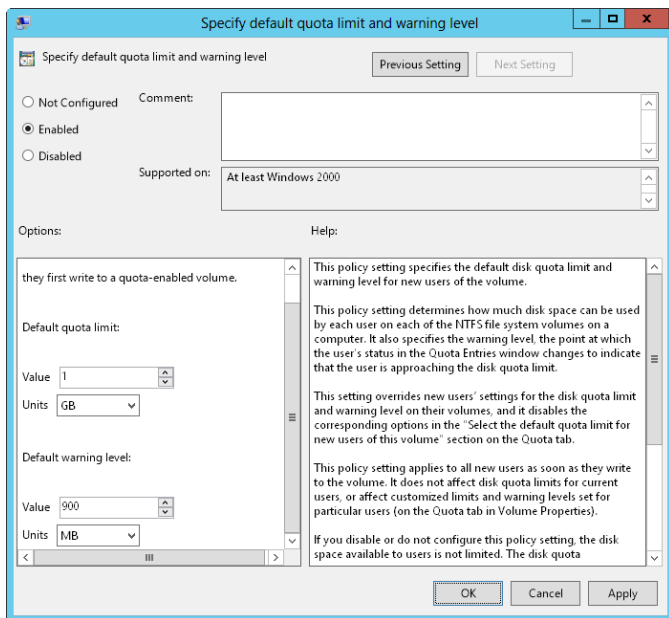
POLICY NAME	DESCRIPTION
Apply Policy To Removable Media	Determines whether quota policies apply to NTFS volumes on removable media. If you don't enable this policy, quota limits apply only to fixed media drives.
Enable Disk Quotas	Turns disk quotas on or off for all NTFS volumes of the computer, and 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.



POLICY NAME	DESCRIPTION
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. Typically, you 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. Access Group Policy for the system (for example, a file server) that you want to work with. Access the Disk Quotas node by expanding Computer Configuration, Administrative Templates, System and then selecting Disk Quotas.
2. Double-tap or double-click Enable Disk Quotas. Select Enabled, and then tap or click OK.
3. Double-tap or double-click Enforce Disk Quota Limit. If you want to enforce disk quotas on all NTFS volumes residing on this computer, tap or click Enabled. Otherwise, tap or click Disabled, and then set specific limits on a per-volume basis. Tap or click OK.
4. Double-tap or double-click Specify Default Quota Limit And Warning Level. In the dialog box shown in Figure 12-18, select Enabled.



**FIGURE 12-18** Enforce disk quotas in the Specify Default Quota Limit And Warning Level dialog box.

5. 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.
6. 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. Tap or click OK.
7. Double-tap or double-click Log Event When Quota Limit Exceeded. Select Enabled so that limit events are recorded in the application log, and then tap or click OK.
8. Double-tap or double-click Log Event When Quota Warning Level Exceeded. Select Enabled so that warning events are recorded in the application log, and then tap or click OK.
9. Double-tap or 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 tap or click OK.

**TIP** To ensure that the policies are enforced immediately, access the Computer Configuration\Administrative Templates\System\Group Policy node, and then double-tap or 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. Tap or click OK.

## Enabling NTFS Disk Quotas on NTFS Volumes

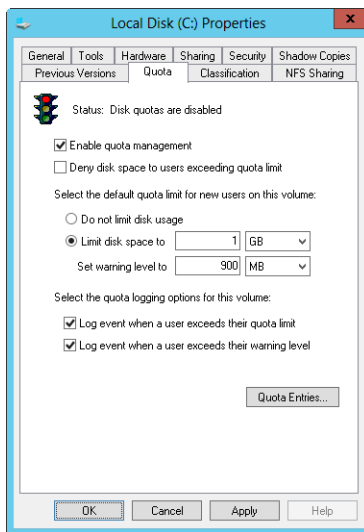
You can set NTFS disk quotas on a per-volume basis. Only NTFS volumes can have disk quotas. After you configure the appropriate group policies, you can use Computer Management to set disk quotas for local and remote volumes.

**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.

To enable NTFS disk quotas on an NTFS volume, follow these steps:

1. Open Computer Management. If necessary, connect to a remote computer.
2. In the console tree, expand Storage, and then select Disk Management. The volumes configured on the selected computer are displayed in the details pane.
3. Using Volume List view or Graphical View, press and hold or right-click the volume you want to work with, and then tap or click Properties.
4. On the Quota tab, select the Enable Quota Management check box, shown in Figure 12-19. 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.



**FIGURE 12-19** After you enable quota management, you can configure a quota limit and quota warning for all users.

**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.

5. 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.

6. 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).

7. To configure logging when users exceed a warning limit or the quota limit, select the Log Event check boxes. Tap or click OK to save your changes.
8. If the quota system isn't currently enabled, you'll see a prompt asking you to enable the quota system. Tap or click OK to allow Windows Server 2012 to 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. Open Computer Management. If necessary, connect to a remote computer.
2. In the console tree, expand Storage and then select Disk Management. The volumes configured on the selected computer are displayed in the details pane.
3. Using Volume List view or Graphical View, press and hold or right-click the volume you want to work with and then tap or click Properties.
4. On the Quota tab, tap or click Quota Entries. This displays the Quota Entries dialog box. Each quota entry is listed according to a status. The status 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.

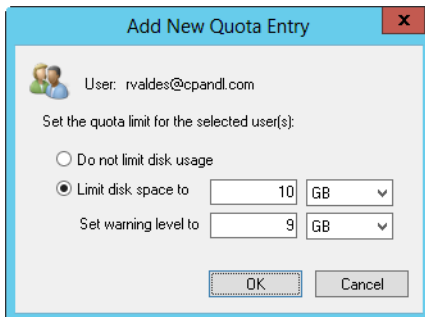
## Creating Disk Quota Entries

You can create disk quota entries for users who haven't yet saved data on a volume. This allows 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. Open the Quota Entries dialog box as discussed in "Viewing Disk Quota Entries" earlier in this chapter. Current quota entries for all users are listed. To refresh the listing, press F5 or choose Refresh from the View menu.
2. 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.
3. In the Select Users dialog box, type the name of a user you want to use in the Enter The Object Names To Select text box and then tap or click Check Names. If a match is found, select the account you want to use and then tap or click OK. If no matches are found, update the name you entered and try searching again. Repeat this step as necessary, and then tap or click OK.
4. After you select a user, the Add New Quota Entry dialog box is displayed, as shown in Figure 12-20. 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. Tap or click OK.



**FIGURE 12-20** In the Add New Quota Entry dialog box, you can customize the user's quota limit and warning level or remove quota restrictions altogether.

## 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. Open the Quota Entries dialog box as discussed in "Viewing Disk Quota Entries" earlier in this chapter. Current quota entries for all users are listed. To refresh the listing, press F5 or choose Refresh from the View menu.
2. 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 using the Shift and Ctrl keys.
3. When prompted to confirm the action, tap or click Yes. This displays the Disk Quota dialog box with a list of current files owned by the selected user or users.
4. 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. 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, tap or click Yes.
  - **Take Ownership Of Files** Select the files you want to take ownership of, and then tap or 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, tap or click Browse to display the Browse For Folder dialog box. Once you find the folder, tap or click Move.
5. Tap or 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 using NTFS. To export and then import disk quota entries, follow these steps:

1. Open the Quota Entries dialog box as discussed in "Viewing Disk Quota Entries" earlier in this chapter. Current quota entries for all users are listed. To refresh the listing, press F5 or choose Refresh from the View menu.
2. Select Export from the Quota menu. This displays the Export Quota Settings dialog box. Choose the save location for the file containing the quota

settings, and then type a name for the file in the File Name text box. Tap or 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.

3. On the Quota menu, tap or click Close to exit the Quota Entries dialog box.
4. Press and hold or right-click Computer Management in the console tree, and then tap or 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.
5. As explained previously, open the Properties dialog box for the target volume. Then tap or click Quota Entries on the Quota tab. This displays the Quota Entries dialog box for the target volume.
6. Tap or click Import on the Quota menu. In the Import Quota Settings dialog box, select the quota settings file you saved previously. Tap or click Open.
7. 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, tap or click Yes to replace an existing entry or tap 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 tapping or clicking Yes or No.

## Disabling NTFS Disk Quotas

You can disable quotas for individual users or all users on a volume. 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 volume, follow these steps:

1. Open Computer Management. If necessary, connect to a remote computer.
2. Open the Properties dialog box for the volume on which you want to disable NTFS quotas.
3. On the Quota tab, clear the Enable Quota Management check box. Tap or click OK. When prompted to confirm, tap or click OK.

## Using, Configuring, and Managing Resource Manager Disk Quotas

---

Windows Server 2012 supports an enhanced quota management system called *Resource Manager disk quotas*. 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 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

When you're working with Windows Server 2012, 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 basis and on a 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 2012 includes the quota templates listed in Table 12-7. 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.



**TABLE 12-7** Disk Quota Templates

QUOTA TEMPLATE	LIMIT	QUOTA TYPE	DESCRIPTION
100 MB Limit	100 MB	Hard	Sends warnings to users as the limit is approached and exceeded
200 MB Limit Reports To User	200 MB	Hard	Sends storage reports to the users who exceed the threshold
200 MB Limit With 50 MB Extension	200 MB	Hard	Uses the DIRQUOTA command to grant an automatic, one-time, 50-MB extension to users who exceed the quota limit
250 MB Extended Limit	250 MB	Hard	Meant to be used by those whose limit has been extended from 200 MB to 250 MB
Monitor 200 GB Volume Usage	200 GB	Soft	Monitors volume usage, and warns when the limit is approached and exceeded
Monitor 500 MB Share	500 MB	Soft	Monitors share usage, and warns when the limit is approached and exceeded

## 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. Table 12-7, shown earlier, provides a summary of the default disk quota templates. Table 12-8, which follows, shows variables that can be used for automatically generated messages and events.

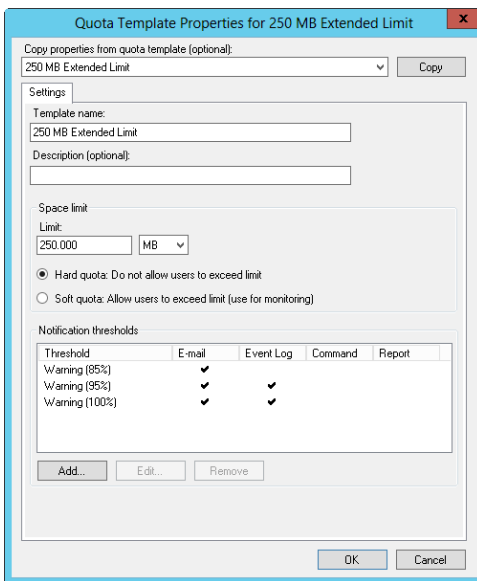
**TABLE 12-8** Key Variables Available for Disk Quota Messages and Event Logging

VARIABLE NAME	DESCRIPTION
[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

VARIABLE NAME	DESCRIPTION
[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-tap or double-click the disk quota template name. This displays a related Properties dialog box, as shown in Figure 12-21.



**FIGURE 12-21** Use disk quota properties to configure the limit, quota type, and notification thresholds.

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 tap or click Edit. To define a new threshold, tap or click Add.

4. When you have finished modifying the quota template, tap or 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, tap or click Create Quota Template. This displays 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, type the limit value and specify whether you are setting the limit in kilobytes, megabytes, gigabytes, or terabytes.
4. Tap or 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 type 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. Table 12-8 lists available variables and their meanings.
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. Table 12-8 lists available variables and their meanings.
7. 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.
8. Repeat steps 5–7 to define additional notification thresholds.
9. Tap or 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. Tap or 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 tapping or clicking Browse and then using the Browse For Folder dialog box to select the path, such as C:\Data. Tap or 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. Tap or click Create.

# Data Backup and Recovery

- Creating a Backup and Recovery Plan 519
- Backing Up Your Data: The Essentials 526
- Performing Server Backups 534
- Managing Encryption Recovery Policy 551
- Backing Up and Restoring Encrypted Data and Certificates 554

Because data is the heart of the enterprise, protecting it is crucial. And to protect your organization's data, you need to implement a data backup and recovery plan. Backing up files can protect against accidental loss of user data, database corruption, hardware failures, and even natural disasters. Your job as an administrator is to make sure that backups are performed and that backups are stored in a secure location.

## Creating a Backup and Recovery Plan

---

Data backup is an insurance plan. Important files are accidentally deleted all the time. Mission-critical data can become corrupt. Natural disasters can leave your office in ruin. With a solid backup and recovery plan, you can recover from any of these events. Without one, you're left with nothing to fall back on.

## Figuring Out a Backup Plan

It takes time to create and implement a backup and recovery plan. You need to figure out what data needs to be backed up, how often the data should be backed up, and more. To help you create a plan, consider the following questions:

- **How important or sensitive is the data on your systems?** Knowing the importance of data can go a long way toward helping you determine whether you need to back it up, as well as when and how it should be backed up. For critical data, such as a database, you should have redundant backup sets that cover several backup periods. For sensitive data, you should be sure that backup data is physically secure or encrypted. For less important data, such as daily user files, you won't need such an elaborate backup plan, but you need to back up the data regularly and ensure that the data can be recovered easily.

- **What type of information does the data contain?** Data that doesn't seem important to you might be very important to someone else. The type of information the data contains can help you determine whether you need to back up the data, as well as when and how the data should be backed up.
- **How often does the data change?** The frequency of change can affect your decision on how often certain data should be backed up. For example, data that changes daily should be backed up daily.
- **Can you supplement backups with shadow copies?** *Shadow copies* are point-in-time copies of documents in shared folders. These point-in-time copies make recovering documents easy because you can quickly go back to an older version in case a document is deleted or overwritten accidentally. You should use shadow copies in addition to standard backups, not to replace backup procedures.
- **How quickly do you need to recover the data?** Recovery time is an important factor in a backup plan. For critical systems, you might need to get back online swiftly. To do this, you might need to alter your backup plan.
- **Do you have the equipment to perform backups?** You must have backup hardware to perform backups. To perform timely backups, you might need several backup devices and several sets of backup media. Backup hardware includes hard disk drives, tape drives, optical drives, and removable disk drives. In most environments, hard disk drives have become the preferred back up media.
- **Who will be responsible for the backup and recovery plan?** Ideally, someone should be a primary contact for the organization's backup and recovery plan. This person might also be responsible for performing the actual backup and recovery of data.
- **What's the best time to schedule backups?** Scheduling backups when system use is as low as possible will speed up the backup process. However, you can't always schedule backups for off-peak hours, so you need to carefully plan when key system data is backed up.
- **Do you need to store backups off-site?** Storing copies of backups off-site is essential to recovering your systems in the event of a natural disaster. In your off-site storage location, you should also include copies of the software you might need to install to reestablish operational systems.

**REAL WORLD** Recovery time objective (RTO) and recovery point objective (RPO) are important factors to consider. RTO represents the time to recover, which might be two hours for one server and four hours for another server. RPO represents your potential data loss, which might be one business day of data with one server or two business days with another server. A high RTO environment is an environment in which you can recover server functionality quickly after an outage. A high RPO environment is an environment in which the data recovered is as up to date as possible.

The frequency of your full server backups will vary according to the speed of your backup system and the amount of data you need to back up. The frequency at which you can create backups controls both the RPO and the RTO available to you. For example, with nightly backups, your RPO will be one business day, meaning that any

server outage will likely result in the loss of an entire business day of data. Meanwhile, your RTO, indicating how long it actually takes to recover, will vary according to the amount of data you have to restore.

## The Basic Types of Backup

There are many techniques for backing up files. The techniques you use depend on the type of data you're backing up, how convenient you want the recovery process to be, and more.

If you view the properties of a file or directory in File Explorer, you'll see an attribute called *archive*. You often use this attribute to determine whether a file or directory should be backed up. If the attribute is on, the file or directory might need to be backed up. You can perform the following basic types of backups:

- **Normal/full backups** All files that have been selected are backed up, regardless of the *archive* attribute's setting. When a file is backed up, the *archive* attribute is cleared. If the file is later modified, this attribute is set, indicating that the file needs to be backed up.
- **Copy backups** All files that have been selected are backed up, regardless of the *archive* attribute's setting. Unlike in a normal backup, the *archive* attribute on files isn't modified. This allows you to perform other types of backups on the files at a later date.
- **Differential backups** Designed to create backup copies of files that have changed since the last normal backup. The presence of the *archive* attribute indicates that the file has been modified, and only files with this attribute set are backed up. However, the *archive* attribute on files isn't modified. This allows you to perform other types of backups on the files at a later date.
- **Incremental backups** Designed to create backups of files that have changed since the most recent normal or incremental backup. The presence of the *archive* attribute indicates that the file has been modified, and only files with this attribute set are backed up. When a file is backed up, the *archive* attribute is cleared. If the file is later modified, this attribute is set, indicating that the file needs to be backed up.
- **Daily backups** Designed to back up files using the modification date on the file itself. If a file has been modified on the same day as the backup, the file will be backed up. This technique doesn't change the *archive* attribute of files.

As part of your backup operations, you'll probably want to perform full backups on a weekly basis and supplement this with daily, differential, or incremental backups. You might also want to create an extended backup set for monthly and quarterly backups that includes additional files that aren't being backed up regularly.

**TIP** You'll often find that weeks or months go by before anyone notices that a file or data source is missing. This doesn't mean the file isn't important. Although some types of data aren't used often, they're still needed. So don't forget that you might also want to create extra sets of backups for monthly or quarterly periods, or for both periods, to ensure that you can recover historical data.

## Differential and Incremental Backups

The difference between differential and incremental backups is extremely important. To understand the distinction, examine Table 13-1. As you can see, with differential backups you back up all the files that have changed since the last full backup (which means that the size of the differential backup grows over time). With incremental backups, you back up only files that have changed since the most recent full or incremental backup (which means the size of the incremental backup is usually much smaller than a full backup).

**TABLE 13-1** Incremental and Differential Backup Techniques

DAY OF WEEK	WEEKLY FULL BACKUP WITH DAILY DIFFERENTIAL BACKUP	WEEKLY FULL BACKUP WITH DAILY INCREMENTAL BACKUP
Sunday	A full backup is performed.	A full backup is performed.
Monday	A differential backup contains all changes since Sunday.	An incremental backup contains changes since Sunday.
Tuesday	A differential backup contains all changes since Sunday.	An incremental backup contains changes since Monday.
Wednesday	A differential backup contains all changes since Sunday.	An incremental backup contains changes since Tuesday.
Thursday	A differential backup contains all changes since Sunday.	An incremental backup contains changes since Wednesday.
Friday	A differential backup contains all changes since Sunday.	An incremental backup contains changes since Thursday.
Saturday	A differential backup contains all changes since Sunday.	An incremental backup contains changes since Friday.

After you determine what data you're going to back up and how often, you can select backup devices and media that support these choices. These are covered in the next section.

## Selecting Backup Devices and Media

Many tools are available for backing up data. Some are fast and expensive. Others are slow but very reliable. The backup solution that's right for your organization depends on many factors, including the following:

- **Capacity** The amount of data you need to back up on a routine basis. Can the backup hardware support the required load given your time and resource constraints?
- **Reliability** The reliability of the backup hardware and media. Can you afford to sacrifice reliability to meet budget or time needs?
- **Extensibility** The extensibility of the backup solution. Will this solution meet your needs as the organization grows?



- **Speed** The speed with which data can be backed up and recovered. Can you afford to sacrifice speed with server or service downtime to reduce costs?
- **Cost** The cost of the backup solution. Does it fit your budget?

## Common Backup Solutions

Capacity, reliability, extensibility, speed, and cost are the issues driving your backup plan. If you understand how these issues affect your organization, you'll be on track to select an appropriate backup solution. Some of the most commonly used backup solutions include the following:

- **Tape drives** Tape drives are the most common backup devices. Tape drives use magnetic tape cartridges to store data. Magnetic tapes are relatively inexpensive but aren't highly reliable. Tapes can break or stretch. They can also lose information over time. The average capacity of tape cartridges ranges from 24 gigabytes (GB) to 160 GB. Compared with other backup solutions, tape drives are slow. Still, the selling point is the low cost.
- **Digital audio tape (DAT) drives** DAT drives are quickly replacing standard tape drives as the preferred backup devices. Many DAT formats are available. The most commonly used formats are Digital Linear Tape (DLT) and Super DLT (SDLT). With SDLT 320 and 600, tapes have a capacity of either 160 GB or 300 GB uncompressed (320 GB or 600 GB compressed). Large organizations might want to look at Linear Tape Open (LTO) tape technologies. LTO-3, LTO-4, and LTO-5 tapes have uncompressed capacity of 400 GB, 800 GB, and 1500 GB, respectively (and compressed capacity of twice that).
- **Autoloader tape systems** Autoloader tape systems use a magazine of tapes to create extended backup volumes capable of meeting an enterprise's high-capacity needs. With an autoloader system, tapes within the magazine are automatically changed as necessary during the backup or recovery process. Most autoloader tape systems use DAT tapes formatted for DLT, SDLT, or LTO. Typical DLT drives can record up to 45 GB per hour, and you can improve that speed by purchasing a tape library system with multiple drives. In this way, you can record on multiple tapes simultaneously. In contrast, most SDLT and LTO drives record over 100 GB per hour, and by using multiple drives in a system you can record hundreds of GB per hour. An example enterprise solution uses 16 LTO drives to achieve data-transfer rates of more than 13.8 terabytes (TB) per hour and can store up to 500 tapes, for a total storage capacity of more than 800 TB.
- **Disk drives** Disk drives provide one of the fastest ways to back up and restore files. With disk drives, you can often accomplish in minutes what takes a tape drive hours. When business needs mandate a speedy recovery, nothing beats a disk drive. The drawback to disk drives is a relatively high cost compared to tape library systems.
- **Disk-based backup systems** Disk-based backup systems provide complete backup and restore solutions using large arrays of disks to achieve high performance. High reliability can be achieved when you use redundant

array of independent disks (RAID) to build in redundancy and fault tolerance. Typical disk-based backup systems use virtual library technology so that Microsoft Windows sees them as autoloader tape library systems. This makes them easier to work with. An example enterprise solution has 128 virtual drives and 16 virtual libraries per node for total storage of up to 7.5 TB per node. When fully scaled, this enterprise solution can store up to 640 TB and transfer up to 17.2 TB per hour.

**NOTE** Disks and disk-based backup systems can be used between the servers you are backing up and an enterprise autoloader. Servers are backed up to disk first (because this is very fast compared to tape) and later backed up to an enterprise autoloader. Having data on tapes also makes it easier to rotate backup sets to off-site storage. That said, tape backups are increasingly being replaced with disk backups. If you back up to disk arrays, you can move data off site by replicating the data to a secondary array at an alternative data center.

Before you can use a backup device, you must install it. When you install backup devices other than standard tape and DAT drives, you need to tell the operating system about the controller card and drivers that the backup device uses.

## Buying and Using Backup Media

Selecting a backup device is an important step toward implementing a backup and recovery plan. But you also need to purchase the tapes, disks, or both that allow you to implement your plan. The number of tapes or disks you need depends on how much data you have to back up, how often you need to back up the data, and how long you need to keep additional data sets.

The typical way to use backup tapes is to set up a rotation schedule whereby you rotate through two or more sets of tapes. The idea is that you can increase tape longevity by reducing tape usage and, at the same time, reduce the number of tapes you need to ensure that you have historic data on hand when necessary.

One of the most common tape-rotation schedules is the 10-tape rotation. With this rotation schedule, you use 10 tapes divided into two sets of 5 (one for each weekday). The first set of tapes is used one week, and the second set of tapes is used the next week. On Fridays, full backups are scheduled. On Mondays through Thursdays, incremental backups are scheduled. If you add a third set of tapes, you can rotate one of the tape sets to an off-site storage location on a weekly basis.

The 10-tape rotation schedule is designed for the 9-to-5 workers of the world. If you're in a 24/7 environment, you'll definitely want extra tapes for Saturday and Sunday. In this case, use a 14-tape rotation with two sets of 7 tapes. On Sundays, schedule full backups. On Mondays through Saturdays, schedule incremental backups.

As disk drives have become more affordable, many organizations have been using disk backup instead of tape backup. With disks, you can use a rotation schedule similar to the one you use with tapes. You will, however, need to modify the way you rotate disks to accommodate the amount of data you are backing up. The key thing to remember is to periodically rotate disks to off-site storage.

## Selecting a Backup Utility

Many backup and recovery solutions are available for use with Windows Server 2012. When selecting a backup utility, you need to keep in mind the types of backups you want to perform and the types of data you are backing up. Windows Server 2012 includes four installable backup and recovery features:

- **Windows Server Backup** A basic and easy-to-use backup and recovery utility. When this feature is installed on a server, you can open the tool using the Tools menu in Server Manager.
- **Backup Command-Line Tools** A set of backup and recovery commands accessible through the Wbadmin command-line tool. You run and use Wbadmin from an elevated, administrator command prompt. Type **wbadmin /?** for a full list of supported commands. Windows PowerShell cmdlets for managing backups are also available.
- **Microsoft Online Backup Service** This service is an add-on that can be downloaded and installed from within Windows Server Backup to schedule backups from a server to Microsoft's Internet cloud-based service. Online backups are possible only for fixed NTFS volumes that don't use BitLocker Drive Encryption. Volumes cannot be shares and must also be configured for read/write access.
- **Repair Your Computer** You can restore a server using repair options if you cannot access recovery options provided by the server manufacturer.

**NOTE** Windows Server Backup and the backup command-line tools are available only for management of backups when you add the Windows Server Backup feature to a server. If you add server administration tools to a server, you might be able to open Windows Server Backup. However, you won't be able to use Windows Server Backup to configure and manage backups.

Windows Server Backup is the feature you'll use the most. You can use Windows Server Backup to perform full or copy backups. You cannot use Windows Server Backup to perform differential backups. Windows Server Backup uses the Volume Shadow Copy Service (VSS) to create fast, block-level backups of the operating system, files and folders, and disk volumes. After you create the first full backup, you can configure Windows Server Backup to automatically run full or incremental backups on a recurring basis.

When you use Windows Server Backup, you need separate, dedicated media for storing archives of scheduled backups. You can back up to external and internal disks, DVDs, and shared folders. Although you can recover full volumes from DVD backups, you cannot recover individual files, folders, or application data from DVD backups.

**NOTE** You cannot back up to tape using Windows Server Backup. If you want to back up to tape, you need a third-party backup utility.

You can use Windows Server Backup to easily recover individual folders and files. Rather than manually restoring files from multiple backups if the files are stored

in incremental backups, you can recover folders and files by choosing the date on which you backed up the version of the item or items you want to restore. Windows Server Backup also works with the Windows Recovery tools, making it easier for you to recover the operating system. You can recover to the same server or to a new server that has no operating system. Because Windows Server Backup uses VSS, you can easily back up data from compliant applications, such as Microsoft SQL Server and Windows SharePoint Services.

Windows Server Backup also includes automatic disk management. You can run backups to multiple disks in rotation simply by adding each disk as a scheduled backup location. Once you configure a disk as a scheduled backup location, Windows Server Backup automatically manages the disk storage, ensuring that you no longer need to worry about a disk running out of space. Windows Server Backup reuses the space of older backups when creating newer backups. To help ensure that you can plan for additional storage needs, Windows Server Backup displays the backups that are available and the current disk usage information.

## Backing Up Your Data: The Essentials

---

Windows Server 2012 provides Windows Server Backup for creating backups. You use Windows Server Backup to archive files and folders, restore archived files and folders, create snapshots of the system state for backup and restore, and schedule automated backups.

### Installing the Windows Backup and Recovery Utilities

The Windows Server backup and recovery tools are available in all editions of Windows Server 2012. However, you cannot install the graphical components of these utilities on core installations of Windows Server 2012. On servers running with a core installation, you need to use the command line or manage backups via a remote session from another computer.

You can install the Windows backup and recovery tools by following these steps:

1. In Server Manager, select 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 tap or click Next.
2. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Tap or 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, tap or click Browse and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, tap or click Next twice.
4. On the Select Features page, select Windows Server Backup. Tap or click Next.

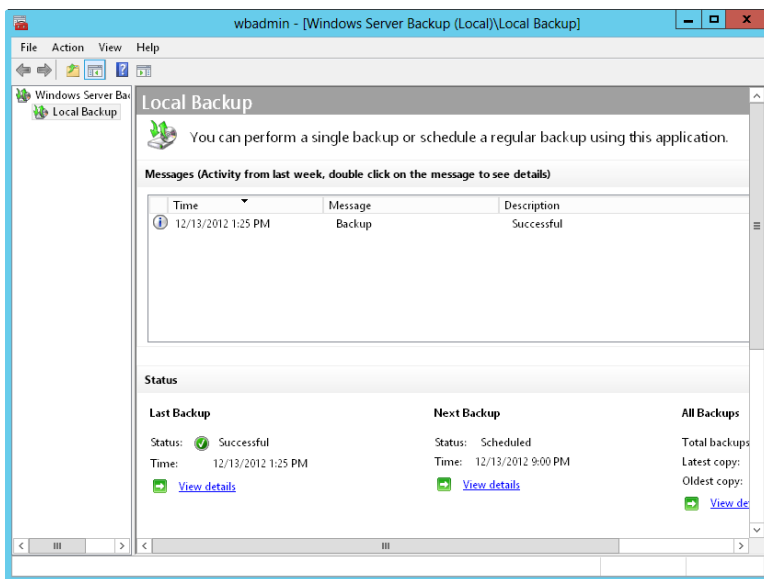
5. Tap or click Install. When the wizard finishes installing the selected features, tap or click Close. From now on, Windows Server Backup and the related command-line tools are available for managing backups.

**REAL WORLD** When you use Windows Server Backup with Microsoft Exchange Server 2010, you can use only full (normal) backups. Using the Windows Server Backup command-line tools with Exchange Server 2010 also is not supported. For more information on backing up Exchange Server 2010, see Chapter 13, “Backing Up and Restoring Exchange Server 2010,” in *Microsoft Exchange Server 2010 Administrator’s Pocket Consultant* (Microsoft Press, 2009).

## Getting Started with Windows Server Backup

You can start Windows Server Backup by selecting the related option on the Tools menu in Server Manager. When you start Windows Server Backup, you’ll see a message about online backup. If you want to use online backups, you need to sign up for the service, register your server, and download the Microsoft Online Backup Service agent. With the Windows Server Backup node selected, you can start this process by clicking the Continue button.

In Windows Server Backup, shown in Figure 13-1, select the Local Backup node to work with backups. The first time you use Windows Server Backup, you’ll see a warning that no backup has been configured for the computer. You clear this warning by creating a backup using the Backup Once feature, located on the Action menu, or by scheduling backups to run automatically using the Backup Schedule feature.



**FIGURE 13-1** Windows Server Backup provides a user-friendly interface for backup and restore.

To perform backup and recovery operations, you must have certain permissions and user rights. Members of the Administrators and Backup Operators groups have full authority to back up and restore any type of file, regardless of who owns the file and the permissions set on it. File owners and those who have been given control over files can also back up files, but they can do so only for files they own or for which they have Read, Read & Execute, Modify, or Full Control permissions.

**NOTE** Keep in mind that although local accounts can work only with local systems, domain accounts have domainwide privileges. Therefore, a member of the local Administrators group can work only with files on the local system, but a member of the Domain Admins group can work with files throughout the domain.

Windows Server Backup provides extensions for working with the following special types of data:

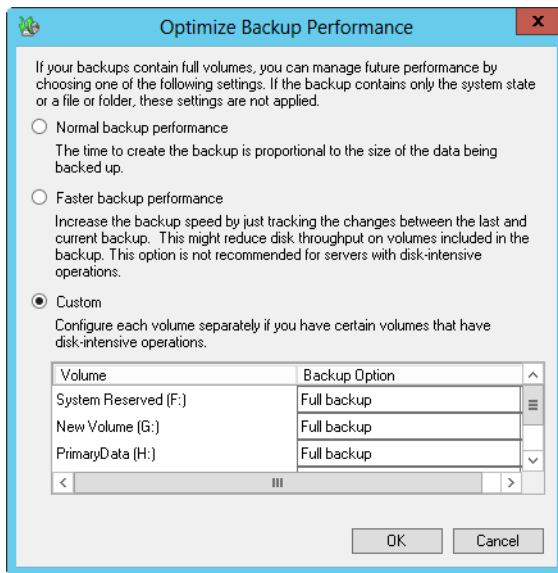
- **System state data** Includes essential system files needed to recover the local system. All computers have system state data, which must be backed up in addition to other files to restore a complete working system.
- **Application data** Includes application data files. You must back up application data if you want to be able to fully recover applications. Windows Server Backup creates block-level backups of application data using VSS.

Windows Server Backup allows you to perform full, copy, and incremental backups. Although you can schedule a full or incremental backup to be performed one or more times each day, you cannot use this feature to create separate run schedules for performing both full and incremental backups. Further, you cannot select the day or days of the week to run backups. This occurs because each server has a single master schedule that is set to run one or more times daily. If your servers have a single master schedule, you can work around this limitation by configuring Windows Server Backup to perform daily incremental backups and then creating a scheduled task via the Task Scheduler that uses Wbadmin to create a full backup on the day of the week or month you want to use.

When you use Windows Server Backup, the first backup of a server is always a full backup. This is because the full-backup process clears the archive bits on files so that Windows Server Backup can track which files are updated subsequently. Whether Windows Server Backup performs full or incremental backups subsequently depends on the default performance settings you configure. You can configure the default performance settings by following these steps:

1. Start Windows Server Backup. In the Actions pane or on the Action menu, tap or click Configure Performance Settings. This displays the Optimize Backup Performance dialog box, shown in Figure 13-2.
2. Do one of the following, and then tap or click OK:
  - Choose Normal Backup Performance to perform full backups of all attached drives.
  - Choose Faster Backup Performance to perform incremental backups of all attached drives.

- Choose Custom. In the option lists provided, specify whether to perform full or incremental backups for individual attached drives.



**FIGURE 13-2** Configure the default backup settings.

3. Once you configure the default performance settings, you can start a full or copy backup by tapping or clicking Backup Once on the Action menu or in the Actions pane. You can configure a backup schedule by tapping or clicking Backup Schedule on the Action menu or in the Actions pane.

## Getting Started with the Backup Command-Line Utility

Wbadmin is the command-line counterpart to Windows Server Backup. You use Wbadmin to manage all aspects of backup configuration that you would otherwise manage in Windows Server Backup. This means you can typically use either tool to manage backup and recovery.

After you install the Backup Command-Line Tools feature as discussed earlier in the chapter, you can use Wbadmin to manage backup and recovery. Wbadmin is located in the %SystemRoot%\System32\ directory. This directory is in your command path by default, so you do not need to add it. You can run Wbadmin by following these steps:

1. Open an elevated, administrator command prompt. One way to do this is to type **cmd** in the Apps Search box, press and hold or right-click Command Prompt in the Apps list, and then tap or click Run As Administrator.
2. In the Command Prompt window, enter the necessary command text or run a script that invokes Wbadmin.

Wbadmin has a number of associated commands, which are summarized in Table 13-2.

**TABLE 13-2** Wbadmin Management Commands

COMMAND	DESCRIPTION
DELETE SYSTEMSTATEBACKUP	Deletes the system state backup or backups from a specified location.
DISABLE BACKUP	Disables scheduled daily backups so that they no longer run.
ENABLE BACKUP	Enables or modifies a scheduled daily backup.
GET DISKS	Lists the disks that are currently online for the local computer. Disks are listed by manufacturer name, type, disk number, GUID, total space, used space, and associated volumes.
GET ITEMS	Lists items contained in a specified backup.
GET STATUS	Reports the status of the currently running backup or recovery job.
GET VERSIONS	Lists details about the available backups stored in a specific location, including the backup time and backup destination.
START BACKUP	Starts a one-time backup using the specified parameters. If no parameters are passed and scheduled backups are enabled, the backup uses the settings for scheduled backups.
START RECOVERY	Initiates a recovery of volumes, applications, or files using the specified parameters.
START SYSTEMSTATEBACKUP	Starts a system state backup using the options specified.
START SYSTEMSTATERECOVERY	Starts a system state recovery using the specified parameters.
STOP JOB	Stops the currently running backup or recovery job. Stopped jobs cannot be restarted from where they were stopped.

When you are working with Wbadmin, you can get help on available commands:

- To view a list of management commands, type **wbadmin /?** at the command prompt.
- To view the syntax for a specific management command, type **wbadmin *Command* /?**, where *Command* is the name of the management command you want to examine, such as **wbadmin stop job /?**.



When you work with Wbadmin, you'll find that just about every command accepts parameters and specific parameter values that qualify what you want to work with. To see more clearly how this works, consider the following syntax example:

```
wbadmin get versions [-backupTarget:{VolumeName | NetworkSharePath}]
[-machine:BackupMachineName]
```

The brackets tell you that *-backupTarget* and *-machine* are optional. Thus, you could type the following to get information on recoverable backups on the local computer:

```
wbadmin get versions
```

You could type the following to get information on recoverable backups stored on the F drive:

```
wbadmin get versions -backupTarget:f:
```

Or you could type the following to get information on recoverable backups stored on the F drive on Server96:

```
wbadmin get versions -backupTarget:f: -machine:server96
```

Many Wbadmin commands use the *-backupTarget* and *-machine* parameters. The backup target is the storage location you want to work with and can be expressed as a local volume name, such as F:, or as a network share path, such as \\FileServer32\backups\Server85. The *-machine* parameter identifies the computer you want to work with for backup or recovery operations.

## Working with Wbadmin Commands

You use Wbadmin commands to manage the backup configuration of your servers. These commands work with a specific set of parameters. The following sections provide an overview of the available commands and their most commonly used syntaxes.

### Using General-Purpose Commands

The following general-purpose commands are provided for getting information about backups and the system you are working with:

- **GET DISKS** Lists the disks that are currently online for the local computer. Disks are listed by manufacturer name, type, disk number, GUID, total space, used space, and associated volumes.

```
wbadmin get disks
```

- **GET ITEMS** Lists items contained in a specified backup.

```
wbadmin get items -version:VersionIdentifier
[-backupTarget:{VolumeName | NetworkSharepath}]
[-machine:BackupMachineName]
```

- **GET STATUS** Reports the status of the currently running backup or recovery job.

```
wbadmin get status
```

- **GET VERSIONS** Lists details about the available backups stored in a specific location, including the backup time and backup destination.

```
wbadmin get versions [-backupTarget:{VolumeName | NetworkSharepath}]
[-machine:BackupMachineName]
```

## Using Backup Management Commands

You can manage backups and their configurations using the following commands and command-line syntaxes:

- **DELETE SYSTEMSTATEBACKUP** Deletes the system state backup or backups from a specified location.

```
wbadmin delete systemstateBackup [-backupTarget:{VolumeName}]
[-machine:BackupMachineName]
[-keepVersions:NumberOfBackupsToKeep | -version:VersionIdentifier |
-delete0]dest]
[-quiet]
```

- **DISABLE BACKUP** Disables scheduled daily backups so that they no longer run.

```
wbadmin disable backup [-quiet]
```

- **ENABLE BACKUP** Enables or modifies a scheduled daily backup.

```
wbadmin enable backup [-addTarget:{BackupTargetDisk}]
[-removeTarget:{BackupTargetDisk}]
[-schedule:TimeToRunBackup]
[-include:VolumesToInclude]
[-allCritical]
[-quiet]
```

- **START BACKUP** Starts a one-time backup using the specified parameters. If no parameters are passed and scheduled backups are enabled, the backup uses the settings for scheduled backups.

```
wbadmin start backup [-backupTarget:{TargetVolume |
TargetNetworkShare}]
[-include:VolumesToInclude]
[-allCritical]
[-noVerify]
[-user:username]
[-password:password]
```

```
[-inheritAc1:InheritAc1]  
[-vssFull]  
[-quiet]
```

- **STOP JOB** Stops the currently running backup or recovery job. Stopped jobs cannot be restarted from where they were stopped.

```
wbadmin stop job [-quiet]
```

## Using Recovery Management Commands

You can recover your computers and data using the following commands and command-line syntaxes:

- **START RECOVERY** Initiates a recovery of volumes, applications, or files using the specified parameters.

```
wbadmin start recovery -version:VersionIdentifier  
-items:VolumesToRecover | AppsToRecover | FilesOrFoldersToRecover  
-itemType:{volume | app | file}  
[-backupTarget:{VolumeHostingBackup | NetworkShareHostingBackup}]  
[-machine:BackupMachineName]  
[-recoveryTarget:TargetVolumeForRecovery | TargetPathForRecovery]  
[-recursive]  
[-overwrite:{Overwrite | CreateCopy | skip}]  
[-notRestoreAc1]  
[-skipBadClusterCheck]  
[-noRollForward]  
[-quiet]
```

- **START SYSTEMSTATEBACKUP** Starts a system state backup using the options specified.

```
wbadmin start systemstateBackup -backupTarget:{VolumeName}  
[-quiet]
```

- **START SYSTEMSTATERECOVERY** Starts a system state recovery using the specified parameters.

```
wbadmin start systemstateRecovery -version:VersionIdentifier  
-showSummary  
[-backupTarget:{VolumeName | NetworkSharePath}]  
[-machine:BackupMachineName]  
[-recoveryTarget:TargetPathForRecovery]  
[-authSysvol]  
[-quiet]
```

## Performing Server Backups

---

As part of your planning for each server you plan to back up, you should consider which volumes you want to back up and whether backups will include system-state recovery data, application data, or both. Although you can manually back up to shared volumes and DVD media, you need a separate, dedicated hard disk for running scheduled backups. After you configure a disk for scheduled backups, the backup utilities automatically manage the disk usage and automatically reuse the space of older backups when creating new backups. Once you schedule backups, you need to check periodically to ensure that backups are being performed as expected and that the backup schedule meets current needs.

When you create or schedule backups, you need to specify the volumes you want to include, and this affects the ways you can recover your servers and your data. You have the following options:

- **Full server (all volumes with application data)** Back up all volumes with application data if you want to be able to fully recover a server, along with its system state and application data. Because you are backing up all files, the system state, and application data, you should be able to fully restore your server using only the Windows backup tools.
- **Full server (all volumes without application data)** Back up all volumes without application data if you want to be able to restore a server and its applications separately. With this technique, you back up the server using the Windows backup tools but exclude locations where applications and application data are stored. Then you back up applications and related data using third-party tools or tools built into the applications. You can fully recover a server by using the Windows backup utilities and then use a third-party utility to restore backups of applications and application data.
- **Critical volumes/bare metal recovery** Back up only critical volumes if you want to be able to recover only the operating system.
- **Noncritical volumes** Back up only individual volumes if you want to be able to recover only files, applications, or data from those volumes.

As part of the backup process, you also need to specify a storage location for backups. Keep the following in mind when you choose storage locations:

- When you use an internal hard disk for storing backups, you are limited in how you can recover your system. You can recover the data from a volume, but you cannot rebuild the entire disk structure.
- When you use an external hard disk for storing backups, the disk will be dedicated for storing your backups and will not be visible in File Explorer. Choosing this option will format the selected disk or disks, removing any existing data.
- When you use a remote shared folder for storing backups, your backup will be overwritten each time you create a new backup. Do not choose this option if you want to store multiple backups for each server.

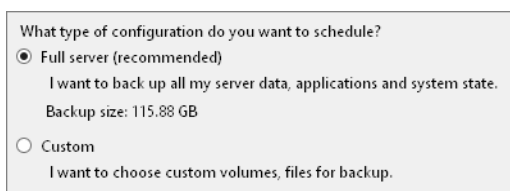
- When you use removable media or DVDs for storing backups, you can recover only entire volumes, not applications or individual files. The media you use must be at least 1 GB in size.

The sections that follow discuss techniques for performing backups. The procedures you use to back up servers with Windows Server Backup and Wbadmin are similar.

## Configuring Scheduled Backups

With Windows Server Backup, you can schedule automated backups for a server by following these steps:

1. In Windows Server Backup, tap or click Backup Schedule on the Action menu or in the Actions pane. This starts the Backup Schedule Wizard. Tap or click Next.
2. On the Select Backup Configuration page, note the backup size listed under the Full Server option, as shown in Figure 13-3. This is the storage space required to back up the server data, applications, and the system state. To back up all volumes on the server, select the Full Server option and then tap or click Next. To back up selected volumes on the server, tap or click Custom and then tap or click Next.

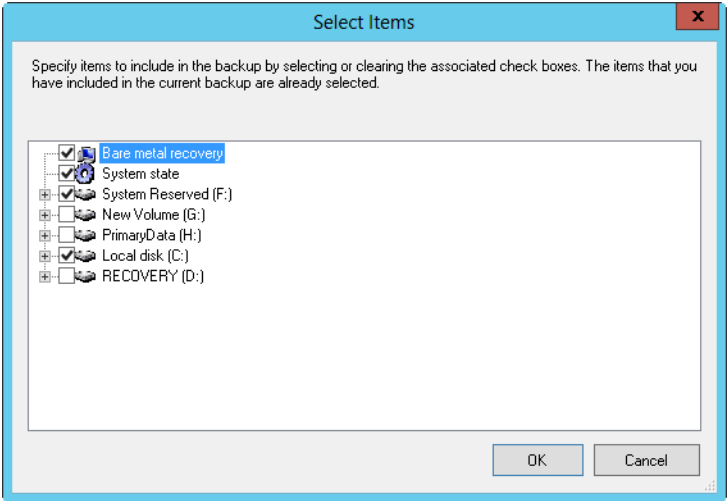


**FIGURE 13-3** Note the backup size.

**NOTE** Volumes that contain operating system files or applications are included in the backup by default and cannot be excluded. Unfortunately, this means that on a server on which you installed Windows Server 2012 on the D drive, you must also back up the entire C drive because the C drive in this case includes the boot manager and other boot files.

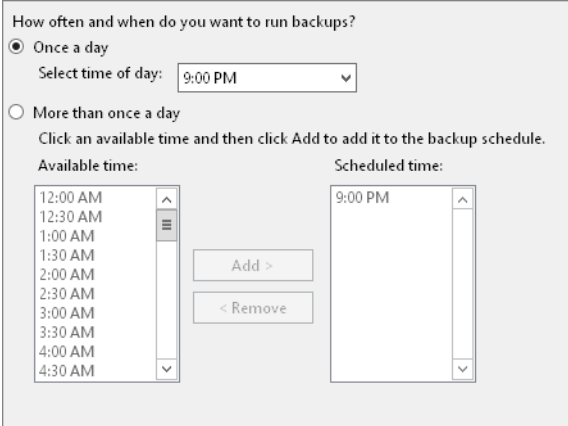
3. If you select Custom, the Select Items For Backup page is displayed. Tap or click Add Items. As shown in Figure 13-4, you can select the check boxes for the volumes you want to back up and clear the check boxes for the volumes you want to exclude. Select the Bare Metal Recovery option if you want to be able to fully recover the operating system. Tap or click OK, and then tap or click Next.

**TIP** After you select items, you might want to tap or click Advanced Settings before continuing. You can then use the options on the Exclusions tab to identify locations and file types that should not be backed up. You also can then use the options on the VSS Settings tab to specify whether you want to create a full backup or a copy backup.



**FIGURE 13-4** Select the items to include in the backup.

4. On the Specify Backup Time page, shown in Figure 13-5, you can specify how often and when you want to run backups. To perform backups daily at a specific time, choose **Once A Day** and then select the time to start running the daily backup. To perform backups multiple times each day, choose **More Than Once A Day**. Next, tap or click a start time under **Available Time**, and then tap or click **Add** to move the time under **Scheduled Time**. Repeat this process for each start time you want to add. Tap or click **Next** when you are ready to continue.



**FIGURE 13-5** Select the time to start running the backup.

5. On the Specify Destination Type page, you have these options:
  - **Back Up To A Hard Disk That Is Dedicated For Backups** Allows you to specify a dedicated hard disk for backups. Although you can use multiple disks for backups, any disk you select will be formatted and then dedicated only to backups. This option is recommended because you'll get the best performance. If you select this option, tap or click Next, select the disk or disks to use, and then tap or click Next again.
  - **Back Up To A Volume** Allows you to write backups to individual volumes on a hard disk. Because any volume you select is not dedicated to backups, it can be used for other purposes. However, the performance of any of the selected volumes is reduced while backups are being written. If you select this option, tap or click Next, use the Add and Remove options to select the volumes to use, and then tap or click Next again.
  - **Back Up To A Shared Network Folder** Allows you to specify a shared network folder for backups. With this option, you can have only one backup at a time because each new backup overwrites the previous backup. If you select this option, tap or click Next. When prompted, tap or click OK. Type the UNC path to the network share, such as `\\FileServer25\Backups\Exchange`. If you want the backup to be accessible to everyone who has access to the shared folder, select Inherit under Access Control. If you want to restrict access to the shared folder to the current user and members of the Administrators and Backup Operators groups, select Do Not Inherit under Access Control. Tap or click Next. When prompted to provide access credentials, type the user name and password for an account authorized to access and write to the shared folder.
6. On the Confirmation page, review the details and then tap or click Finish. The wizard formats the disk. The formatting process might take several minutes or considerably longer depending on the size of the disk.
7. On the Summary page, tap or click Close. Your backups are now scheduled for the selected server.

With Wbadmin, you can schedule backups using the ENABLE BACKUP command. ENABLE BACKUP accepts the following parameters:

- **-addTarget** Sets the storage location for backups according to the GUID of the disk you want to use. The GUID of a disk is listed as the disk identifier in the output of the Wbadmin GET DISKS command.
- **-removeTarget** Sets the storage location to remove from the backup schedule according to the GUID of the disk you want to use. The GUID of a disk is listed as the disk identifier in the output of the Wbadmin GET DISKS command.
- **-include** Sets a comma-delimited list of volume drive letters, volume mount points, and GUID volume names to back up.
- **-allCritical** Includes all operating system volumes in the backup automatically.
- **-quiet** Specifies that you want to run the command with no prompts to the user.

To see how ENABLE BACKUP is used, consider the following examples:

#### Schedule a backup for C and D at 9:00 P.M. daily

```
wbadmin enable backup -addTarget:{06d88776-0000-0000-0000-000000000000}  
-schedule:18:00 -include:c:,d:
```

#### Schedule a backup for all operating system volumes at 6:00 A.M. and 9:00 P.M. daily

```
wbadmin enable backup -addTarget:{06d88776-0000-0000-0000-000000000000}  
-schedule:06:00,18:00 -allCritical
```

## Modifying or Stopping Scheduled Backups

Once you've configured scheduled backups on a server, you can modify or stop the scheduled backups by following these steps:

1. Start Windows Server Backup. Tap or click Backup Schedule on the Action menu or in the Actions pane. This starts the Backup Schedule Wizard. Tap or click Next.
2. On the Modify Scheduled Backup Settings page, tap or click Modify Backup if you want to add or remove backup items, times, or targets, and then follow steps 3 to 10 in the section "Configuring Scheduled Backups" earlier in the chapter. If you want to stop the scheduled backups from running, tap or click Stop Backup, tap or click Next, and then tap or click Finish. When prompted to confirm, tap or click Yes and then tap or click Close.

**NOTE** Stopping backups releases backup disks for normal use. Backup archives are not deleted from the backup disks and remain available for use in recovery.

With Wbadmin, you can modify scheduled backups using the ENABLE BACKUP command. For targets, you must use the *-addTarget* and *-removeTarget* parameters to modify the target disks. For the run schedule and included volumes, you simply set the new values you want to use. Consider the following examples:

#### Adding a new target to scheduled backups

```
wbadmin enable backup -addTarget:{41cd2567-0000-0000-0000-000000000000}
```

#### Removing a target from scheduled backups

```
wbadmin enable backup -removeTarget:{06d88776-0000-0000-0000-000000000000}
```

#### Modifying the run schedule and included volumes

```
wbadmin enable backup -schedule:03:00 -include:c:,d:,e:
```

## Creating and Scheduling Backups with Wbadmin

One way to create manual backups is to use the Wbadmin START BACKUP command. START BACKUP accepts the following parameters:

- **-backupTarget** Sets the storage location for the backup as either a drive letter or UNC path to a shared folder on a remote server.



- **-include** Sets a comma-delimited list of volume drive letters, volume mount points, and GUID volume names to back up.
- **-allCritical** Includes all operating system volumes in the backup automatically.
- **-inheritAcl** Specifies that you want the backup folder at the remote shared folder to inherit the security permissions of the shared folder. If you do not specify this parameter, the backup folder is accessible only to the user you specify in the `-user` parameter, administrators, and backup operators.
- **-noVerify** Specifies that you do not want to verify backups written to removable media. If you do not specify this parameter, backups written to removable media are verified.
- **-password** Sets the password to use when connecting to the remote shared folder.
- **-quiet** Specifies that you want to run the command with no prompts to the user.
- **-user** Sets the user name to use when connecting to the remote shared folder.
- **-vssFull** Specifies that you want to perform a full backup using VSS, which ensures that all server and application data is backed up. Do not use this parameter if you are using a third-party backup utility to back up application data.

To see how START BACKUP is used, consider the following examples:

#### Performing a full backup of the server

```
wbadmin start backup -backupTarget:f: -vssfull
```

#### Backing up C and D to F

```
wbadmin start backup -backupTarget:f: -include:c:,d:
```

#### Backing up all critical volumes

```
wbadmin start backup -backupTarget:f: -allCritical
```

#### Backing up C and D to a remote shared folder

```
wbadmin start backup -backupTarget:\\filesrv27\backups -include:c:,d:
-user:williams
```

If you want to create a schedule to run backups at different times on different days, you can use Task Scheduler to create the necessary tasks to run this command on the schedule you set. You can use Task Scheduler and Wbadmin to schedule tasks to run backups by following these steps:

1. In Computer Management, tap or click Task Scheduler. You are connected to the local computer by default. As necessary, connect to the computer that you want to access.
2. Press and hold or right-click the Task Scheduler node, and then tap or click Create Task. This opens the Create Task dialog box.

3. On the General tab, type the name of the task and then set security options for running the task.
  - If the task should run under a user other than the current user, tap or click Change User Or Group. In the Select User Or Group dialog box, select the user or group under which the task should run, and then provide the appropriate credentials when prompted later.
  - Set other run options as necessary. By default, tasks run only when a user is logged on. If you want to run the task regardless of whether a user is logged on, select Run Whether User Is Logged On Or Not. You can also elect to run with highest privileges and configure the task for earlier releases of Windows.
4. On the Triggers tab, tap or click New. In the New Trigger dialog box, select On A Schedule in the Begin The Task list. Use the options provided to set the run schedule, and then tap or click OK.
5. On the Actions tab, tap or click New. In the New Action dialog box, select Start A Program in the Action list.
6. In the Program/Script text box, type **%windir%\System32\wbadmin.exe**.
7. In Add Arguments, type the START BACKUP command you want to use along with its parameters, such as the following:
 

```
start backup -backupTarget:f: -include:c:,d:,e:\mountpoint,
  \\?\volume{be345a23-32b2-432d-43d2-7867ff3e3432}\
```
8. Tap or click OK to close the New Action dialog box.
9. On the Conditions tab, specify any limiting conditions for starting or stopping the task.
10. On the Settings tab, choose any additional optional settings for the task.
11. Tap or click OK to create the task.

## Running Manual Backups

You can use Windows Server Backup to back up servers manually by following these steps:

1. Start Windows Server Backup. Tap or click Backup Once on the Action menu or in the Actions pane. This starts the Backup Once Wizard.
2. If you want to back up the server using the same options you use for the Backup Schedule Wizard, select Scheduled Backup Options, tap or click Next, and then tap or click Backup to perform the backup. Skip the remaining steps.
3. If you want to back up the server using different options, select Different Options and then tap or click Next.
4. On the Select Backup Configuration page, note the backup size listed under the Full Server option. This is the storage space required to back up the server data, applications, and the system state. To back up all volumes on the server, select the Full Server option and then tap or click Next. To back up selected volumes on the server, tap or click Custom and then tap or click Next.

5. If you select Custom, the Select Items For Backup page is displayed. Tap or click Add Items. Select the check boxes for the volumes you want to back up, and clear the check boxes for the volumes you want to exclude. Select the Bare Metal Recovery option if you want to be able to fully recover the operating system. Tap or click OK, and then tap or click Next.

**TIP** After you select items, you might want to tap or click Advanced Settings before continuing. You can then use the options on the Exclusions tab to identify locations and file types that should not be backed up. You also can then use the options on the VSS Settings tab to specify whether you want to create a full backup or a copy backup.

6. On the Specify Destination Type page, do one of the following:
  - If you want to back up to local drives, select Local Drives and then tap or click Next. On the Backup Destination page, select the internal or external disk or DVD drive to use as the backup target. When stored on a DVD, backups are compressed and you can recover only full volumes. As a result, the size of the backup on a DVD might be smaller than the volume on the server. Tap or click Next.
  - If you want to back up to a remote shared folder, select Remote Shared Folder and then tap or click Next. On the Specify Remote Folder page, type the UNC path to the remote folder, such as \\FileServer43\Backups. If you want the backup to be accessible to everyone who has access to the shared folder, select Inherit under Access Control. If you want to restrict access to the shared folder to the current user, administrators, and backup operators, select Do Not Inherit under Access Control. Tap or click Next. When prompted to provide access credentials, type the user name and password for an account authorized to access and write to the shared folder.
7. Tap or click Next, and then tap or click Backup. The Backup Progress dialog box shows you the progress of the backup process. If you tap or click Close, the backup will continue to run in the background.

## Recovering Your Server from Hardware or Startup Failure

Windows Server 2012 includes an extensive diagnostics and resolution architecture. These features can help you recover from many types of hardware, memory, and performance issues and either resolve them automatically or help users through the process of resolving them.

Windows Server 2012 includes more reliable and better-performing device drivers to prevent many common causes of hangs and crashes. Improved input/output (I/O) cancellation for device drivers ensures that the operating system can recover gracefully from blocking calls and that there are fewer blocking disk I/O operations.

To reduce downtime and restarts required for application installations and updates, Windows Server 2012 can use the update process to mark in-use files for update and then automatically replace the files the next time the application is started. In some cases, Windows Server 2012 can save the application's data, close

the application, update the in-use files, and then restart the application. To improve overall system performance and responsiveness, Windows Server 2012 uses memory efficiently, provides ordered execution for groups of threads, and provides several process-scheduling mechanisms. By optimizing memory and process usage, Windows Server 2012 ensures that background processes have less performance impact on system performance.

Windows Server 2012 provides improved guidance on the causes of unresponsive conditions. By including additional error-reporting details in the event logs, Windows Server 2012 makes it easier to identify and resolve issues. To automatically recover from service failures, Windows Server 2012 uses service recovery policies more extensively than its predecessors. When recovering a failed service, Windows Server 2012 automatically handles both service and nonservice dependencies. Any necessary dependent services and system components are started prior to starting the failed service.

In earlier versions of Windows, an application crash or hang is marked as not responding, and it is up to the user to exit and then restart the application. Windows Server 2012 attempts to resolve the issue of unresponsive applications by using Restart Manager. Restart Manager can shut down and restart unresponsive applications automatically. Thanks to Restart Manager, you might not have to intervene to try to resolve issues with frozen applications.

Failed installation and nonresponsive conditions of applications and drivers are also tracked through Action Center, and the built-in diagnostics display a warning message. By tapping or clicking the Action Center icon in the system tray, you can view recent messages. If you tap or click a message, Windows Server 2012 opens the Message Details page in Action Center, which might provide a solution for the problem.

You also can view a list of current problems at any time by following these steps:

1. In Control Panel, under the System And Security heading, tap or click Review Your Computer's Status.
2. In Action Center, a list of known problems is displayed. For some issues, you'll be able to select a related View Message Details button to display a Message Details page. If a solution is available, tap or click the link provided to download the solution or visit a related website to get more information.

While you are working with Action Center, you can have Windows Server check for solutions to problems by tapping or clicking the Check For Solutions link on the Maintenance panel.

Windows Server 2012 attempts to resolve issues related to running out of virtual memory by providing Resource Exhaustion Detection And Recovery. This feature monitors the systemwide virtual memory commit limit and alerts you if the computer is running low on virtual memory. To help you to correct this issue, it also identifies the processes consuming the largest amount of memory, allowing you to close any or all of these high resource-consuming applications directly from the Close Programs To Prevent Information Loss dialog box. The resource exhaustion alert is also logged in the system event log.

In early versions of Windows, corrupted system files are one of the most common causes of startup failure. Windows Server 2012 includes built-in diagnostics to automatically detect corrupted system files during startup and guide you through automated or manual recovery. To resolve startup problems, Windows Server 2012 uses the Startup Repair tool (StR), which is installed automatically and started when a system fails to boot. Once started, StR attempts to determine the cause of the startup failure by analyzing startup logs and error reports. Then StR attempts to fix the problem automatically. If StR is unable to resolve the problem, it restores the system to the last known working state and then provides diagnostic information and support options for further troubleshooting.

Hardware problems addressed by built-in diagnostics include error detection and disk failure detection. If a device is having problems, hardware diagnostics can detect error conditions and either repair the problem automatically or guide the user through a recovery process. With disk drives, hardware diagnostics can use fault reports provided by disk drives to detect potential failures and alert you before they happen. Hardware diagnostics can also help guide you through the backup process after alerting you that a disk might be failing.

Performance problems addressed by built-in diagnostics include slow application startup, slow boot, slow standby/resume, and slow shutdown. If a computer is experiencing degraded performance, performance diagnostics can detect the problem and provide possible solutions for resolving the problem. For advanced performance issues, you can track related performance and reliability data in the Performance Diagnostics console, which includes a performance monitor and a reliability monitor. (This is discussed in Chapter 3, “Monitoring Processes, Services, and Events.”)

Memory problems addressed by built-in diagnostics include both memory leaks and failing memory. A memory leak occurs if an application or a system component doesn't completely free areas of physical memory after it is done with them. If you suspect that a computer has a memory problem that is not being automatically detected, you can run Windows Memory Diagnostics manually during startup by selecting the related option. If the Windows Memory Diagnostics option is not provided during startup, you can run the program by following these steps:

1. Start Windows Memory Diagnostics. One way to do this is to type **mdsched.exe** in the App Search box and then press Enter.
2. Choose whether to restart the computer now and run the tool immediately or schedule the tool to check for problems at the next restart.
3. Windows Memory Diagnostics runs automatically after the computer restarts, using the standard test mix and performing two test passes by default.

You can change the run options using the F1 key. Three different levels of memory testing can be performed, including Basic, Standard, and Extended. Use a basic test to quickly check the memory. Use the standard test to perform a standard test of the memory. Use the extended test when you want to perform more extensive testing. Set the number of test passes using the Pass Count option.

To detect system crashes possibly caused by failing memory, memory diagnostics work with the Microsoft Online Crash Analysis tool. If a computer crashes because of

failing memory, and memory diagnostics detect this, you are prompted to schedule a memory test the next time the computer is restarted.

## Recovering from a Failed Start

Windows Server 2012 enters Windows Error Recovery mode automatically if Windows fails to start. In this mode, you'll see a Recovery screen the next time you try to start the server. The options will include the following:

- **Continue** Exits the repair menu, and continues to load the operating system
- **Use Another Operating System** Exits the repair menu, and allows you to select the operating system to load (if multiple operating systems are installed)
- **Troubleshoot** Displays the Advanced Options menu
- **Turn Off Your PC** Exits the repair menu, and shuts down the server

The Advanced Options menu has three options:

- **System Image Recovery** Allows you to recover the server using a system image file. The image file can come from a remote computer.
- **Command Prompt** Allows you to access a command prompt and work with the commands and tools available in the recovery environment.
- **Startup Settings** Allows you to change the startup behavior and start the server in safe mode. Here, you click Restart to restart the computer in safe mode so that you can disable driver signature enforcement, early-launch antimalware protection, and automatically restart on system failure. It also allows you to enable low-resolution video mode, debugging mode, boot logging, and safe mode.

## Starting a Server in Safe Mode

Many startup problems occur because something on the system has changed; for example, a device might have been incorrectly installed. The system configuration or registry might have been updated improperly, causing a conflict. Often you can resolve startup issues using safe mode to recover or troubleshoot system problems. When you have finished using safe mode, be sure to restart the server using a normal startup. You will then be able to use the server as you normally would.

In safe mode, Windows Server 2012 loads only basic files, services, and drivers. The drivers loaded include those for the mouse, monitor, keyboard, mass storage, and base video. The monitor driver sets the basic settings and modes for the server's monitor; the base video driver sets the basic options for the server's graphics card. No networking services or drivers are started unless you choose the Safe Mode With Networking option. Because safe mode loads a limited set of configuration information, it can help you troubleshoot problems.

You can start a server in safe mode by following these steps:

1. If the computer won't start normally, the Recovery screen is displayed during startup. On the Recovery screen, tap or click Troubleshoot.

2. On the Advanced Options screen, tap or click Startup Settings. Next, on the Windows Startup Settings screen, tap or click Restart.
3. Use the arrow keys to select the safe mode you want to use, and then press Enter. The safe mode option you use depends on the type of problem you're experiencing. The key options are as follows:
  - **Repair Your Computer** Loads the Startup Repair tool. Choose this option to restart the server and go back to the Recovery screen.
  - **Safe Mode** Loads only basic files, services, and drivers during the initialization sequence. The drivers loaded include the mouse, monitor, keyboard, mass storage, and base video. No networking services or drivers are started.
  - **Safe Mode With Networking** Loads basic files, services, and drivers, as well as services and drivers needed to start networking.
  - **Safe Mode With Command Prompt** Loads basic files, services, and drivers, and then starts a command prompt instead of the Windows graphical interface. No networking services or drivers are started.

**TIP** In Safe Mode With Command Prompt, you can start the Explorer shell from the command-line interface by pressing **Ctrl+Shift+Esc** and typing **explorer.exe** in the New Process window on the File menu of Task Manager.

- **Enable Boot Logging** Allows you to create a record of all startup events in a boot log.
- **Enable Low-Resolution Video** Allows you to start the system in low-resolution 640 by 480 display mode, which is useful if the system display is set to a mode that can't be used with the current monitor.
- **Last Known Good Configuration** Starts the computer in safe mode using registry information that Windows saved at the last shutdown, including the HKEY\_CURRENT\_CONFIG (HKCC) hive. This registry hive stores information about the hardware configuration with which you previously and successfully started the computer.
- **Debugging Mode** Starts the system in debugging mode, which is useful only for troubleshooting operating system bugs.
- **Directory Services Restore Mode** Starts the system in safe mode, and allows you to restore the directory service. This option is available on Windows Server 2008 R2 and later domain controllers.
- **Disable Automatic Restart On System Failure** Prevents Windows Server from automatically restarting after an operating system crash.
- **Disable Driver Signature Enforcement** Starts the computer in safe mode without enforcing digital signature policy settings for drivers. If a driver with an invalid or missing digital signature is causing startup failure, this option resolves the problem temporarily so that you can start the computer and resolve the problem by getting a new driver or changing the driver signature enforcement settings.

- **Disable Early Launch Anti-Malware Driver** Starts the computer in safe mode without running the boot driver for the computer's antimalware software. If the boot driver for the computer's antimalware software is preventing startup, you need to check the software developer's website for an update that resolves the boot problem or configure the software without boot protection.
  - **Start Windows Normally** Starts the computer with its regular settings.
4. If a problem doesn't reappear when you start in safe mode, you can eliminate the default settings and basic device drivers as possible causes. If a newly added device or updated driver is causing problems, you can use safe mode to remove the device or reverse the update.

## Backing Up and Restoring the System State

In Windows Server 2012, there are approximately 50,000 system state files, which use approximately 4 GB of disk space in the default installation of an x64-based computer. The fastest and easiest way to back up and restore a server's system state is to use Wbadmin. With Wbadmin, you can use the `START SYSTEMSTATEBACKUP` command to create a backup of the system state for a computer and the `START SYSTEMSTATERECOVERY` command to restore a computer's system state.

**TIP** When you select a system state restore on a domain controller, you have to be in the Directory Services Restore mode. To learn how to restore Active Directory, see the next section.

To back up a server's system state, type the following at an elevated command prompt:

```
wbadmin start systemstatebackup -backupTarget:VolumeName
```

Here *VolumeName* is the storage location for the backup, such as F:.

To restore a server's system state, type the following at an elevated command prompt:

```
wbadmin start systemstaterecovery -backupTarget:VolumeName
```

Here *VolumeName* is the storage location that contains the backup you want to recover, such as F:. Additionally, you can do the following:

- Use the `-recoveryTarget` parameter to restore to an alternate location.
- Use the `-machine` parameter to specify the name of the computer to recover if the original backup location contains backups for multiple computers.
- Use the `-authSysvol` parameter to perform an authoritative restore of the SYSVOL.

You can also recover the system state by using a backup that includes the system state or by performing a recovery.



## Restoring Active Directory

When restoring system state data to a domain controller, you must choose whether you want to perform an authoritative or nonauthoritative restore. The default is nonauthoritative. In this mode, Active Directory and other replicated data are restored from backup and any changes are replicated from another domain controller. Thus, you can safely restore a failed domain controller without overwriting the latest Active Directory information. On the other hand, if you're trying to restore Active Directory throughout the network using archived data, you must use an authoritative restore. With an authoritative restore, the restored data is restored on the current domain controller and then replicated to other domain controllers.

**CAUTION** An authoritative restore overwrites all Active Directory data throughout the domain. Before you perform an authoritative restore, you must be certain that the archive data is the correct data to propagate throughout the domain and that the current data on other domain controllers is inaccurate, outdated, or otherwise corrupted.

To restore Active Directory on a domain controller and enable the restored data to be replicated throughout the network, follow these steps:

1. Make sure the domain controller server is shut down.
2. Restart the domain controller server, and enter safe mode.
3. Select Directory Services Restore Mode.
4. When the system starts, use the Backup utility to restore the system state data and other essential files.
5. After restoring the data but before restarting the server, use the Ntdsutil.exe tool to mark objects as authoritative. Be sure to check the Active Directory data thoroughly.
6. Restart the server. When the system finishes startup, the Active Directory data should begin to replicate throughout the domain.

## Restoring the Operating System and the Full System

As discussed previously, Windows Server 2012 includes startup repair features that can recover a server in case of corrupted or missing system files. The startup repair process can also recover from some types of boot failures involving the boot manager. If these processes fail and the boot manager is the reason you cannot start the server, you can use the Windows Server 2012 installation disc or system recovery options to restore the boot manager and enable startup.

System recovery options are available only with full server installations and not with Server Core installations. With Server Core installations, you need to use the installation disc to initiate recovery.

System recovery options include the following tools:

- **System Image Recovery** Allows you to recover a server's operating system or perform a full system recovery. With an operating system or full system recovery, make sure your backup data is available and that you can log on

with an account that has the appropriate permissions. With a full system recovery, keep in mind that data that was not included in the original backup will be deleted when you recover the system, including any in-use volumes that were not included in the backup.

- **Windows Memory Diagnostics Tools** Allows you to diagnose a problem with the server's physical memory. Three different levels of memory testing can be performed: basic, standard, and exhaustive.

You can also access a command prompt. This command prompt gives you access to the command-line tools available during installation as well as to these additional programs:

- **Startup Repair Wizard (X:\Sources\Recovery\StartRep.exe)** Normally, this tool is started automatically on boot failure if Windows detects an issue with the boot sector, boot manager, or Boot Configuration Data (BCD) store.
- **Startup Recovery Options (X:\Sources\Recovery\Recenv.exe)** Allows you to start the Startup Recovery Options Wizard. If you previously entered the wrong recovery settings, you can provide different options.

As an administrator, you can perform command-line troubleshooting by following these steps:

1. If the computer won't start normally, the Recovery screen is displayed during startup. On the Recovery screen, tap or click Troubleshoot.
2. On the Advanced Options screen, tap or click Command Prompt.
3. When prompted to choose an account, tap or click the Administrator account. Next, enter the password for the Administrator account and tap or click Continue.
4. Use the command prompt to perform troubleshooting. For example, you could run the Startup Repair Wizard by entering **x:\sources\recovery\startrep.exe**.

You can recover a server's operating system or perform a full system recovery by using a backup image you created earlier with Windows Server Backup. With an operating system recovery, you recover all critical volumes but do not recover nonsystem volumes. If you recover your full system, Windows Server Backup reformats and repartitions all disks that were attached to the server. Therefore, you should use this method only when you want to recover the server data onto separate hardware or when all other attempts to recover the server on the existing hardware have failed.

**NOTE** When you recover the operating system or the full system, make sure that your backup data is available and that you can log on with an account that has the appropriate permissions. With a full system recovery, keep in mind that existing data that was not included in the original backup will be deleted when you recover the system. This includes any volumes that are currently used by the server but were not included in the backup.

You can recover the operating system using a backup image by following these steps:

1. If the computer won't start normally, the Recovery screen is displayed during startup. On the Recovery screen, tap or click Troubleshoot.
2. On the Advanced Options screen, tap or click System Image Recovery.
3. When prompted to choose an account, tap or click the Administrator account. Next, enter the password for the Administrator account and tap or click Continue. This starts the Re-Image Your Computer Wizard.
4. On the Select A System Image Backup page, tap or click Use The Latest Available System Image (Recommended) and then tap or click Next. Or tap or click Select A System Image, and then tap or click Next.
5. If you select an image to restore, do one of the following on the Select The Location Of The Backup page:
  - Tap or click the location that contains the system image you want to use, and then tap or click Next. Afterward, tap or click the system image you want to use, and then tap or click Next.
  - To browse for a system image on the network, tap or click Advanced and then tap or click Search For A System Image On The Network. When you are prompted to confirm that you want to connect to the network, tap or click Yes. In the Network Folder text box, specify the location of the server and shared folder in which the system image is stored, such as **\\FileServer22\Backups**, and then tap or click OK.
  - To install a driver for a backup device that doesn't show up in the location list, tap or click Advanced and then tap or click Install A Driver. Insert the installation media for the device, and then tap or click OK. After Windows installs the device driver, the backup device should be listed in the location list.
6. On the Choose Additional Restore Options page, do the following optional tasks and then tap or click Next:
  - Select the Format And Repartition Disks check box to delete existing partitions and reformat the destination disks to be the same as the backup.
  - Select Only Restore System Drives to restore only the drives from the backup that are required to run Windows: the boot, system, and recovery volumes. If the server has data drives, they will not be restored.
  - Tap or click Install Drivers to install device drivers for the hardware to which you are recovering.
  - Tap or click Advanced to specify whether the computer is restarted and the disks are checked for errors immediately after the recovery operation is completed.
7. On the Confirmation page, review the details for the restoration and then tap or click Finish. The wizard then restores the operating system or the full server as appropriate for the options you selected.

# Restoring Applications, Nonsystem Volumes, and Files and Folders

Windows Server 2012 provides separate processes for system state and full server recovery and the recovery of individual volumes and files and folders. You can use the Recovery Wizard in Windows Server Backup to recover nonsystem volumes and files and folders from a backup. Before you begin, you should be sure that the computer you are recovering files to is running Windows Server 2012. If you want to recover individual files and folders, you should be sure that at least one backup exists on an internal or external disk or in a remote shared folder. You cannot recover files and folders from backups saved to DVDs or removable media.

With this in mind, you can recover nonsystem volumes, files and folders, or application data by following these steps:

1. Start Windows Server Backup. In the Actions pane or on the Action menu, tap or click Recover. This starts the Recovery Wizard.
2. On the Getting Started page, specify whether you will recover data from the local computer or from another location and then tap or click Next.
3. If you are recovering data from another location, specify whether the backup you want to restore is on a local drive or in a remote shared folder, tap or click Next, and then specify location-specific settings. When you are recovering from a local drive, on the Select Backup Location page, select the location of the backup from the drop-down list. When you are recovering from a remote shared folder, on the Specify Remote Folder page, type the path to the folder that contains the backup. In the remote folder, the backup should be stored at `\\BackupServer\WindowsImageBackup\ComputerName`.
4. If you are recovering from another location, on the Select Server page, select which server's data you would like to recover. Tap or click Next.
5. On the Select Backup Date page, select the date and time of the backup you want to restore using the calendar and the time list. Backups are available for dates shown in bold. Tap or click Next.
6. On the Select Recovery Type page, do one of the following:
  - To restore individual files and folders, tap or click Files And Folders and then tap or click Next. On the Select Items To Recover page, under Available Items, tap or click the plus sign (+) to expand the list until the folder you want is visible. Tap or click a folder to display the contents of the folder in the adjacent pane, tap or click each item you want to restore, and then tap or click Next.
  - To restore noncritical, nonoperating system volumes, tap or click Volumes and then tap or click Next. On the Select Volumes page, you'll see a list of source and destination volumes. Select the check boxes associated with the source volumes you want to recover, and then select the location to which you want to recover the volumes by using the Destination Volume lists. Tap or click Next. If prompted to confirm the recovery operation, tap or click Yes. Skip steps 7 and 8.

- To restore application data, tap or click Applications and then tap or click Next. On the Select Application page, under Applications, tap or click the application you want to recover. If the backup you are using is the most recent, you might see a check box labeled Do Not Perform A Roll-Forward Recovery Of The Application Databases. Select this check box if you want to prevent Windows Server Backup from rolling forward the application database that is currently on your server. Tap or click Next. Because any data on the destination volume will be lost when you perform the recovery, make sure that the destination volume is empty or does not contain information you will need later.
7. Next, you can specify whether you want to restore data to its original location (nonsystem files only) or an alternate location. For an alternate location, type the path to the restore location or tap or click Browse to select it. With applications, you can copy application data to an alternate location. You cannot, however, recover applications to a different location or computer.
  8. For file and folder recovery, choose a recovery technique to apply when files and folders already exist in the recovery location. You can create copies so that you have both versions of the file or folder, overwrite existing files with recovered files, or skip duplicate files and folders to preserve existing files. You also can restore the original security permissions to files and folders being recovered.
  9. On the Confirmation page, review the details and then tap or click Recover to restore the specified items.

## Managing Encryption Recovery Policy

---

If you're an administrator for an organization that uses the Encrypting File System (EFS), your disaster-recovery planning must include additional procedures and preparations. You need to consider how to handle issues related to personal encryption certificates, EFS recovery agents, and EFS recovery policy. These issues are discussed in the sections that follow.

### Understanding Encryption Certificates and Recovery Policy

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.

Every file that's encrypted has a unique encryption key. This means that encrypted files can be copied, moved, and renamed just like any other file—and in most cases these actions don't affect the encryption of the data. The user who encrypted the file always has access to the file, provided that the user's private key is available in the user's profile on the computer or the user has credential roaming with Digital Identification Management Service (DIMS). 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 using a public/private key that EFS generates automatically on a per-user basis. By default, Windows XP SP1 and later releases of Windows use the Advanced Encryption Standard (AES) algorithm for encrypting files with EFS. AES is not supported on Windows 2000 or Windows XP versions prior to SP1, and AES-encrypted files viewed on these computers can appear to be corrupted when in fact they are not. Internet Information Services 7 and later can use an AES provider for encrypting passwords by default.

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.

**TIP** An alternative to a roaming profile is to copy the user's encryption certificate to the computers the user uses. You can do this by using the certificate backup and restore process discussed in "Backing Up and Restoring Encrypted Data and Certificates" later in this chapter. 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 in which this occurs is when a user leaves the company and the associated user account is deleted. Although a manager might have been able to log on to the user's account, check files, and save important files to other folders, encrypted files will be accessible afterward only if the encryption is removed by the manager acting as the user who encrypted the files or, if while logged on as the user, the manager moves the files to a 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 that's necessary to unlock data in encrypted files. However, to protect sensitive data, recovery agents don't have access to a user's private key or any private key information.

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 2012 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 stand-alone configuration, the recovery agent is the administrator of the local computer by default. You can designate additional recovery agents. Further, if you want local recovery agents in a domain environment rather than domain-level recovery agents, you must delete the recovery policy from the Group Policy for the domain.

You can delete recovery policies if you don't want them to be available.

## Configuring the EFS Recovery Policy

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.

Through Group Policy, you can view, assign, and delete recovery agents. Follow these steps:

1. Access the Group Policy console for the local computer, site, domain, or organizational unit you want to work with. For details on working with Group Policy, see Chapter 4, "Automating Administrative Tasks, Policies, and Procedures."
2. Expand Computer Configuration, Windows Settings, Security Settings, and Public Key Policies, and then tap or click Encrypting File System to access the configured Recovery Agents in Group Policy.
3. The pane at the right lists the recovery certificates currently assigned. Recovery certificates are listed according to who they are issued to, who issued them, their expiration date and purpose, and more.
4. To designate an additional recovery agent, press and hold or right-click Encrypting File System and then tap or 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. Tap or click Next. On the Select Recovery Agents page, tap or click Browse Directory, and in the Find Users, Contacts, And Groups dialog box, select the user you want to work with. Tap or click OK, and then tap or click Next. Tap or click Finish to add the recovery agent.

**NOTE** Before you can designate additional recovery agents, you should set up a root certificate authority (CA) in the domain. Afterward, you must 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. You can also use Cipher.exe to generate the EFS recovery agent key and certificate.

5. 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, tap or click Yes to permanently and irrevocably delete the certificate. If the recovery policy is empty (meaning it has no other designated recovery agents), EFS is turned off so that users can no longer encrypt files.

## Backing Up and Restoring Encrypted Data and Certificates

---

You can back up and restore encrypted data like you can any other data. The key thing to remember is that you must use backup software that understands EFS, such as the built-in backup and restore tools. You must be careful when using this type of software, however.

The backup or restore process doesn't necessarily back up or restore the certificate needed to work with the encrypted data. The user's profile data contains that certificate. If the user's account exists and the profile still contains the necessary certificate, the user can still work with the encrypted data.

If the user's account exists and you previously backed up the user's profile and then restored the profile to recover a deleted certificate, the user can still work with the encrypted data. Otherwise, there's no way to work with the data, and you need to have a designated recovery agent access the files and then remove the encryption.

Being able to back up and restore certificates is an important part of any disaster-recovery plan. The next sections examine the techniques you can use to perform these tasks.

### Backing Up Encryption Certificates

You use the Certificates snap-in to back up and restore personal certificates. Personal certificates are saved with the Personal Information Exchange (.pfx) format.

To back up personal certificates, follow these steps:

1. Log on as the user to the computer where the personal certificate you want to work with is stored. Tap or click Start, type **mmc** in the Search box, and then press Enter. This opens the Microsoft Management Console (MMC).
2. In the MMC, select File, and then select Add/Remove Snap-In. This opens the Add Or Remove Snap-Ins dialog box.
3. In the Available Snap-Ins list, select Certificates, and then tap or click Add. Select My User Account, and then tap or click Finish. This adds the Certificates snap-in to the Selected Snap-Ins list. The focus for the snap-in is set to the currently logged-on user account.
4. Tap or click OK to close the Add Or Remove Snap-Ins dialog box.
5. Expand Certificates—Current User, expand Personal, and then select Certificates. Press and hold or right-click the certificate you want to save, tap or click All Tasks, and then tap or click Export. This starts the Certificate Export Wizard. Tap or click Next.
6. Select Yes, Export The Private Key. Tap or click Next twice.
7. On the security page, use the options provided to specify security principals that should have access to the certificate. The default security principal is the Administrator account. Afterward, type and confirm a password for opening the certificate. Tap or click Next.



8. Tap or click Browse. Use the dialog box provided to specify a file location for the certificate file and then tap or click Save. Be sure that this location is secure, because you don't want to compromise system security. The file is saved with the .pfx extension.
9. Tap or click Next, and then tap or click Finish. If the export process is successful, you'll see a message box confirming this. Tap or click OK to close the message box.

## Restoring Encryption Certificates

When you have a backup of a certificate, you can restore the certificate to any computer on the network—not just the original computer. The backup and restore process is, in fact, how you move certificates from one computer to another.

Follow these steps to restore a personal certificate:

1. Copy the Personal Information Exchange (.pfx) file onto removable media, such as a flash drive or a floppy disk, and then log on as the user to the computer where you want to use the personal certificate.

**NOTE** Log on to the target computer as the user whose certificate you're restoring. If you don't do this, the user won't be able to work with his encrypted data.

2. Access the Certificates snap-in for My User Account as described previously.
3. Expand Certificates—Current User, and then press and hold or right-click Personal. Tap or click All Tasks, and then tap or click Import. This starts the Certificate Import Wizard.
4. Tap or click Next, and then insert the removable media.
5. Tap or click Browse. In the Open dialog box, locate the personal certificate on the removable media. Be sure to select Personal Information Exchange as the file type. After you locate the file, select it, and then tap or click Open.
6. Tap or click Next. Type the password for the personal certificate, and then tap or click Next again.
7. The certificate should be placed in the Personal store by default. Accept the default by tapping or clicking Next. Tap or click Finish. If the import process is successful, you'll see a message box confirming this. Tap or click OK.



PART IV

# Windows Server 2012 Network Administration

**CHAPTER 14** Managing TCP/IP Networking **559**

**CHAPTER 15** Running DHCP Clients and Servers **573**

**CHAPTER 16** Optimizing DNS **615**



# Managing TCP/IP Networking

- Navigating Networking in Windows Server 2012 **559**
- Managing Networking in Windows 8 and Windows Server 2012 **562**
- Installing TCP/IP Networking **565**
- Configuring TCP/IP Networking **566**
- Managing Network Connections **571**

As an administrator, you enable networked computers to communicate by using the basic networking protocols built into Microsoft Windows Server 2012. The key protocol you use is TCP/IP. TCP/IP is a suite of protocols and services used for communicating over a network and is the primary protocol used for inter-network 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 4, “Automating Administrative Tasks, Policies, and Procedures.”

## Navigating Networking in Windows Server 2012

---

Windows Server 2012 has an extensive set of networking tools:

- **Network Explorer** Provides a central console for browsing computers and devices on the network
- **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

Before I describe how these networking tools are used, let's first look at the following Windows Server 2012 features on which these tools rely:

- **Network Discovery** A feature of Windows Server 2012 that controls the ability to see other computers and devices
- **Network Awareness** A feature of Windows Server 2012 that reports changes in network connectivity and configuration

**REAL WORLD** Computers running Windows Vista with SP1 or later, as well as later releases of Windows, support extensions to network awareness. These extensions allow a computer connected to one or more networks via two or more interfaces (regardless of whether they are wired or wireless) to select the route with the best performance for a particular data transfer. As part of selecting the best route, Windows chooses the best interface (either wired or wireless) for the transfer. This mechanism improves the selection of wireless over wired networks when both interfaces are present.

Network discovery settings for the computer you are working with determine the computers and devices you can browse or view in Windows Server 2012 networking tools. Discovery settings work in conjunction with a computer's Windows Firewall settings to block or allow the following:

- Discovery of network computers and devices
- Discovery of your computer by others

Network discovery settings are meant to provide the appropriate level of security for each of the various categories of networks to which a computer can connect. Three categories of networks are defined:

- **Domain network** Designates a network in which computers are connected to the corporate domain they are joined to
- **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

Because a computer saves settings separately for each category of network, different block and allow settings can be used for each network category. 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. Based on the network category, Windows Server 2012 automatically configures settings that turn discovery on or off. The On (Enabled) state means

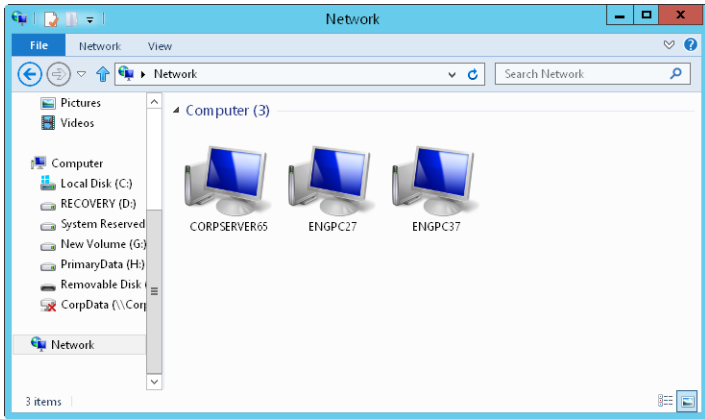
- The computer can discover other computers and devices on the network.
- Other computers on the network can discover the computer.

The Off (Disabled) state means

- The computer cannot discover other computers and devices on the network.
- Other computers on the network cannot discover the computer.

Typically, you will find that a network adapter is set as public before you join a computer to the domain. Network Explorer, shown in Figure 14-1, displays a list of

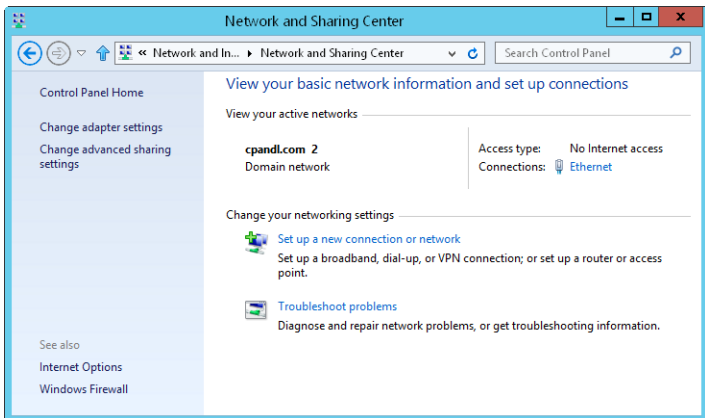
discovered computers and devices on the network. To access Network Explorer, tap or click File Explorer on the Start screen. In File Explorer, tap or click the location path selection button and then tap or click Network.



**FIGURE 14-1** Use Network Explorer to browse network resources.

The computers and devices listed in Network Explorer depend on the network discovery settings of the computer, the operating system, and whether the computer is a member of a domain. If discovery is blocked and a server running Windows Server 2012 is not a member of a domain, you'll see a note about this. When you tap or click the warning message and then select Turn On Network Discovery And File Sharing, you enable network discovery, file sharing, and printer sharing. This opens related Windows Firewall ports as well.

Network And Sharing Center, shown in Figure 14-2, provides the current network status, as well as an overview of the current network configuration. In Control Panel, you can access Network And Sharing Center by tapping or clicking View Network Status And Tasks under the Network And Internet heading.



**FIGURE 14-2** View and manage network settings with Network And Sharing Center.

Network And Sharing Center provides an overview of the network. The value below the network name shows the category of the current network as Domain Network, Private Network, or Public 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 tap or click the name of a network connection, you can display the related status dialog box.

Tapping or clicking Change Adapter Settings displays the Network Connections page, which you can use to manage network connections. Tapping or clicking Change Advanced Sharing Settings provides options for configuring the computer's sharing and discovery settings for each network profile. To manage a profile, expand the profile's view panel by tapping or clicking the Expand button (showing a down arrow), tap or click the setting you want to work with, and then tap or click Save Changes. To turn on or off network discovery, tap or click Turn On Network Discovery or Turn Off Network Discovery as appropriate, and then tap or click Save Changes.

From Network And Sharing Center, you can attempt to diagnose a networking problem. To do this, tap or click Troubleshoot Problems and then tap or 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.

## Managing Networking in Windows 8 and Windows Server 2012

---

In Group Policy, you'll find network management policies for both wired networks (IEEE 802.3) and wireless networks (IEEE 802.11) under Computer Configuration\Windows Settings\Security Settings. Only one wired policy and one wireless policy can be created and applied at a time. This means you can establish both a wired policy and a wireless policy for computers running Windows Vista and later releases of Windows. You also can create a wireless policy for computers running Windows XP.

If you press and hold or right-click the Wired Network (IEEE 802.3) node, you can create a policy for Windows Vista and later releases that specifies whether the Wire AutoConfig service is used to configure and connect these clients to 802.3 wired Ethernet networks. For Windows 7 and later releases of Windows, you have options for preventing the sharing of user credentials and for specifying whether to prohibit computers from making autoconnection attempts to the network for a specified amount of time.

If you press and hold or right-click the Wireless Network (IEEE 802.11) node, you can create separate policies for Windows XP computers and computers running later releases that enable WLAN autoconfiguration, define the specific networks that can be used, and set network permissions. For Windows 7 and later releases of Windows, you have options for preventing the sharing of user credentials, for specifying whether to prohibit computers from making autoconnection attempts



to the network for a specified amount of time, and for preventing the use of hosted networks.

Windows Vista with SP1 or later and later releases of Windows support several wired and wireless enhancements. These changes allow users to change their password when connecting to a wired or wireless network (as opposed to using the Winlogon change password feature), to correct a wrong password entered during sign on, and to reset an expired password—all as part of the network logon process.

Network security enhancements include the following:

- Secure Socket Tunneling Protocol (SSTP)
- Secure Remote Access (SRA)
- CryptoAPI Version 2 (CAPI2)
- Online Certificate Status Protocol (OCSP) extensions
- Port preservation for Teredo
- Remote Desktop Protocol (RDP) file signing

SSTP allows data transmission at the data link layer over a Hypertext Transfer Protocol over Secure Sockets Layer (HTTPS) connection. SRA enables secure access to remote networks over HTTPS. Together these technologies enable users to securely access a private network using an Internet connection. SSTP and SRA represent improvements over the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol/Internet Protocol Security (L2TP/IPsec) protocols because they use the standard TCP/IP ports for secure web traffic, and this allows them to traverse most firewalls as well as Network Address Translation (NAT) and web proxies.

SSTP uses HTTP over Secure Sockets Layer (SSL), which is also known as Transport Layer Security (TLS). HTTP over SSL (TCP port 443) is commonly used for protected communications with websites. Whenever users connect to a web address that begins with *https://*, they are using HTTP over SSL. Using HTTP over SSL solves many of the virtual private network (VPN) protocol connectivity problems. Because SSTP supports both IPv4 and IPv6, users can establish secure tunnels using either IP technology. Essentially, you get VPN technology that works everywhere, which should mean that you receive far fewer support calls.

CAPI2 extends support for public key infrastructure (PKI) and X.509 certificates and implements additional functionality for certificate path validation, certificate stores, and signature verification. One of the steps during certificate path validation is revocation checking, which involves verifying the certificate status to ensure that it has not been revoked by its issuer; this is where Online Certificate Status Protocol (OCSP) comes into the picture.

OCSP is used to check the revocation status of certificates. CAPI2 also supports independent OCSP signer chains and specifying additional OCSP download locations on a per-issuer basis. Independent OCSP signer chains modify the original OCSP implementation so that it can work with OCSP responses that are signed by trusted OCSP signers that are separate from the issuer of the certificate being validated. Additional OCSP download locations make it possible to specify OCSP download locations for issuing CA certificates as URLs that are added as a property to the CA certificate.

To ensure IPv4/IPv6 coexistence, Windows allows applications to use IPv6 on an IPv4 network, and it also supports related technologies such as port preservation for Teredo. Teredo is a User Datagram Protocol (UDP)–based tunneling technology that can traverse NATs. This feature allows Teredo communications between “port preserving” symmetric NATs and other types of NATs. A NAT is port preserving if it chooses to use the same external port number as the internal port number.

Current releases of Windows Server support TCP Chimney offloading. This feature enables the networking subsystem to offload the processing of a TCP/IP connection from a server’s processors to its network adapters, as long as the network adapters support TCP/IP offload processing. Both TCP/IPv4 connections and TCP/IPv6 connections can be offloaded. By default, TCP connections are offloaded on 10–gigabits per second (Gbps) network adapters but are not offloaded on 1-Gbps network adapters. To modify the related settings, you can use Netsh.

Network Diagnostic Framework (NDF) simplifies network troubleshooting by automating many common troubleshooting steps and solutions. When you run Windows Network Diagnostics, each diagnostic session generates a report with diagnostics results, and you can view this information in Action Center by tapping or clicking the Troubleshooting link and then tapping or clicking View History. On the Troubleshooting History page, each diagnostic session is listed by type and date run. To get detailed information, select the session you want to review and then tap or click View Details.

The diagnostic information shown in Action Center comes from an Event Trace Log (ETL) file created as part of diagnostics. If you press and hold or right-click a diagnostic session and then select Open File Location, you can see the files generated as part of diagnostics for the selected diagnostic session.

You can use the Netsh Trace context to perform comprehensive tracing as well as network packet capturing and filtering. You perform traces using predefined or custom scenarios and providers. Tracing scenarios are collections of providers. Providers are the actual components in the network protocol stack that you want to work with, such as TCP/IP, Windows Filtering Platform and Firewall, Wireless LAN Services, Winsock, or NDIS. Typically, you use Network Monitor (Netmon) to analyze trace data. If you collect trace data on a computer where Netmon isn’t installed, you can simply copy the trace file to a computer where Netmon is installed so that you can analyze the data.

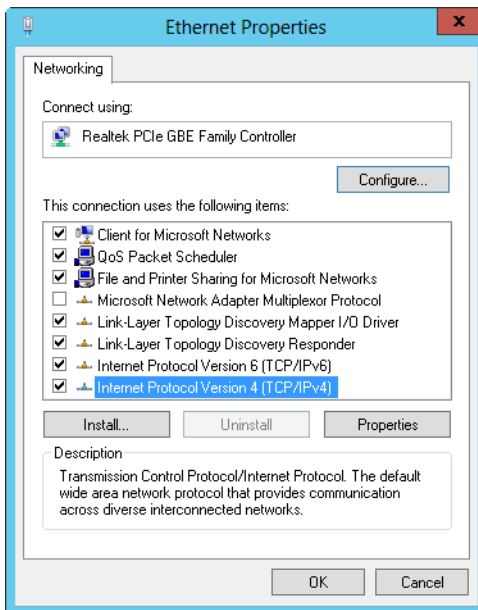
Windows Vista with SP1 or later and later releases of Windows use an RDP 6.1 or later compatible client. Here, RDP files can be digitally signed to prevent users from opening or running potentially dangerous RDP files from unknown sources. Administrators can sign RDP files using a signing tool provided by Microsoft. Three related settings can be configured through Group Policy or through the registry. These include a comma-separated list of certificate hashes that are trusted by the administrator (known as the *trusted publishers list*), an option to allow users to decide to accept untrusted publishers (enabled by default), and an option to allow users to accept unsigned files (enabled by default).

# Installing TCP/IP Networking

To install networking on a computer, you must install TCP/IP networking and a network adapter. Windows Server 2012 uses TCP/IP as the default wide area network (WAN) protocol. Normally, you install networking during Windows Server 2012 setup. You can also install TCP/IP networking through network connection properties.

To install TCP/IP after installing Windows Server 2012, log on to the computer using an account with administrator privileges and then follow these steps:

1. In Control Panel, access Network And Sharing Center by tapping or clicking View Network Status And Tasks under the Network And Internet heading.
2. In Network And Sharing Center, tap or click Change Adapter Settings.
3. In Network Connections, press and hold or right-click the connection you want to work with, and then tap or click Properties. This displays a Properties dialog box for the connection, shown in Figure 14-3.



**FIGURE 14-3** Install and configure TCP/IP protocols.

4. If Internet Protocol Version 6 (TCP/IPv6), Internet Protocol Version 4 (TCP/IPv4), or both aren't shown in the list of installed components, you need to install them. Tap or click Install. Tap or click Protocol, and then tap or click Add. In the Select Network Protocol dialog box, select the protocol to install, and then tap or click OK. If you are installing both TCP/IPv6 and TCP/IPv4, repeat this procedure for each protocol.

5. 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, and then tap or click OK.
6. As necessary, follow the instructions in the next section for configuring network connections for the computer.

## Configuring TCP/IP Networking

---

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 and is connected to a network, one network connection is created for each adapter. If no network connection is available, you should connect the computer to the network or create a different type of connection.

Computers use IP addresses to communicate over TCP/IP. Windows Server 2012 provides the following ways to configure IP addresses:

- **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 DHCPv4 and no DHCPv4 server is available, Windows Server 2012 assigns an alternate private IP address automatically. By default, the alternate IPv4 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 IPv4 address, which is particularly useful for laptop users.

## 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 normally 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 14-1 summarizes private network IPv4 addresses.

**TABLE 14-1** Private IPv4 Network Addressing

PRIVATE NETWORK ID	SUBNET MASK	NETWORK ADDRESS RANGE
10.0.0.0	255.0.0.0	10.0.0.0–10.255.255.255
172.16.0.0	255.240.0.0	172.16.0.0–172.31.255.255
192.168.0.0	255.255.0.0	192.168.0.0–192.168.255.255

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 to see whether an address is in use. Open a command prompt and type **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 local area network (LAN) 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, tap or click Change Adapter Settings. In Network Connections, press and hold or right-click the connection you want to work with, and then tap or click Properties.
2. Double-tap or 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:
  - Tap or click Use The Following IPv6 Address, and then type 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 2012 should insert a default value for the subnet prefix into the Subnet Prefix Length 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.
4. For an IPv4 address, do the following:
  - Tap or click Use The Following IP Address, and then type 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 2012 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. Type 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. Type a preferred address and an alternate DNS server address in the text boxes provided.
7. When you have finished, tap or 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:

1. In Network And Sharing Center, tap or 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 see a LAN connection for an installed adapter, check the driver for the adapter. It might be installed incorrectly. Press and hold or right-click the connection you want to work with, and then tap or click Properties.

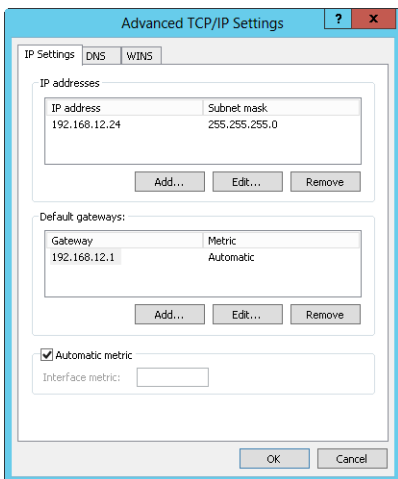
2. Double-tap or double-click Internet Protocol Version 6 (TCP/IPv6) or Internet Protocol Version 4 (TCP/IPv4) as appropriate for the type of IP address you are configuring.
3. Select Obtain An IPv6 Address Automatically or Obtain An IP Address Automatically as appropriate for the type of IP address you are configuring. You can select Obtain DNS Server Address Automatically, or you can select Use The Following DNS Server Addresses and then type a preferred and alternate DNS server address in the text boxes provided.
4. 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. Tap or click OK, tap or click Close, and then skip the remaining step.
5. To use a manual configuration, on the Alternate Configuration tab, select User Configured and then type 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, as shown earlier in Table 14-1, 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, tap or click OK, and then tap or click Close.

## Configuring Multiple Gateways

To provide fault tolerance in case of a router outage, you can choose to configure Windows Server 2012 computers so that they use multiple default gateways. When you assign multiple gateways, Windows Server 2012 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 2012 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, tap or click Change Adapter Settings. In Network Connections, press and hold or right-click the connection you want to work with, and then tap or click Properties.
2. Double-tap or 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. Tap or click Advanced to open the Advanced TCP/IP Settings dialog box, shown in Figure 14-4.



**FIGURE 14-4** Configure multiple IP addresses and gateways in 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.
  - a. Tap or click Add, and then type the gateway address in the Gateway text box.
  - b. By default, Windows Server 2012 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 tap or click Add.
  - c. Repeat steps a–c for each gateway you want to add.
5. Tap or click OK, and then tap or click Close.

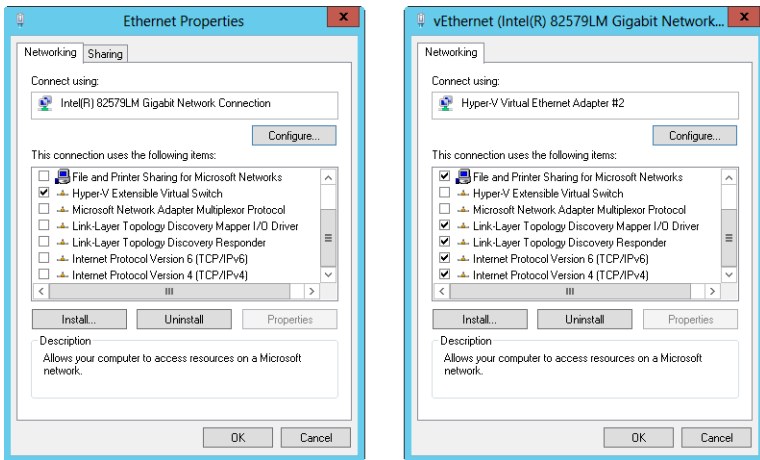
## Configuring Networking for Hyper-V

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 see 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 allows 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. Hyper-V binds the virtual network service to a physical network adapter only when an external virtual network is created.



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. As shown in Figure 14-5, the server has 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.



**FIGURE 14-5** Use switched virtual networking to ensure communications with hosted virtual machines.

## Managing Network Connections

Network connections make it possible for computers to access resources on the network and the Internet. One network connection is created automatically for each network adapter installed on a computer. This section examines techniques you can use to manage these connections.

### Checking the Status, Speed, and Activity for Network Connections

To check the status of a network connection, follow these steps:

1. In Network And Sharing Center, tap or click Change Adapter Settings. In Network Connections, press and hold or right-click the connection you want to work with, and then tap or click Status.
2. This displays the Status dialog box for the network connection. If the connection is disabled or the media is unplugged, you won't be able to access this dialog box. Enable the connection or connect the network cable to resolve the problem, and then try to display the Status dialog box again.

## Enabling and Disabling Network Connections

Network connections are created and connected automatically. If you want to disable a connection so that it cannot be used, follow these steps:

1. In Network And Sharing Center, tap or click Change Adapter Settings. In Network Connections, press and hold or right-click the connection, and then tap or click Disable to deactivate the connection and disable it.
2. If you want to enable the connection later, press and hold or right-click the connection in Network Connections and then tap or click Enable.

If you want to disconnect from a network, follow these steps:

1. In Network And Sharing Center, tap or click Change Adapter Settings. In Network Connections, press and hold or right-click the connection and then tap or click Disconnect. Typically, only remote access connections have a Disconnect option.
2. If you want to activate the connection later, press and hold or right-click the connection in Network Connections and then tap or click Connect.

## Renaming Network Connections

Windows Server 2012 initially assigns default names to network connections. In Network Connections, you can rename a connection at any time by pressing and holding or right-clicking the connection, tapping or clicking Rename, and then typing a new name. If a computer has multiple network connections, a descriptive name can help you and others better understand the uses of a particular connection.

# Running DHCP Clients and Servers

- Understanding DHCP **573**
- Installing a DHCP Server **579**
- Configuring DHCP Servers **584**
- Managing DHCP Scopes **592**
- Managing the Address Pool, Leases, and Reservations **605**
- Backing Up and Restoring the DHCP Database **610**

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. Once DHCP is installed, you rely on the DHCP server to supply the basic information necessary for TCP/IP networking, which 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, an 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 boot 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 then, 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 2012 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:

1. Install and configure two DHCP servers. The servers should be on the same physical network.
2. Create a DHCPv4 scope on one of the servers. Scopes are pools of IPv4 or IPv6 addresses you can assign to clients through leases.
3. 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 1, “Windows Server 2012 Administration Overview,” and Chapter 14, “Managing 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 as well as its network configuration parameters through DHCPv6.
- **DHCPv6 stateless mode** In DHCPv6 stateless mode, a client uses auto-configuration 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. IPv6 clients running Microsoft Windows XP or Windows Server 2003 do not include a DHCPv6 client and therefore ignore 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 typing the following command at an elevated 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 typing the following command at an elevated 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 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 non-link-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 Option-Data 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 Link-Address 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 Peer-Address 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, type the command `ipconfig /all` at the command prompt. If the IP address has been assigned automatically, you'll see an entry for Autoconfiguration IP Address. In the following example, the autoconfiguration IPv4 address is 169.254.98.59:

### Windows IP Configuration

```
Host Name .....: DELTA
Primary DNS Suffix .....: microsoft.com
Node Type .....: Hybrid
IP Routing Enabled.....: No
WINS Proxy Enabled.....: No
DNS Suffix Search List.....: microsoft.com
```

### 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.



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. They 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 class C IP addresses.

**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.

## Installing a DHCP Server

---

Dynamic IP addressing is available only if a DHCP server is installed on the network. 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. Only authorized DHCP servers can provide dynamic IP addresses to clients.

## Installing DHCP Components

On a server running Windows Server 2012, 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 they will service and are connected to. Be sure that the server has static IPv4 and IPv6 addresses.
2. In Server Manager, tap or click Manage and then tap or 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 tap or click Next.

3. On the Installation Type page, Role-Based Or Feature-Based Installation is selected by default. Tap or 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, tap or click Browse and then use the Browse For Virtual Hard Disks dialog box to locate the VHD. When you are ready to continue, tap or click Next.

**NOTE** Only servers running Windows Server 2012 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 see an additional dialog box. Tap or click Add Features to close the dialog box, and add the required features to the server installation. When you are ready to continue, tap or 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, type that alternate path in the box provided, and then tap or 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*.

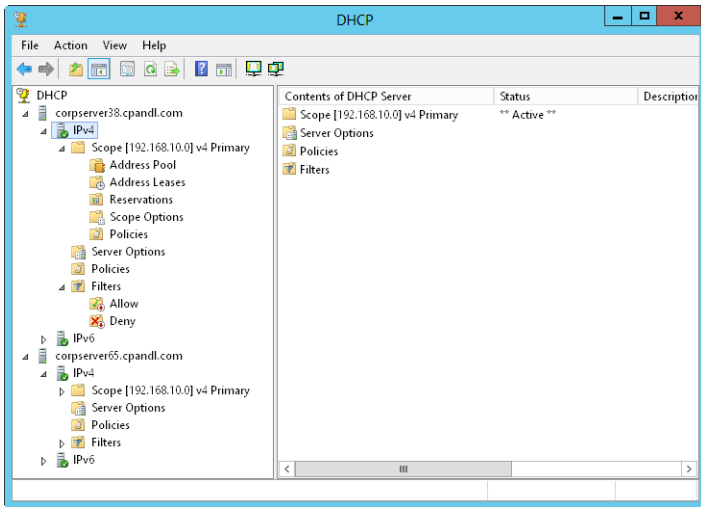
7. After you review the installation options and save them as necessary, tap or click Install to begin the installation process. The Installation Progress page tracks the progress of the installation. If you close the wizard, tap or click the Notifications icon in Server Manager and then tap or click 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. As stated in the Post-Deployment Configuration task panel, additional configuration is required for the DHCP server. Tap or click the Complete DHCP Configuration link. This starts the DHCP Post-Install Configuration Wizard.
10. The Description page states 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. Tap or 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, tap or click Commit to attempt to authorize the server using these credentials.

- If you want to use alternate credentials or if you are unable to authorize the server using your current credentials, select Use Alternate Credentials and then tap or click Specify. In the Windows Security dialog box, enter the user name and password for the authorized account and then tap or click OK. Tap or click Commit to attempt to authorize the server using these credentials.
  - If you want to authorize the DHCP server later, select Skip AD Authorization and then tap or click Commit. Keep in mind that in domains, only authorized DHCP servers can provide dynamic IP addresses to clients.
- 12.** When the wizard finishes the post-install configuration, review the installation details to ensure tasks were completed successfully and then tap or 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, tap or 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, press and hold or right-click the entry for the DHCP Server and then tap or 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 clients, including 003 Router, 006 DNS Servers, 015 DNS Domain Name, and 044 WINS/NBNS Servers. See “Setting Scope Options” later in the 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” later in the chapter.
    - Create and activate any DHCP scopes that the server will use, as discussed in “Creating and Managing Scopes” later in the chapter.

## 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, tap or click Tools and then tap or click DHCP to open the DHCP console. The main window for the DHCP console is shown in Figure 15-1. As you can see, the main window 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 see the scopes and options defined for the related IP version. The right pane shows the expanded view of the current selection.



**FIGURE 15-1** Use the DHCP console to create and manage DHCP server configurations.

Icons on the various nodes show the current status of the nodes. For server and IP nodes, you might see 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.

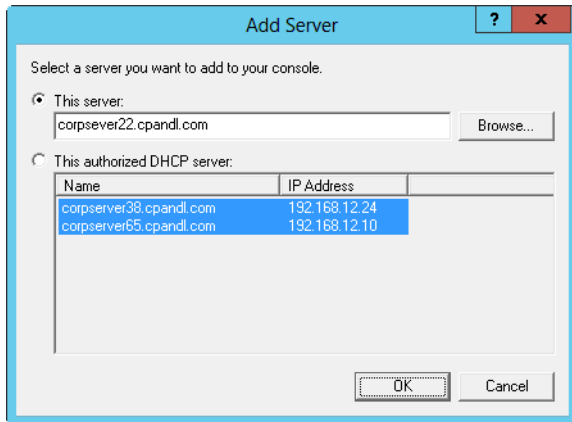
For scopes, you might see the following icons:

- 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 see entries for remote DHCP servers. You can connect to remote servers by following these steps:

1. Press and hold or right-click DHCP in the console tree, and then tap or click Add Server. This opens the dialog box shown in Figure 15-2.
2. Select This Server, and then type the IP address or computer name of the DHCP server you want to manage.
3. Tap or click OK. An entry for the DHCP server is added to the console tree.



**FIGURE 15-2** If your DHCP server isn't listed, you need to add it to the DHCP console by using the Add Server dialog box.

**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: press and hold or 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. Press and hold or right-click the server you want to manage in the DHCP console, point to All Tasks, and then tap or click Start, Stop, Pause, Resume, or Restart, as appropriate.

**NOTE** You also can use Server Manager to start and stop a DHCP server. Tap or 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, press and hold or right-click entry for the DHCP Server and then tap or 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 2012 requires authorization to prevent unauthorized DHCP servers from serving domain clients. This in turn ensures that network operations can run smoothly.

In the DHCP console, you authorize a DHCP server by pressing and holding or right-clicking the server entry in the tree view and then selecting Authorize. To remove the authorization, press and hold or right-click the server and then select Unauthorize.

# 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, press and hold or right-click the server you want to work with and then tap or click Add/Remove Bindings.
2. Select the IPv4 or IPv6 tab as appropriate for the type of binding you want to work with.
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 you are working with, ensure the server has a static address for that protocol.
4. Tap or 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 you want to work with, pressing and holding or right-clicking IPv4 or IPv6 as appropriate for the type of address you want to work with, and then tapping or clicking Display Statistics.

By default, these statistics are updated only when you start the DHCP console or when you select the server and then tap or click the Refresh button on the toolbar. If you monitor DHCP routinely, you might want these statistics to be updated automatically. To do that, follow these steps:

1. In the DHCP console, expand the node for the server you want to work with, press and hold or right-click IPv4 or IPv6 as appropriate for the type of address you want to work with, and then tap or click Properties.
2. On the General tab, select Automatically Update Statistics Every and enter an update interval in hours and minutes. Tap or click OK.

## DHCP Auditing and Troubleshooting

Windows Server 2012 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. Although you can enable and configure logging separately for IPv4 and IPv6, the two protocols use the same log files by default. The default location for DHCP logs is %SystemRoot%\System32\DHCP. In this directory, you'll find a different log file for each day of the week. The log file for Monday is named DhcpSrvLog-Mon.log, the log file for Tuesday is named DhcpSrvLog-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 necessarily clear a log file. Log data is cleared only when a log hasn't been written to in the last 24 hours. 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 you want to work with, press and hold or right-click IPv4 or IPv6 as appropriate for the type of address you want to work with, and then tap or click Properties.
2. On the General tab, select or clear the Enable DHCP Audit Logging check box and then tap or 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 you want to work with, press and hold or right-click IPv4 or IPv6 as appropriate for the type of address you want to work with, and then tap or click Properties.
2. Tap or click the Advanced tab. Audit Log File Path shows the current folder location for log files. Enter a new folder location, or tap or click Browse to select a new location.
3. Tap or click OK. Windows Server 2012 now needs to restart the DHCP Server service. When prompted to restart the service, tap or 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. Normally, 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\SYSTEM\CurrentControlSet\Services\DHCPServer\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 using the default setup for DHCP, current DHCP clients automatically update their own DNS records after receiving an IP address lease, and the DHCP server updates records for early clients after issuing a 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 2012. 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.

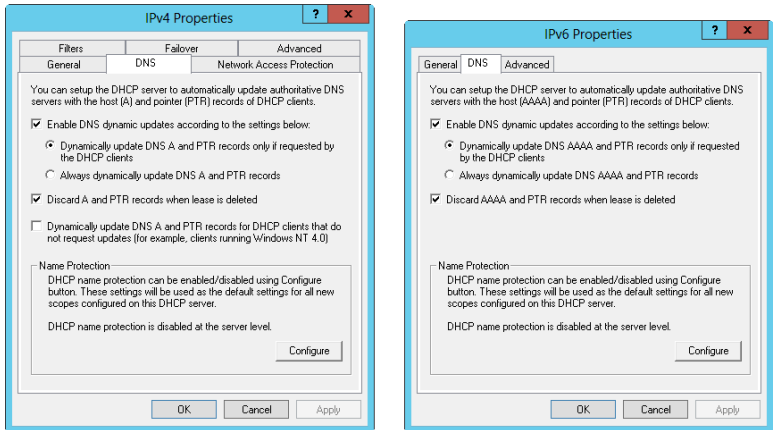
Name protection is designed to prevent name squatting. Name squatting occurs when a non-Windows-based computer registers a name in DNS that is already registered to a computer running a Windows operating system. By enabling name protection, you can prevent name squatting by non-Windows-based computers. 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 you want to work with, press and hold or right-click IPv4 or IPv6, and then tap or click Properties.
2. Tap or click the DNS tab. Figure 15-3 shows the default DNS integration settings for DHCP. Because these settings are configured by default, you usually don't need to modify the configuration.



**FIGURE 15-3** The DNS tab shows the default settings for DNS integration with DHCP.

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, tap or click Configure. In the Name Protection dialog box, select or clear Enable Name Protection and then tap or 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 you want to work with, and then expand IPv4 or IPv6.
2. Press and hold or right-click the scope you want to work with, and then tap or click Properties.
3. Tap or click the DNS tab. The options available are the same as those shown in Figure 15-3. 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. Tap or click Configure. In the Name Protection dialog box, select or clear Enable Name Protection and then tap or click OK.

# Integrating DHCP and NAP

Network Access Protection (NAP) is designed to protect the network from clients that do not have the appropriate security measures in place. The easiest way to enable NAP with DHCP is to set up the DHCP server as a Network Policy Server. To do this, you need to install the Network Policy And Access Services role, configure a compliant policy for NAP and DHCP integration on the server, and then enable NAP for DHCP. This process enables NAP for network computers that use DHCP, but it does not fully configure NAP for use.

You can create a NAP and DHCP integration policy by following these steps:

1. On the server that you want to act as the Network Policy Server, use the Add Roles And Features Wizard to install the Network Policy And Access Services role. You should install the Network Policy Server role service at a minimum.
2. In the Network Policy Server Console, available from the Tools menu in Server Manager, select the NPS (Local) node in the console tree, and then tap or click Configure NAP in the main pane. This starts the Configure NAP Wizard.
3. In the Network Connection Method list, choose Dynamic Host Configuration Protocol (DHCP) as the connection method you want to deploy on your network for NAP-capable clients. As shown in Figure 15-4, the policy name is set to NAP DHCP by default. Tap or click Next.

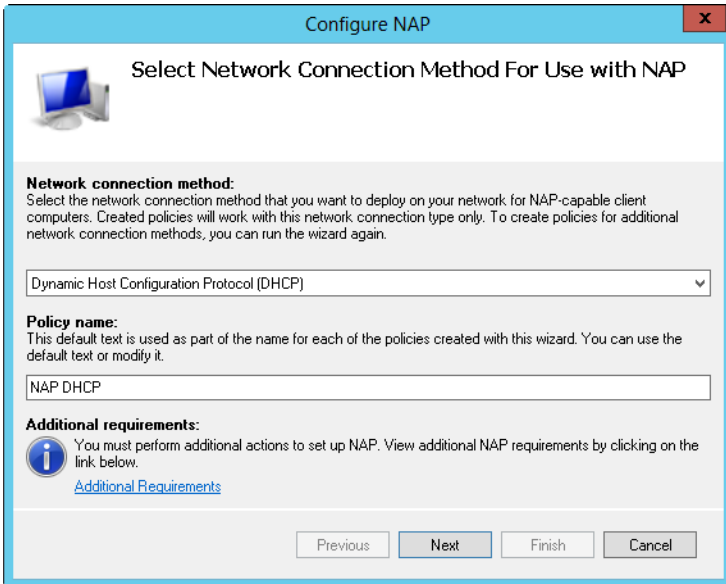


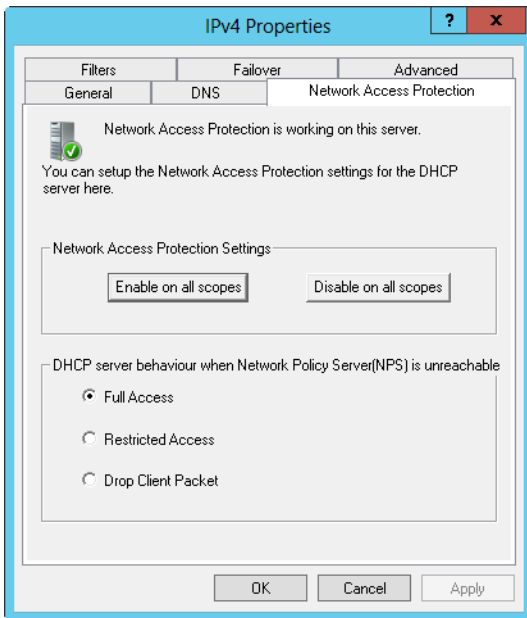
FIGURE 15-4 Configure Network Access Protection policy for the local DHCP server.

4. On the Specify NAP Enforcement Servers Running DHCP Server page, you need to identify all remote DHCP servers on your network by doing the following:
  - Tap or click Add. In the New RADIUS Client dialog box, type a friendly name for the remote server in the Friendly Name text box. Then type the DNS name of the remote DHCP server in the Address text box. Tap or click Verify to ensure that the DNS name is valid.
  - In the Shared Secret panel, select Generate, and then tap or click the Generate button to create a long shared-secret keyphrase. You need to enter this keyphrase in the NAP DHCP policy on all remote DHCP servers. Be sure to write down this keyphrase. Alternatively, copy the keyphrase to Notepad and save it in a file stored in a secure location. Tap or click OK.
5. Tap or click Next. On the Specify DHCP Scopes page, you can identify the DHCP scopes to which this policy should apply. If you do not specify any scopes, the policy applies to all NAP-enabled scopes on the selected DHCP servers. Tap or click Next twice to skip the Configure Machine Groups page.
6. On the Specify A NAP Remediation Server Group And URL page, select a Remediation Server, or tap or click New Group to define a remediation group and specify servers to handle remediation. Remediation servers store software updates for NAP clients that need them. In the text box provided, type a URL for a webpage that provides users with instructions on how to bring their computers into compliance with NAP health policy. Be sure that all DHCP clients can access this URL. Tap or click Next to continue.
7. On the Define NAP Health Policy page, use the options provided to determine how NAP health policy works. In most cases, the default settings work fine. With the default settings, NAP-ineligible clients are denied access to the network, and NAP-capable clients are checked for compliance and automatically remediated, which allows them to get needed software updates that you've made available. Tap or click Next, and then tap or click Finish.

You can modify NAP settings globally for each DHCP server or on a per-scope basis. To view or change the global NAP settings, follow these steps:

1. In the DHCP console, expand the node for the server you want to work with, press and hold or right-click IPv4, and then tap or click Properties.
2. On the Network Access Protection tab, shown in Figure 15-5, tap or click Enable On All Scopes or Disable On All Scopes to enable or disable NAP for all scopes on the server.

**NOTE** When the local DHCP server is also a Network Policy Server, the Network Policy Server should always be reachable. If you haven't configured the server as a Network Policy Server or the DHCP server is unable to contact the designated Network Policy Server, you'll see an error stating this on the Network Access Protection tab.



**FIGURE 15-5** The Network Access Protection tab controls the protection options for DHCP.

3. Choose one of the following options to specify how the DHCP server behaves if the Network Policy Server is unreachable, and then tap or click OK to save your settings:
  - **Full Access** Gives DHCP clients full (unrestricted) access to the network. This means clients can perform any permitted actions.
  - **Restricted Access** Gives DHCP clients restricted access to the network. This means clients can work only with the server to which they are connected.
  - **Drop Client Packet** Blocks client requests, and prevents the clients from accessing the network. This means clients have no access to resources on the network.

You can view and change the NAP settings for individual scopes by following these steps:

1. In the DHCP console, expand the node for the server you want to work with, and then expand IPv4.
2. Press and hold or right-click the scope you want to work with, and then tap or click Properties.
3. On the Network Access Protection tab, tap or click Enable For This Scope or Disable For This Scope to enable or disable NAP for this scope.
4. If you're enabling NAP and want to use a NAP profile other than the default, tap or click Use Custom Profile on the Network Access Protection tab, and then type the name of the profile, such as **Alternate NAP DHCP**.

5. Tap or click OK to save your settings.

## 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 is allowed to 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 you want to work with, press and hold or right-click IPv4, and then tap or 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 see 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. Once you create this script, you can restore the configuration by entering the following command at the command prompt:

```
netsh exec dhcpconfig.dmp
```

**TIP** You can also use this technique to set up another DHCP server with the same configuration. Simply 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, press and hold or right-click the DHCP server entry, tap or click Backup, use the dialog box provided to select the folder for the backup, and then tap or click OK. To restore the configuration, press and hold or right-click the DHCP server entry, tap or click Restore, use the dialog box provided to select the backup folder, and then tap or click OK. When prompted to confirm, tap or click Yes.

## Managing DHCP Scopes

---

After you install a DHCP server, you need to configure the scopes that the DHCP server will use. Scopes are pools of IP addresses you can lease to clients. As explained earlier in "Understanding Scopes," 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 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 and 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. 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.

#### 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 you want to work with, press and hold or right-click IPv4, and then tap or click New Superscope. This starts the New Superscope Wizard. Tap or click Next.
2. Type a name for the superscope, and then tap or click Next.
3. Select scopes to add to the superscope. Select individual scopes by tapping or clicking their entry in the Available Scopes list. Select multiple scopes by tapping or clicking entries while holding down Shift or Ctrl.
4. Tap or click Next, and then tap or click Finish.

#### Adding Scopes to a Superscope

You can add scopes to a superscope when you create it, or you can do it later. To add a scope to a superscope, follow these steps:

1. Press and hold or right-click the scope you want to add to a superscope, and then tap or click Add To Superscope.
2. In the Add Scope To A Superscope dialog box, select a superscope.
3. Tap or click OK. The scope is then added to the superscope.

## Removing Scopes from a Superscope

To remove a scope from a superscope, follow these steps:

1. Press and hold or right-click the scope you want to remove from a superscope, and then tap or click Remove From Superscope.
2. Confirm the action by tapping or 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, press and hold or right-click the superscope and then select Activate. To deactivate a superscope, press and hold or 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, press and hold or right-click the superscope and then select Delete. When prompted, tap or click Yes to confirm the action.

## 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 you want to work with, and then select and press and hold or right-click IPv4. If you want to add the new scope to a superscope automatically, select and then press and hold or right-click the superscope instead.
2. On the shortcut menu, tap or click New Scope. This starts the New Scope Wizard. Tap or click Next.
3. Type a name and description for the scope, and then tap or 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.

- When you enter an IP address range, the bit length and subnet mask are filled in for you automatically (as shown in Figure 15-6). Unless you use subnets, you should use the default values.

**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: 192 . 168 . 15 . 1

End IP address: 192 . 168 . 15 . 254

Configuration settings that propagate to DHCP Client

Length: 24

Subnet mask: 255 . 255 . 255 . 0

< Back   Next >   Cancel

**FIGURE 15-6** In the New Scope Wizard, enter the IP address range for the scope.

- Tap or click Next. If the IP address range you entered is on multiple networks, 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, tap or click Back, and then modify the IP address range you entered.
- 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, type a start address and an end address in the Exclusion Range's Start IP Address and End IP Address boxes and then tap or 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 tap or click Remove.
- Tap or click Next. Specify the duration of leases for the scope using the Day(s), Hour(s), and Minutes boxes. The default duration is eight days. Tap or click Next.



**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.
10. Tap or 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 tap or 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. Tap or click Next. As shown in Figure 15-7, 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.

**New Scope Wizard**

**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" value="corpserver65"/>	<input type="text" value="192.168.12.10"/>	<input type="button" value="Add"/>
<input type="button" value="Resolve"/>		<input type="button" value="Remove"/>
		<input type="button" value="Up"/>
		<input type="button" value="Down"/>

**FIGURE 15-7** Use the Domain Name And DNS Servers page to configure default DNS settings for DHCP clients.

13. In the IP Address box, enter the IP address of the primary DNS server, and then tap or 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. Tap or click Next.

**TIP** If you know the name of a server instead of its IP address, enter the name in the Server Name box, and then tap or click Resolve. The IP address is then entered in the IP Address box, if possible. Add the server by tapping or clicking Add.

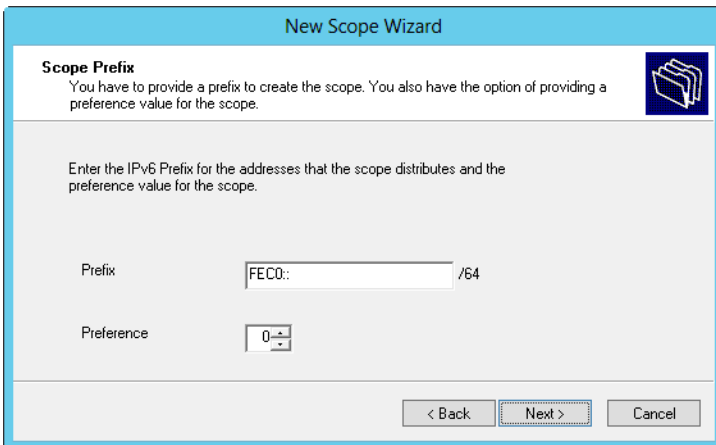
14. Configure default WINS settings for the DHCP clients. The techniques you use are the same as those previously described. Tap or click Next.
15. If you want to activate the scope, select Yes, I Want To Activate This Scope Now and then tap or click Next. Otherwise, select No, I Will Activate This Scope Later and then tap or click Next.
16. Tap or 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 you want to work with.
2. Select and then press and hold or right-click IPv6. On the shortcut menu, tap or click New Scope. This starts the New Scope Wizard. Tap or click Next.
3. Type a name and description for the scope, and then tap or click Next.
4. On the Scope Prefix page, shown in Figure 15-8, enter the 64-bit network prefix and then set a preference value. Tap or click Next.



The screenshot shows the 'New Scope Wizard' dialog box with the 'Scope Prefix' page selected. The title bar reads 'New Scope Wizard'. Below the title, the text says 'Scope Prefix' and 'You have to provide a prefix to create the scope. You also have the option of providing a preference value for the scope.' There is a folder icon in the top right corner. The main area contains the instruction 'Enter the IPv6 Prefix for the addresses that the scope distributes and the preference value for the scope.' Below this, there are two input fields: 'Prefix' with the value 'FE00::' and a '/64' suffix, and 'Preference' with a spinner box showing '0'. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

**FIGURE 15-8** In the New Scope Wizard, enter the network prefix and preference value.

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, type a start address and an end address in the Exclusion Range's Start IPv6 Address and End IPv6 Address boxes, and then tap or click Add. To exclude a single IPv6 address, use that address as the start IPv6 address and then tap or 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 tap or click Remove.

6. Tap or click Next. Dynamic IPv6 addresses can be temporary or nontemporary. A nontemporary address is similar to a reservation. On the Scope Lease page, shown in Figure 15-9, specify the duration of leases for nontemporary addresses 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. Tap or click Next.

**New Scope Wizard**

**Scope Lease**  
The lease duration specifies how long a client can use an IPv6 address obtained from this scope.

Lease durations should typically be equal to the average time the computer is connected to the same physical network.

Non Temporary Address(ANA)

Preferred LifeTime  
Days: 8 Hours: 0 Minutes: 0

Valid LifeTime  
Days: 12 Hours: 0 Minutes: 0

< Back Next > Cancel

**FIGURE 15-9** Specify the duration of nontemporary leases.

**NOTE** A lease lifetime 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 nontemporary leases is from 8 to 30 days.

7. If you want to activate the scope, select Yes under Activate Scope Now and then tap or click Finish. Otherwise, select No under Activate Scope Now, and then tap or click Finish.

## Creating Multicast Scopes

To create a multicast scope, follow these steps:

1. In the DHCP console, expand the node for the server you want to work with. Select and then press and hold or right-click IPv4. If you want to add the new scope to a superscope, select and then press and hold or right-click the superscope instead.
2. On the shortcut menu, tap or click New Multicast Scope. This starts the New Multicast Scope Wizard. Tap or click Next.
3. Enter a name and description for the scope, and then tap or click Next.
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 using Class D IP addresses. This means the valid IP address range is 224.0.0.0 to 239.255.255.255.
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. Tap or click Next. If you make a mistake, tap or 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:
  - 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 tap or 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 tap or click Remove.
8. Tap or click Next. Specify the duration of leases for the scope using the Day(s), Hour(s), and Minutes boxes. The default duration is 30 days. Tap or 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 tap or click Next. Otherwise, select No, and then tap or click Next.
10. Tap or click Finish to complete the process.

## Setting Scope Options

Scope options allow 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.

## VIEWING AND 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:

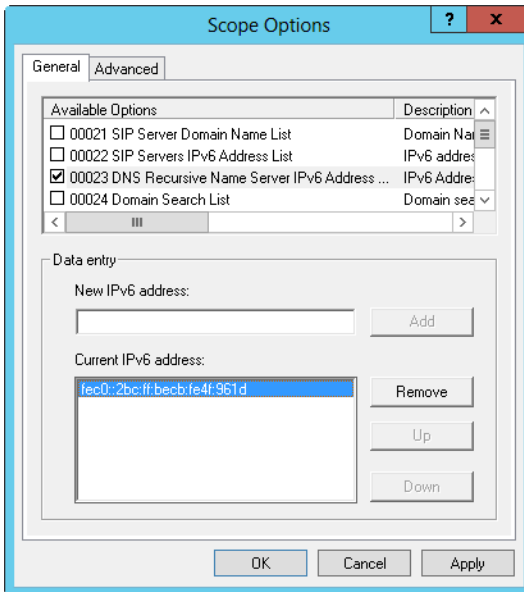
1. In the DHCP console, double-tap or double-click the server you want to work with, and then expand its IPv4 and IPv6 folders in the tree view.
2. To view current settings, select the Server Options node under IPv4 or IPv6, depending on the type of address you want to work with. Currently configured options are displayed in the right pane.
3. To assign new settings, press and hold or right-click Server Options, and then tap or click Configure Options. This opens 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. Tap or click OK to save your changes.

## VIEWING AND 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 you want to work with.
2. To view current settings, select Scope Options. Currently configured options are displayed in the right pane.
3. To assign new settings, press and hold or right-click Scope Options, and then tap or 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, as shown in Figure 15-10. Repeat this step to configure other options.



**FIGURE 15-10** Select the option you want to configure in the Scope Options dialog box, and then enter the required information on the Data Entry panel.

4. Tap or click OK.

## VIEWING AND 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 you want to work with.
2. Double-tap or double-click the Reservations folder for the scope.
3. To view current settings, tap or click the reservation you want to examine. Currently configured options are displayed in the right pane.
4. To assign new settings, press and hold or right-click the reservation, and then tap or click Configure Options. This opens 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-tap or double-click the server you want to work with, and then expand its IPv4 and IPv6 folders in the tree view. This should display the currently configured scopes for the server.
2. Press and hold or right-click the scope you want to modify, and then tap or 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, tap or 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 a normal folder icon.

You can activate an inactive scope by pressing and holding or right-clicking it in the DHCP console and then selecting Activate. You can deactivate an active scope by pressing and holding or right-clicking it in the DHCP console, and then selecting Deactivate.

**TIP** Deactivating turns off a scope but doesn't terminate current client leases. If you want to terminate leases, follow the instructions in "Releasing Addresses and Leases" 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. Press and hold or right-click the normal scope for IPv4 addresses that you want to modify, and then tap or click Properties.
2. On the Advanced tab, tap or click Both to support DHCP and BOOTP clients.
3. As necessary, set a lease duration for BOOTP clients, and then tap or click OK.

## 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, press and hold or right-click the scope you want to remove, and then tap or click Delete.
2. When prompted to confirm that you want to delete the scope, tap or 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 act as the primary server by splitting an existing scope or superscope. During the failover-scope



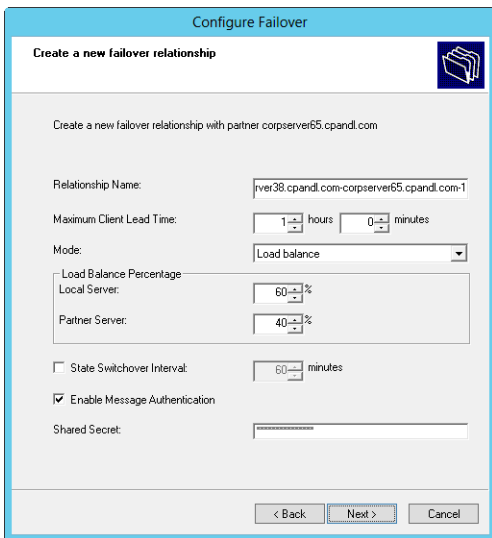
creation process, you need to specify the partner server with which you want to split the primary server's scope. This additional server acts as the secondary server for the scope. Because failover scopes are a server-side enhancement, no additional configuration is required for DHCP clients.

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 allows the fastest server to respond to and accept a DHCP OFFER 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 allows the primary DHCP server to respond to and accept the DHCP OFFER 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-tap or double-click the entry for the primary server, and then expand its IPv4 folder in the tree view.
2. The scope you want to work with must already be defined. Press and hold or right-click the scope or superscope that you want to configure for failover, and then tap or click Configure Failover. This starts the Configure Failover Wizard. Tap or click Next.
3. Next, you need to specify the partner server to use for failover. Tap or click Add Server. Use the options in the Add Server dialog box to select the secondary DHCP server for the failover scope, and then tap or click OK. Clear the Reuse Existing Failover Relationships check box and then tap or click Next to continue.
4. On the Create A New Failover Relationship page, shown in Figure 15-11, use the Mode list to set the failover mode as Load Balance or Hot Standby.



**FIGURE 15-11** Specify the percentage of the split.

5. 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 servers.
6. 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. By default, 5 percent of the IP addresses are reserved for the standby server.
7. Type 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. When you are ready to continue, tap or click Next.
8. Tap or click Finish. Review the summary of the failover scope configuration. If any errors were encountered, you might need to take corrective action. Tap or 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, press and hold or 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, press and hold or right-click the scope and then select Replicate Partnership.
- If you suspect the DHCP database the partners share is out of sync, press and hold or right-click the scope and then select Replicate Scope.
- If you no longer want to the scope to fail over, you can deconfigure failover by pressing and holding or right-clicking the scope and then selecting Deconfigure Failover.

You can't modify the failover settings once the partnership is established. However, you can deconfigure failover and then reconfigure failover.

## Managing the Address Pool, Leases, and Reservations

---

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, press and hold or 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 allows 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, press and hold or right-click the IPv4 node, and then tap or click Properties.
2. On the Filters tab, tap or 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. Tap or click OK to save your changes.

Before you can configure MAC address filtering, you must do 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, provided that 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, press and hold or right-click the IPv4 node, and then tap or click Properties.
2. On the Filters tab, you'll see the current filter configuration details. To use an allow list, select Enable Allow List. To use a deny list, select Enable Deny List.
3. Tap or click OK to save your changes.

**NOTE** As an alternative, you can simply press and hold or right-click the Allow or Deny node, under the Filters node, and then select Enable to enable allow or deny lists. If you press and hold or right-click the Allow or Deny node and then select Disable, you disable allow or deny lists.

Once you enable filtering, you define your filters using the MAC address for the client computer or device's network adapter. On a client computer, you can obtain the MAC address by typing the command **ipconfig /all** at the command prompt. The Physical Address entry shows the client's MAC address. You must type 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 normally 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-tap or double-click the IPv4 node, and then double-tap or double-click the Filters node.
2. Press and hold or right-click Allow or Deny as appropriate for the type of filter you are creating, and then tap or click New Filter.
3. Enter the MAC address to filter, and then enter a comment in the Description text box if you want to. Tap or click Add. Repeat this step to add other filters.
4. Tap or 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 you want to work with, and then press and hold or right-click the Address Pool folder or Exclusions folder. On the shortcut menu, tap or 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 tap or 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. Tap or click Close when you have finished.

To define an exclusion range for a scope with IPv6 addresses, follow these steps:

1. In the DHCP console, expand the scope you want to work with, and then press and hold or right-click the Exclusions folder. On the shortcut menu, tap or 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 tap or 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. Tap or 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, press and hold or right-click the exclusion, select Delete, and then tap or 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.

To reserve an IPv4 address for a client, follow these steps:

1. In the DHCP console, expand the scope you want to work with, and then press and hold or right-click the Reservations folder. On the shortcut menu, tap or click New Reservation.
2. In the Reservation Name text box, type 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.

4. The MAC Address box specifies the MAC address for the client computer's NIC. You can obtain the MAC address by typing the command **ipconfig /all** at the command prompt on the client computer. The Physical Address entry shows the client's MAC address. You must type 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. Tap or click Add to create the address reservation. Repeat this step to add other address reservations.
8. Tap or click Close when you have finished.

To reserve an IPv6 address for a client, follow these steps:

1. In the DHCP console, expand the scope you want to work with, and then press and hold or right-click the Reservations folder. On the shortcut menu, tap or click New Reservation.

2. In the Reservation text box, type 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 typing the command **ipconfig /all** at the command prompt on the client computer. The Physical Address entry shows the client's MAC address. You must type 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. Tap or click Add to create the address reservation. Repeat this step to add other address reservations.
8. Tap or click Close when you have finished.

## Releasing Addresses and Leases

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 typing the command **ipconfig /release** at the 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 typing the command **ipconfig /renew** at the command prompt.

## Modifying Reservation Properties

You can modify the properties of reservations by following these steps:

1. In the DHCP console, expand the scope you want to work with, and then tap or click the Reservations folder.
2. Press and hold or right-click a reservation, and then tap or 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 you want to work with, and then tap or click the Address Leases folder or Reservations folder, as appropriate.
2. Press and hold or right-click the lease or reservation you want to delete, and then tap or click Delete.
3. Confirm the deletion by tapping or 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 type the command **ipconfig /release** at the 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
- **J500000A.log** A reserved log file for the DHCP server
- **J500000B.log** A reserved log file for the DHCP server
- **J500000C.log** A reserved log file for the DHCP server
- **J500000D.log** A reserved log file for the DHCP server
- **J500000E.log** A reserved log file for the DHCP server
- **J500000F.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, press and hold or right-click the server you want to back up, and then tap or click Backup.
2. In the Browse For Folder dialog box, select the folder that will contain the backup DHCP database and then tap or click OK.

Registry keys that control the location and timing of DHCP backups, as well as other DHCP settings, are located under HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\DHCPServer\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. Tap or 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, press and hold or right-click the server you want to restore, and then tap or click Restore.
2. In the Browse For Folder dialog box, select the folder that contains the backup you want to restore and then tap or 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 DHCPServer 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 this chapter.

## 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. Once 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, press and hold or right-click the scope you want to work with, and then tap or click Reconcile.
2. In the Reconcile dialog box, tap or click Verify.
3. Inconsistencies are reported in the status window. Select the displayed addresses, and then tap or click Reconcile to repair inconsistencies.
4. If no inconsistencies are found, tap or click OK.

To reconcile all scopes on a server, follow these steps:

- 1.** In the DHCP console, expand the server entry, press and hold or right-click the IPv4 node, and then tap or click Reconcile All Scopes.
- 2.** In the Reconcile All Scopes dialog box, tap or click Verify.
- 3.** Inconsistencies are reported in the status window. Select the displayed addresses, and then tap or click Reconcile to repair inconsistencies.
- 4.** If no inconsistencies are found, tap or click OK.



# Optimizing DNS

- Understanding DNS **615**
- Configuring Name Resolution on DNS Clients **620**
- Installing DNS Servers **621**
- Managing DNS Servers **631**
- Managing DNS Records **636**
- Updating Zone Properties and the SOA Record **641**
- Managing DNS Server Configuration and Security **645**

This chapter discusses the techniques you use to set up and manage Domain Name System (DNS) on a network. DNS is a name-resolution service that resolves computer names to IP addresses. When you use DNS, a fully qualified host name—*omega.microsoft.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. Fully integrating DNS with these Windows networking features allows you to optimize DNS for Microsoft 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 *omega.microsoft.com*, *omega* represents the host name for an individual computer, *microsoft* 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. Normal 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 `jacob.hr.microsoft.com`. Here, *jacob* is the host name, *hr* is the child domain, and *microsoft.com* is the parent domain.

## Integrating Active Directory and DNS

As stated in Chapter 6, “Using Active Directory,” 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 (AD DS).

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 allows 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 allows 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 see 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 for Windows Server 2008 R2 and later has been enhanced to load 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 running Windows Server 2008 R2 and later 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. Using these addresses, clients can communicate with DNS servers anywhere on the network, even if the servers are on different subnets.

**NOTE** Configuring a DNS client is explained in Chapter 14, “Managing TCP/IP Networking.” Configuring a DNS server is explained in the next section of this chapter.

The DNS client built into computers running Microsoft Windows 7 and later, as well as Windows Server 2008 R2 or later, supports DNS traffic over Internet Protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6). By default, IPv6 configures the well-known site-local addresses of DNS servers at FEC0:0:0:FFF::1, FEC0:0:0:FFF::2, and FEC0:0:0:FFF::3. 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 DNS
```

DNS servers running Windows Server 2008 R2 or later 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 15, “Running DHCP Clients and Servers.” 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 7 or later, as well as Windows Server 2008 R2 or later, 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 Windows 8 and Windows Server 2012 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 has been increased to 410 milliseconds (msec) for the first retry and 410 msec for the second retry, making the total timeout value 820 msec instead of 300 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, as well as 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 2008 and later support 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 using its DNS suffix search lists, if any.
3. If none of those names can be resolved, the client requests resolution 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.

The GlobalNames zone provides single-label name resolution only when all authoritative DNS servers are running Windows Server 2008 R2 and later. However, other DNS servers that are not authoritative for any zone can be running other operating systems. 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. Open Network And Sharing Center, and then tap or click Change Adapter Settings.
2. In Network Connections, press and hold or right-click the connection you want to work with and then tap or click Properties.
3. Double-tap or 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 type primary and alternate DNS server addresses in the text boxes provided.
5. Tap or click Advanced to display the Advanced TCP/IP Settings dialog box. In this dialog box, tap or 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. Tap or click Add if you want to add a server IP address to the list. Tap or click Remove to remove a selected server address from the list. Tap or 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 button.
- **Append Primary And Connection Specific DNS Suffixes** Normally, 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. (Tap or click System And Security\System in Control Panel, tap or 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 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. Tap or click Add if you want to add a domain suffix to the list. Tap or click Remove to remove a selected domain suffix from the list. Tap or 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 button to change its position.
- **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 higher as well as Windows 2000 Server, Windows Server 2003, and later 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 2012 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:

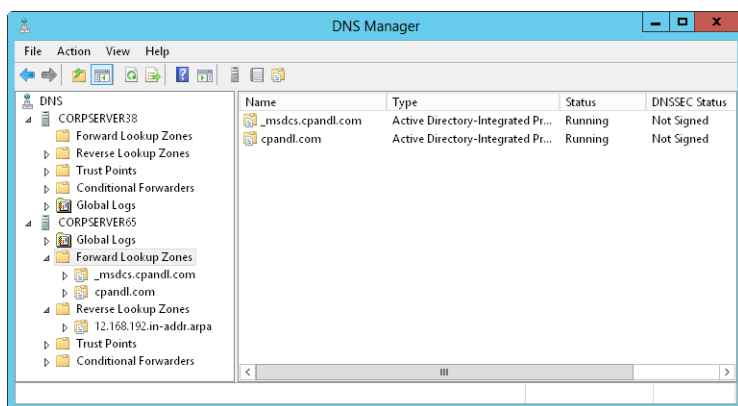
1. In Server Manager, tap or click **Manage** and then tap or click **Add Roles And Features**, or select **Add Roles And Features** in the Quick Start pane. This starts the Add Roles And Features Wizard. If the wizard displays the **Before You Begin** page, read the Welcome text and then tap or click **Next**.
2. On the **Installation Type** page, **Role-Based Or Feature-Based Installation** is selected by default. Tap or 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, tap or click **Browse** and then use the **Browse For Virtual Hard Disks** dialog box to locate the VHD. When you are ready to continue, tap or click **Next**.

**NOTE** Only servers running Windows Server 2012 and that have been added for management in Server Manager are listed.

4. On the **Server Roles** page, select **DNS Server**. If additional features are required to install a role, you'll see an additional dialog box. Tap or click **Add Features** to close the dialog box, and add the required features to the server installation. When you are ready to continue, tap or 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, type that alternate path in the box provided, and then tap or 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. Tap or click Install to begin the installation process. The Installation Progress page tracks the progress of the installation. If you close the wizard, tap or click the Notifications icon in Server Manager and then tap or click the link provided to reopen the wizard.
7. When Setup finishes installing the DNS Server role, the Installation Progress page will be updated to reflect this. Review the installation details to ensure the installation was successful.
8. From now on, the DNS Server service should start automatically each time you reboot the server. If it doesn't start, you need to start it manually. (See "Starting and Stopping a DNS Server" later in this chapter.)
9. After you install a DNS server, you use the DNS console to configure and manage DNS. In Server Manager, tap or click Tools and then tap or click DNS to open the DNS Manager console, shown in Figure 16-1.



**FIGURE 16-1** Use the DNS Manager console to manage DNS servers on the network.

10. If the server you want to configure isn't listed in the tree view, you need to connect to the server. Press and hold or right-click DNS in the tree view, and then tap or 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 tap or click OK.
  - If you're trying to connect to a remote server, select The Following Computer, type the server's name or IP address, and then tap or click OK.

11. An entry for the DNS server should be listed in the tree view pane of the DNS Manager console. Press and hold or right-click the server entry, and then tap or click Configure A DNS Server. This starts the Configure A DNS Server Wizard. Tap or click Next.
12. On the Select Configuration Action page, shown in Figure 16-2, select Configure Root Hints Only to specify that only the base DNS structures should be created at this time.

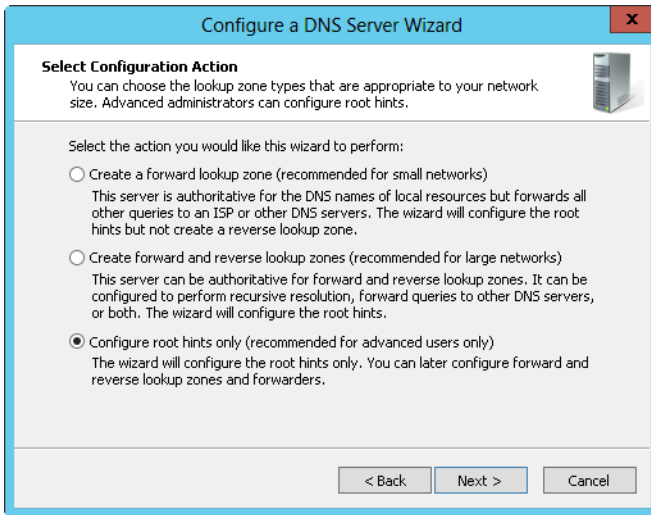


FIGURE 16-2 Configure the root hints only to install the base DNS structures.

13. Tap or click Next. The wizard searches for existing DNS structures and modifies them as necessary.
14. Tap or click Finish to complete the process.

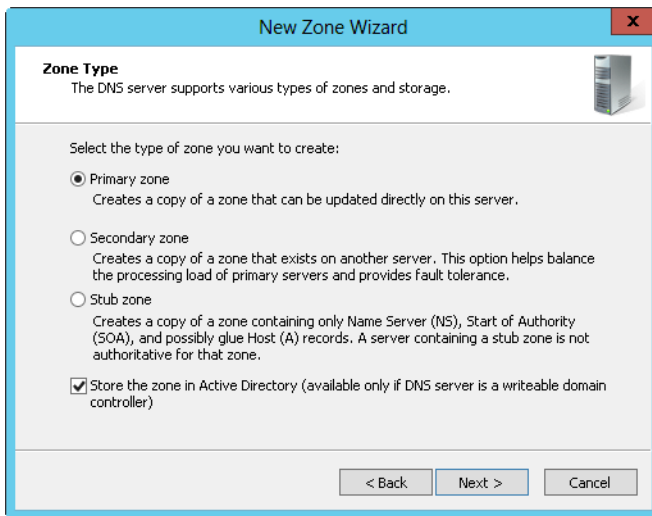
**REAL WORLD** If the wizard wasn't able to configure the root hints, you might need to configure the root hints manually or copy them from another server. However, a default set of root hints is included with DNS Server, and these root hints should be added automatically. To confirm, press and hold or 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.
2. An entry for the DNS server should be listed in the tree view pane of the DNS Manager console. Press and hold or right-click the server entry, and then tap or click New Zone. This starts the New Zone Wizard. Tap or click Next.
3. As Figure 16-3 shows, 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. Tap or click Next.



**FIGURE 16-3** In the New Zone Wizard, select the zone type.

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 (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 and supports compatibility with Windows 2000, not every domain controller is a DNS server as well (nor do you need to configure every domain controller as a DNS server).

5. Tap or click Next. Select Forward Lookup Zone, and then tap or click Next.
6. Type 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 microsoft.com domain, you would type **microsoft.com** as the zone name. Tap or 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 type a new file name. Tap or 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** Choosing this option disables dynamic updates in DNS. You should use this option only when the zone isn't integrated with Active Directory.
9. Tap or click Next, and then tap or 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" later in this chapter.

**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 workstations 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 adatum.com as your organization's DNS name and private IP addresses, as discussed in Chapter 14.

## 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 medium-size 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.

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. Press and hold or right-click the server entry, and then tap or click New Zone. This starts the New Zone Wizard. Tap or click Next.
3. For Zone Type, select Secondary Zone. Tap or 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 tap or click Next.
5. Type the full DNS name for the zone, and then tap or click Next.
6. Tap or click in the Master Servers list, type the 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. If you want to copy zone data from other servers in case the first server isn't available, repeat this step.
7. Tap or click Next, and then tap or 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 type 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. Press and hold or right-click the server entry, and then tap or click New Zone. This starts the New Zone Wizard. Tap or 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. Tap or click Next.
4. If you're configuring a reverse lookup zone for a secondary server, select Secondary Zone and then tap or 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 tap or click Next.

7. Choose whether you want to create a reverse lookup zone for IPv4 or IPv6 addresses, and then tap or click Next. Do one of the following:
  - If you are configuring a reverse lookup zone for IPv4, type the network ID for the reverse lookup zone. The values you enter set the default name for the reverse lookup zone. Tap or click Next.
  - 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, in this case you'd use 168.192.in-addr.arpa and allow the DNS Manager console to create the necessary subnet zones when needed.
  - If you are configuring a reverse lookup zone for IPv6, type the network prefix for the reverse lookup zone. The values you enter are used to automatically generate the related zone names. Depending on the prefix you enter, you can create up to eight zones. Tap or click Next.
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 type a new file name. Tap or 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. Tap or click Next, and then tap or 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. When all the DNS servers for your zones are running Windows Server 2008 or later releases, deploying a GlobalNames zone creates static, global records with single-label names, without relying on WINS. This allows users to access hosts 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.
2. Press and hold or right-click the Forward Lookup Zones node, and then tap or click New Zone. In the New Zone Wizard, tap or 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 tap or click Next. On the Zone Name page, enter **GlobalNames** as the zone name. Tap or click Next twice, and then tap or click Finish.
3. On every authoritative DNS server in the forest now and in the future, you need to type the following at an elevated command prompt: **dnscmd ServerName /enableglobalnamesupport 1**, where *ServerName* is the name of the DNS server that hosts the GlobalNames zone. To specify the local computer, use a period (.) instead of the server name, such as **dnscmd . /enableglobalnamesupport 1**.
4. For each server that you want users to be able to access using a single-label name, add an alias (CNAME) record to the GlobalNames zone. In the DNS Manager console, press and hold or right-click the GlobalNames node, select New Alias (CNAME), and then use the dialog box provided to create the new resource record.

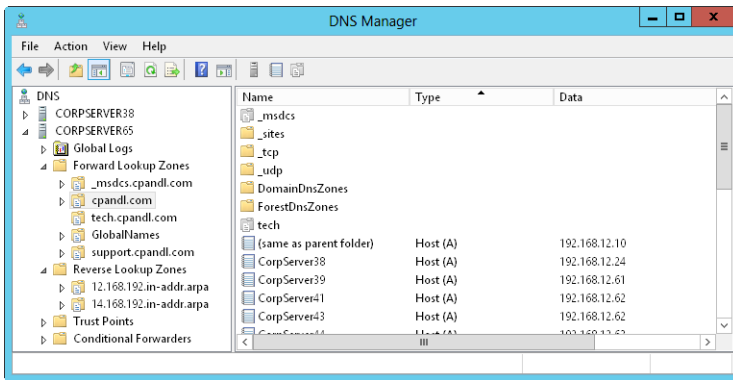
**NOTE** An authoritative DNS server tries to resolve queries in the following order: using local zone data, using the GlobalNames zone, using DNS suffixes, using WINS. For dynamic updates, an authoritative DNS server checks the GlobalNames zone before checking the local zone data.

**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.

## Managing DNS Servers

The DNS Manager console is the tool you use to manage local and remote DNS servers. As shown in Figure 16-4, the DNS Manager console's main window is divided into two panes. The left pane allows 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-tap or 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.
- Press and hold or right-click an entry to display a context menu.



**FIGURE 16-4** Manage local and remote DNS servers using the DNS Manager console.

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. Press and hold or right-click DNS in the console tree, and then tap or 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 type the IP address or fully qualified host name of the remote computer you want to connect to.
3. Tap or click OK. Windows Server 2012 attempts to contact the server. If it does, 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 tapping or clicking Yes when prompted.

In the DNS Manager console, you can delete a server by selecting its entry and then pressing Delete. When prompted, tap or 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. Press and hold or right-click the server you want to manage in the DNS Manager console, point to All Tasks, and then tap or click Start, Stop, Pause, Resume, or Restart as appropriate.

**NOTE** In Server Manager, under the DNS Server node, expand the DNS node and then press and hold or right-click the server you want to work with. On the shortcut menu, select Start Service, Stop Service, Pause Service, Resume Service, or Restart Service as appropriate.

## Using DNSSEC and Signing Zones

Windows 7 or later versions, as well as Windows Server 2008 R2 or later, 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. Normally, 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 2012 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 (NSEC3 and RSA/SHA-2 standards). Previously, you could not sign records with NSEC3 and RSA/SHA-2.

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 2008 R2 or later DNS 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 2008 R2 or later 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 2008 R2 or later 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 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, press and hold or 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 tap or click Next.
2. On the Signing Options page, select Customize Zone Signing Parameters and then tap or 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, tap or click Next twice.
4. On the Key Signing Key page, configure a KSK by tapping or clicking Add, accepting or changing the default values for key properties and rollover, and then tapping or clicking OK. When you are ready to continue, tap or click Next twice.
5. On the Zone Signing Key page, configure a ZSK by tapping or clicking Add, accepting or changing the default values for key properties and rollover, and then tapping or clicking OK. When you are ready to continue, tap or click Next five times.
6. 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, press and hold or 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 tap or click Next.
2. On the Signing Options page, select Sign The Zone With Parameters Of An Existing Zone. Type the name of an existing signed zone, such as **cpandl.com**. Tap or 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. Tap or click Next twice.
4. After the wizard signs the zone, click Finish.

## Creating Child Domains Within Zones

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 you want to work with.



2. Press and hold or right-click the parent domain entry, and then tap or click New Domain.
3. Enter the name of the new domain, and then tap or click OK. For hr.microsoft.com, you would enter **hr**. For mis.microsoft.com, you would enter **mis**.

## Creating Child Domains in Separate Zones

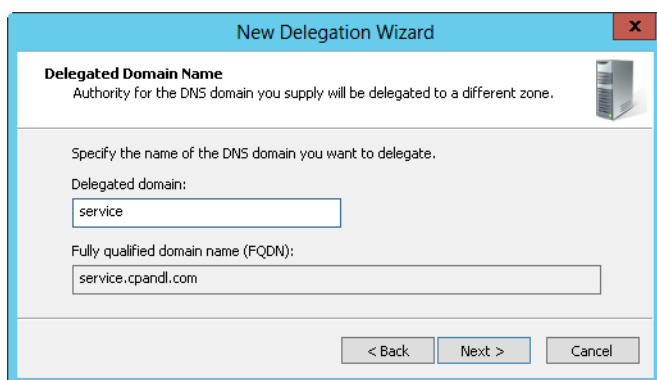
As your organization grows, you might want to organize the DNS namespace into separate zones. At your corporate headquarters, you could have a zone for the parent domain microsoft.com. At branch offices, you could have zones for each office, such as memphis.microsoft.com, newyork.microsoft.com, and la.microsoft.com.

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 allows 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 you want to work with.
2. Press and hold or right-click the parent domain entry, and then tap or click New Delegation. This starts the New Delegation Wizard. Tap or click Next.
3. As shown in Figure 16-5, type the name of the delegated domain, such as **service**, and then tap or click Next. The name you enter updates the value in the Fully Qualified Domain Name text box.



**FIGURE 16-5** Entering the name of the delegated domain sets the fully qualified domain name (FQDN).

4. Tap or 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.memphis.adatum.com**, and then tap or 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, tap or click OK to close the New Name Server Record dialog box.
7. Tap or click Next, and then tap or 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, press and hold or right-click the domain or subnet entry.
2. On the shortcut menu, tap or click Delete, and then confirm the action by tapping or clicking Yes.
3. If the domain or subnet is integrated with Active Directory, you'll see a warning prompt. Confirm that you want to delete the domain or subnet from Active Directory by tapping or 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, 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 allows email messages to be delivered to the correct mail servers in the domain.
- **NS (name server)** Specifies a name server for the domain, which allows 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.

## 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.

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 you want to work with.
2. Press and hold or right-click the domain you want to update, and then tap or click **New Host (A Or AAAA)**. This opens the dialog box shown in Figure 16-6.

**FIGURE 16-6** Create address records and pointer records simultaneously with the New Host dialog box.

3. Type the single-part computer name, such as **servicespc85**, 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” earlier in this chapter. The Allow Any Authenticated User option is available only when a DNS server is configured on a domain controller.

5. Tap or click Add Host, and then tap or click OK. Repeat these steps as necessary to add other hosts.
6. Tap or 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 you want to work with.
2. Press and hold or right-click the subnet you want to update, and then tap or click New Pointer (PTR).
3. Type the host IP address, such as **192.168.1.95**, and then type the host name, such as **servicespc54**. Tap or click OK.

### Adding DNS Aliases with CNAME

You specify host aliases using CNAME records. Aliases allow a single host computer to appear to be multiple host computers. For example, the host `gamma.microsoft.com` can be made to appear as `www.microsoft.com` and `ftp.microsoft.com`.

To create a CNAME record, follow these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server you want to work with.
2. Press and hold or right-click the domain you want to update, and then tap or click New Alias (CNAME).
3. In the Alias Name text box, type 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, type the full host name of the computer for which the alias is to be used.
5. Tap or 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.

You create an MX record by following these steps:

1. In the DNS Manager console, expand the Forward Lookup Zones folder for the server you want to work with.
2. Press and hold or right-click the domain you want to update, and then tap or 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 **cpandl.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.cpandl.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.

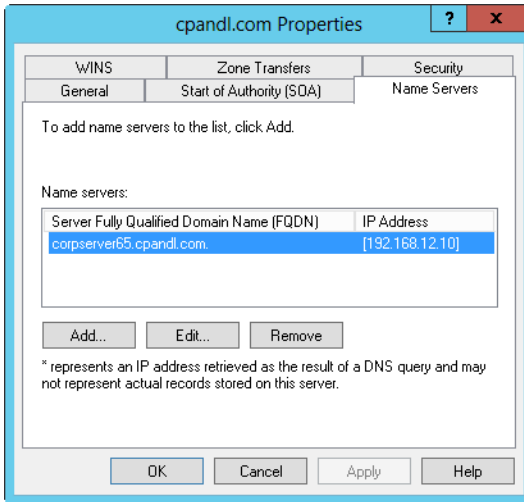
4. Tap or click OK.

## 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 you want to work with.
2. Display the DNS records for the domain by selecting the domain folder in the tree view.
3. Press and hold or right-click an existing NS record in the view pane, and then tap or click Properties. This opens the Properties dialog box for the domain with the Name Servers tab selected, as shown in Figure 16-7.



**FIGURE 16-7** Configure name servers for the domain through the domain's Properties dialog box.

4. Tap or click Add. This displays the New Name Server Record dialog box.
5. In the Server Fully Qualified Domain Name text box, type the name of a DNS server for the child domain, such as **corpserver01.cpandl.com** and then tap or 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 using the Up and Down buttons. When you are ready to continue, tap or click OK to close the New Name Server Record dialog box.
7. Tap or click OK to save your changes.

## Viewing and Updating DNS Records

To view or update DNS records, follow these steps:

1. Double-tap or double-click the zone you want to work with. Records for the zone should be displayed in the right pane.
2. Double-tap or 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 tap or 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:

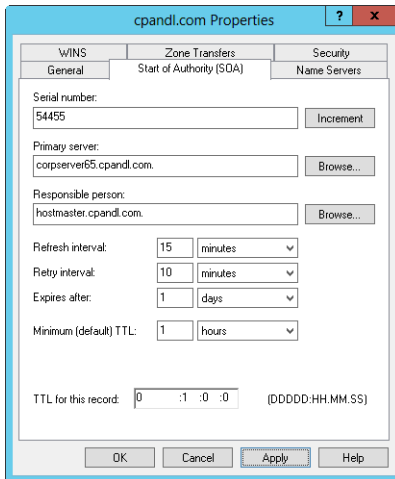
- Press and hold or right-click the zone you want to update, and then tap or click Properties.
- Select the zone, and then tap or click Properties on the Action menu.

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, press and hold or right-click the zone you want to update and then tap or click Properties.
2. Tap or click the Start Of Authority (SOA) tab, and then update the text boxes shown in Figure 16-8.



The screenshot shows the 'cpandl.com Properties' dialog box with the 'Start of Authority (SOA)' tab selected. The dialog has three tabs: 'WINS General', 'Zone Transfers Start of Authority (SOA)', and 'Security Name Servers'. The 'Start of Authority (SOA)' tab contains the following fields and controls:

- Serial number: 54455 (with an 'Increment' button)
- Primary server: corpserver65.cpandl.com (with a 'Browse...' button)
- Responsible person: hostmaster.cpandl.com (with a 'Browse...' button)
- Refresh interval: 15 minutes (with a dropdown menu)
- Retry interval: 10 minutes (with a dropdown menu)
- Expires after: 1 days (with a dropdown menu)
- Minimum (default) TTL: 1 hours (with a dropdown menu)
- TTL for this record: 0 :1 :0 :0 (DDDD:HH.MM.SS)

At the bottom of the dialog are buttons for 'OK', 'Cancel', 'Apply', and 'Help'.

**FIGURE 16-8** In the zone's Properties dialog box, set general properties for the zone and update the SOA record.

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, 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. If the interval is set to 60 minutes, NS record changes might not be propagated to a secondary server for up to an hour. You reduce network traffic by increasing this value.
- **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 allows 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.



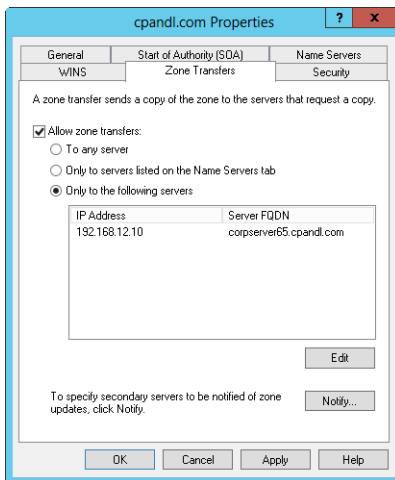
## Allowing and Restricting Zone Transfers

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 2012 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 allows 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, press and hold or right-click the domain or subnet you want to update and then tap or click Properties.
2. Tap or click the Zone Transfers tab, as shown in Figure 16-9.



**FIGURE 16-9** Use the Zone Transfers tab to allow zone transfers to any server or to designated servers.

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 tap or click Edit as appropriate to display the Allow Zone Transfers dialog box. Tap or click in the IP Address list, type the IP address of the secondary server for the zone, and then press Enter. Windows then attempts to validate the server. If

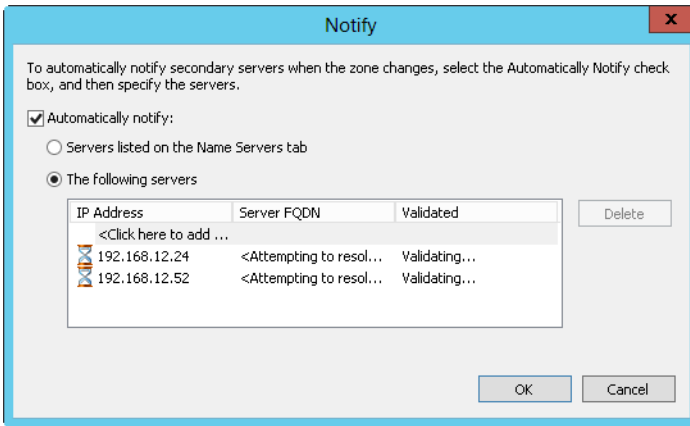
an error occurs, make sure the server is connected to the network and that you've entered the correct IP address. If you want to copy zone data from other servers in case the first server isn't available, you can add IP addresses for other servers as well. Tap or click OK.

5. Tap or 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, press and hold or right-click the domain or subnet you want to update and then tap or click Properties.
2. On the Zone Transfers tab, tap or click Notify. This displays the dialog box shown in Figure 16-10.



**FIGURE 16-10** In the Notify dialog box, notify all secondaries listed on the Name Servers tab or specific servers that you designate.

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. Tap or click in the IP Address list, type 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. Tap or 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, press and hold or right-click the domain or subnet you want to update and then tap or click Properties.
2. Under Type on the General tab, tap or 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. Tap or click OK twice.

## Enabling and Disabling Dynamic Updates

Dynamic updates allow 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, press and hold or right-click the domain or subnet you want to update and then tap or 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. Tap or click OK.

**NOTE** DNS integration settings must also be configured for DHCP. See "Integrating DHCP and DNS" in Chapter 15.

## 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 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, press and hold or right-click the server you want to configure and then tap or 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. Tap or click OK.

## Controlling Access to DNS Servers Outside the Organization

Restricting access to zone information allows 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 it. 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, press and hold or right-click the server you want to configure and then tap or click Properties.

2. Tap or click the Advanced tab. To configure the server as a nonforwarder, ensure that the Disable Recursion check box is cleared, tap or 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, tap or click Edit. This displays the Edit Forwarders dialog box.
4. Tap or 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. Tap or 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 allows 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 press and hold or right-click the Conditional Forwarders folder for the server you want to work with. Tap or 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 **adatum.com**.
3. Tap or 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).
5. 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. Tap or click OK.
  6. 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:

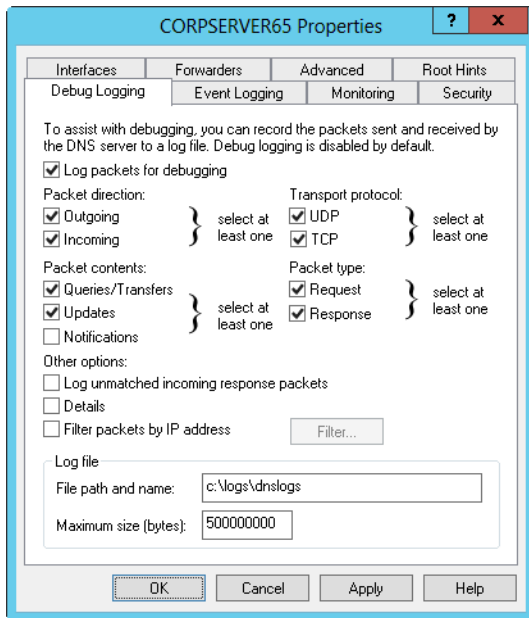
1. In the DNS Manager console, press and hold or right-click the server you want to configure and then tap or click Properties.
2. Use the options on the Event Logging tab to configure DNS logging. To disable logging altogether, choose No Events.
3. Tap or click OK.

## Using Debug Logging to Track DNS Activity

You normally 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, press and hold or right-click the server you want to configure and then tap or click Properties.
2. On the Debug Logging tab, shown in Figure 16-11, select the Log Packets For Debugging check box and then select the check boxes for the events you want to track temporarily.



**FIGURE 16-11** Use the Debug Logging tab to select the events you want to log.

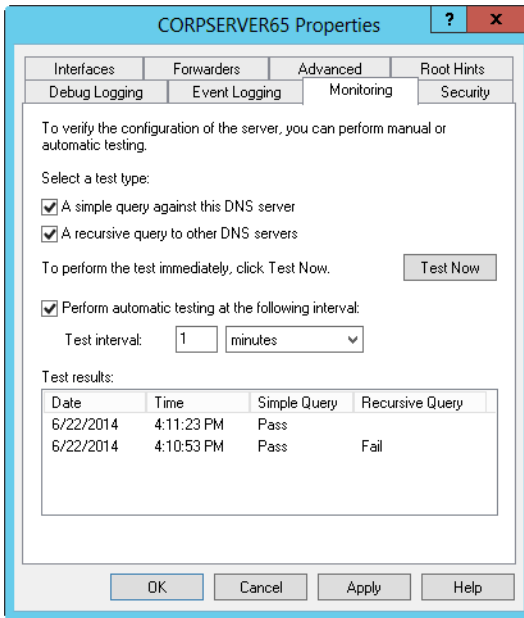
3. In the File Path And Name text box, enter the name of the log file, such as **dns.logs**. Logs are stored in the %SystemRoot%\System32\Dns directory by default.
4. Tap or click OK. When finished debugging, turn off logging by clearing the Log Packets For Debugging check box.

## Monitoring a DNS Server

Windows Server 2012 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, press and hold or right-click the server you want to configure and then tap or click Properties.
2. Tap or click the Monitoring tab, shown in Figure 16-12. 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.



**FIGURE 16-12** Configure a DNS server for manual or automatic monitoring on the Monitoring tab.

3. You can perform a manual test by tapping or 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 see 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 normally indicate a DNS resolution problem.

**NOTE** If all recursive query tests fail, the advanced server option Disable Recursion might be selected. Tap or 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.



# Index

## Symbols & Numbers

- \$ (special share symbol), 474–475
- 32-bit processes, 92–93
- 64-bit systems, 6, 43–44

## A

- access controls, 297–300. *See also* permissions
- Account Lockout Policy, 194–195, 321–322, 325–326
- account policies
  - Default Domain Policy GPO, 147
  - lockout, 194–195, 321–322, 325–326
  - security templates for, 193–195
  - setting, 321–322
- accounts
  - Administrator, 307–308, 317
  - computer. *See* computer accounts
  - default and predefined, 306–308
  - deleting, 373
  - disk quotas with, 504
  - domain. *See* domain accounts
  - expiration dates, changing, 375, 378
  - exporting, 372
  - group. *See* group accounts
  - group memberships of, managing, 339–340
  - Guest account, 308
  - importing, 372
  - locked out accounts, 373–374
  - lockout policies, 194–195, 321–322, 325–326
  - logon rights for, 312–313
  - managed service accounts, 341–345
  - naming policies, 319–320
  - passwords for. *See* passwords
  - security options, 360–361
  - tools for creating, 318
  - updating, 368–369
  - user. *See* user accounts
  - virtual accounts, 346
- ACEs (access control entries), 297
- ACLs (access control lists), 252
- ACPI (Advanced Configuration and Power Interface), 8–11
- Act As Part Of The Operating System privilege, 310
- Action Center, 542, 564
- activation, 45, 73–74, 221
- active cooling mode, 8–9
- Active Directory
  - 2008 R2 features, 220–221
  - activation feature, 221
  - Administration Tool, 252. *See also* LDAP (Lightweight Directory Access Protocol)
  - Administrative Center. *See* Active Directory Administrative Center
  - Adprep command, 250
  - ADSI use for communications, 242
  - auditing objects of, 502–503
  - authentication mechanism assurance, 220
- CAs (certificate authorities), 14–15, 33, 75
- Certificate Services, 14–15, 33
- claims-based policy controls, 221
- client computer requirements for, 230
- command-line tool summary, 250–251
- Configuration Wizard, 218–219
- data stores, 14
- deferred index creation, 221
- DHCP integration options, 19–20. *See also* DHCP (Dynamic Host Configuration Protocol)
- diagnosing problems, tool for, 252, 289–290
- Directory Services Restore Mode, 16–17
- distribution of updated data. *See* replication
- DNS integration with, 18–20, 217–218, 616–617, 625–626, 628
- Domain Services, 15, 33
- Domains And Trusts snap-in, 226–227
- domains, structural overview of, 223–224. *See also* domains
- Ds commands, list of, 251

### Active Directory, *continued*

- Enhanced Fine-Grained Password Policy, 222
- exporting accounts, 372
- features, table of key, 221–223
- Federation Services, 15, 33
- forests, domain, 225–227, 317
- FSMO (flexible single-master operations) roles, 71, 273, 276
- full integration with DNS, 19
- functional levels of domains, 225–227, 231–237
- Group Managed Service Accounts, 222
- Group Policy, relation to, 138
- importing accounts, 372
- installing, 218, 270–272
- interoperability with other directory services, 242
- Kerberos with. *See* Kerberos
- LDAP for. *See* LDAP (Lightweight Directory Access Protocol)
- Lightweight Directory Services, 15, 33
- maintenance while stopped, 16–17
- managed accounts, 220
- mechanism of, overview, 237–238
- multiple domain controllers with, 14
- Ntdsutl command, 251
- offline domain join feature, 220
- off-premises domain join, 222
- partial integration with DNS, 18–19
- permissions, setting, 380–381
- ports used by, 292–293
- PowerShell, module for, 220–221, 259
- publishing shares, 469
- read-only domain controllers, 16
- recovery of deleted objects, 245–248
- Recycle Bin, 220, 222, 244–248
- Relative ID soft ceilings and warnings, 223
- Restartable Active Directory Domain Services, 16–17
- restoring, 547
- Rights Management Services, 15, 33
- RODC with DNS issues, 19
- schemas of. *See* schemas
- searching for directory objects, 254–256, 258
- searching for users and groups, 350–351
- Server Manager integration, 223
- servers running. *See* domain controllers sites, 138
- structures provided by, 223

- support tools, table of, 251–252
- tools for, snap-ins, 249–250
- trees, domain, 225–227
- troubleshooting, 292–294
- Users And Computers tool. *See* Active Directory Users And Computers
- virtual domain controller cloning, 223
- Web Services, 221, 259
- Windows Firewall issues with, 250

### Active Directory Administrative Center

- adding members to groups, 339–340
- capabilities of, 221, 256
- central access rules, 299–300
- contact information, setting, 349–350
- creating computer accounts with, 260–261
- creating domain user accounts, 333–334
- deleting, disabling, and enabling computer accounts, 264
- domains, connecting to, 257–258
- editing properties of computer accounts, 263
- enabling disabled accounts, 374
- global group creations, 336–337
- moving computer accounts, 266
- multiple user accounts, managing, 375–378
- Organization panel, 349
- OUs (organizational units), managing, 281
- primary group selection, 340
- resetting locked computer accounts, 264–266
- searching with, 258–259
- unlocking accounts, 373–374

### Active Directory Domain Services, 230

### Active Directory Domain Services Installation Wizard, 282

### Active Directory Installation Wizard, 270–272

### Active Directory Sites And Services, 229–230, 280–289

### Active Directory Users And Computers

- adding members to groups, 339–340
- Advanced Features, 253
- auditing objects of, 502–503
- contact information, setting for user accounts, 347–349
- copying domain user accounts, 371
- creating computer accounts with, 261–263
- creating domain user accounts, 330–334

- deleting, disabling, and enabling computer accounts, 264
- domains and controllers, connecting to, 254
- editing properties of computer accounts, 263
- enabling disabled accounts, 374
- folder set of, 252–253
- global group creations, 336
- home directory specification, 354–355
- listing accounts or resources, 368–369
- Logon Hours settings, 355–357
- moving computer accounts, 266
- multiple accounts, working with, 373
- multiple user accounts, managing, 375–378
- opening, 252
- operations masters, viewing, 273–274
- OUs (organizational units), managing, 228, 253, 281–282
- permissions, setting, 380–381
- primary group selection, 340
- renaming user accounts, 369–371
- resetting locked computer accounts, 264–265
- scope, selecting for a new group, 336, 337
- searching for directory objects, 254–256
- searching for users and groups, 350–351
- user profiles, setting for multiple accounts, 376–377
- active partitions, 402, 445
- AD CS (Active Directory Certificate Services), 14–15, 33
- AD DS (Active Directory Domain Services). *See also* Active Directory
  - adding to a server, 218
  - centrality to Active Directory, 15
  - Configuration Wizard, 218–219, 271–272
  - DNS, installing first, 616
  - Group Policy with. *See* GPMC (Group Policy Management Console)
  - installing domain controllers, 270–272
  - Restartable AD Domain Services, 16–17
  - role description, 33
- AD FS (Active Directory Federation Services), 15, 33, 297
- AD LDS (Active Directory Lightweight Directory Services), 15, 33
- AD RMS (Active Directory Rights Management Services), 15, 33
- Add Roles And Features Wizard, 70–72, 526–527
- Add Workstations To Domain privilege, 310
- Add Workstations To The Domain right, 267–268
- Adjust Memory Quotas For A Process privilege, 310
- administrative shares, 474–476
- administrative templates, 151–152
- Administrative Tools program group, 23
- Administrator account, 307–308, 317
- administrators, 315–317, 486–487
- Administrators group, 315–317, 486
- ADMX format, 141–142
- Adprep command, 232–235, 250–251
- ADSI (Active Directory Service Interface), 242
- ADSI Edit tool, 252, 289–290
- Adsiedit.msc, 252
- Advanced Configuration and Power Interface (ACPI), 8–11
- Advanced tab, System Properties dialog box, 76–85
- ADWS (Active Directory Web Services), 259
- Aero, Windows, 5
- AES (Advanced Encryption Standard), 361, 552
- alerts, performance counter, 129–130
- allocation units, 409, 411, 425, 442
- alternate credentials, 58
- alternate IP addressing, 568–569
- Anonymous Logon identity, 317
- antimalware software, 546
- Apple computers, 340, 463, 479–481
- application directory partitions, 238
- Application Server, 33
- applications
  - Applications log, 108
  - Apps, 7, 137, 363
  - backups of data files, 528, 534
  - performance settings, 77
  - restoring, 550–551
  - Server Core installations not for running, 40
  - updates for, 541–542
- applications and services logs, 108–109
- Apps, 7, 137, 363
- archive attribute, 521
- archives, event logs, of, 116–118
- ARP command, 48
- ARP test, 574
- attached storage, standards-based techniques, 419

audio, 37

auditing

- Active Directory objects, 502–503
- DHCP server, 585–586
- file or folder specific, 500–502
- Generate Security Audits privilege, 311
- logon failures, 379
- registry, of the, 502
- rights required to change, 498
- setting policies for, 210, 498–503

authentication. *See also* accounts; passwords

- Active Directory authentication mechanism assurance, 220
- Active Directory Domain Services role in, 15
- Authenticated Users identity, 317
- claims-based access controls, 496–498
- Kerberos for. *See* Kerberos protocols for, 296–297
- remote server management issues, 26–27
- Rights Management Services, 15
- security policies for, 209–210

autoenrollment, 186–187

autoloader tape systems, 523

Automatic Updates, 135–136, 187–190

- WSUS (Windows Server Update Services), 189–190

## B

Background Intelligent Transfer Service (BITS), 35

background processes, 87

backups

- application data, 528, 534
- archive attribute, 521
- Back Up Files And Directories privilege, 310
- cloud-based, 525, 527
- configuring, 527–533
- copy backups, 521
- critical volumes only option, 534
- daily, 521, 532
- differential backups, 521–522, 525
- DVDs for, 525, 535, 541
- encryption certificates, of, 554–555
- full, 521, 525
- of GPOs (Group Policy Objects), 165
- hardware considerations, 520, 522–524
- incremental backups, 521–522, 525, 528–529

LTO drives, 523

manual, 540–541

media options for, 524

Microsoft Online Backup Service, 525, 527

normal, 521

- off-site backups, 520
- permissions for, 528
- planning considerations, 519–521
- remote shared folders for, 534, 539, 541
- restoring data from. *See* restoring scheduling automatic, 535–540
- scheduling considerations, 520
- shadow copies as supplements to, 481–484, 520
- storage location specification, 534–535
- system state, 528, 546
- utilities for, list of, 525
- volume options for, 534
- Wbadmin. *See* Wbadmin backup command
- Windows Server Backup for. *See* Windows Server Backup

Balanced power plan, 8

baseline metrics, establishing, 119

basic disks, 395, 401–404, 421

basic permissions, 488–493

batch scripts, 49, 178–179, 317

BCD Editor, 433–434

Best Practices Analyzer, 60

binaries, 24, 57, 63–65, 72, 218

biometric security, 38

BitLocker Drive Encryption, 35, 136, 404

BITS (Background Intelligent Transfer Service), 35

boot logging, enabling, 545

booting

- BCD Editor, 433–434
- boot partitions, 402, 433–434, 443
- setting default startup OS, 83–84
- Startup Repair (StR), 543
- Windows Boot Manager, 83–84, 447
- Windows Preboot Environment, 5, 447–449

BOOTP (Bootstrap Protocol), 601

BranchCache, 36, 386

BranchCache for Network Files, 34

bridgehead servers, 291–292, 293

built-in accounts, 306

built-in capabilities for groups, 313–315

built-in groups, 308–309

built-in local groups, 303

built-in user accounts, 307  
 Bypass Traverse Checking privilege, 310

## C

- Cached statistic, 98
- CALL command, 49
- callback parameters, 359
- CAPI2 (CryptoAPI Version 2), 563
- CAs (certificate authorities), 14–15, 33, 75
- cellular network policies, 168
- central access rules, 298–300
- certificate authorities (CAs), 14–15, 33, 75
- certificates
  - AD CS (Active Directory Certificate Services), 14–15
  - autoenrollment, 186–187
  - backing up, 554–555
  - CAs (certificate authorities), 14–15, 33, 75
  - DNSSEC, 632–634
  - file encryption, 414
  - recovery of, 551–553
  - Rights Management Services, 15
  - roaming profiles with, 362
  - role services for, 33
- Change Permissions special permission, 490–491
- Change share permission, 470
- Change Time privileges, 310
- Character Map, 5–6
- charms, 6–7
- Check Disk, 49, 450–453
- child domains, 18, 218, 616, 634–636
- child objects, 487
- chkdsk, 49, 450–453
- CHKNTFS command, 49
- claims policy, 221
- Claims-Aware Agent, 33
- claims-based access controls, 297–299, 496–498
- clean installations, 43–47
- Client for NFS, 36
- clustering, 36
- cmdlets, 24–25. *See also* Windows PowerShell 3.0
- CNAME (canonical name) records, 636, 638
- commands. *See also specific commands*
  - console for. *See* Windows PowerShell 3.0
  - mini Windows PC, table of, 48–52
  - NET HELP command, 23
  - opening new command prompts, 41
  - safe mode with prompt, 545
  - Search box execution of, 137
- Comma-Separated Value Directory Exchange (CSVDE), 372
- components, removability of, 5
- compound identities, 297–298
- compression, 49, 410–412, 425
- computer accounts
  - auditing, 503
  - claims-based access controls, 496–498
  - Create Account Objects privilege, 268
  - creating, 259–263
  - editing properties of, 263
  - group memberships of, 496
  - groups, adding computers to, 339–340
  - listing, 368–369
  - managed computer accounts, 262–263
  - managing, 264
  - moving, 266
  - name protection, 586
  - objects for, 231
  - passwords for, 264–266
  - permissions, setting, 380–381
  - protected, 264
  - resetting locked, 264–266
- computer groups, 339–340, 503
- Computer Management
  - archiving event logs, 116–118
  - clearing event logs, 116
  - closing file shares, 478
  - configuring share permissions, 470–472
  - creating shared folders, 463–466
  - disabling services, 107
  - event logs options, setting, 115–116
  - Event Viewer, 111–112, 116–117
  - file shares, viewing, 462
  - file sharing, viewing sessions of, 476–478
  - joining computers to domains or workgroups, 267–270
  - Open Files node, 477–478
  - opening, 102, 267
  - pausing services, 103
  - publishing shares, 469
  - remote servers, working with, 102
  - service configuration, 101–107
  - sessions, ending, 477
  - starting services, 103
  - stopping services, 103
  - stopping sharing on folders, 478–479
- Computer Name tab, 74–75

## contact information

- contact information, setting for user accounts, 347–350
- contacts, adding to groups, 340
- containers, 138–139
- Control Panel, 7, 23, 151, 542
- Convert utility, 445–447
- cooling modes, 8–9
- copy backups, 521
- COPY command, 49
- core installations. *See* Server Core installations
- cores, multiple, 6, 10–11
- costed networks, 168
- counters, performance, 122–126, 129–133
- crashes
  - automatic rebooting on fatal errors, 84
  - crash dump partitions, 402
  - improvements in recovery mechanisms for, 541–542
  - Microsoft Online Crash Analysis, 543–544
  - Stop errors, 84–85
- Create Files/Write Data special permission, 490–491
- Create Folders/Append Data special permission, 490–491
- Create privileges, 310–311
- Creator identities, 317
- credentials, preventing delegation of, 361
- Critical events, 112
- CryptoAPI Version 2 (CAPI2), 563
- Cscript command, 42
- c-states, 10–11
- CSVDE (Comma-Separated Value Directory Exchange), 372

## D

- D states, 11
- daily backups, 521, 532
- DAT (digital audio tape) drives, 523
- Data Center Bridging, 36
- data collector sets, 124–130
- Data Deduplication service, 34, 386
- Data Execution Prevention (DEP), 80–81, 92
- data scrubbers, 393
- data stores, 14, 238–239
- DATE command, 49
- Dcgpofix, 172–173
- DCOM Server Process Launcher, 106
- Debug Programs privilege, 311
- decrypting files and directories, 418
- deduplication, 34, 386, 436, 441
- Default Domain Controller Policy GPO, 147, 172–173
- Default Domain Controllers GPO, 497
- Default Domain Policy GPO, 172–173
- default gateways, 568–570
- default group accounts, 315–318
- defragmenting drives, 453–455
- DEL command, 49
- delegation, 311, 361
- Delete special permission, 490–491
- Delete Subfolders And Files special permission, 490–491
- DEP (Data Execution Prevention), 80–81, 92
- deploying software using Group Policy, 181–186
- descriptions, user, 348
- desktop, 7, 151
- Desktop Experience, 4–6, 39, 56, 136
- device drivers, 42, 53, 75–76, 311, 545
- devices
  - accounts for. *See* computer accounts
  - claims, 495–498
  - power states of, 11
- DFS (Distributed File System), 34, 109, 142, 151, 252, 386–387
- DHCP (Dynamic Host Configuration Protocol)
  - DHCIDs, 586
  - DNS integration with, 586–587, 616, 618
  - DNS with, 19–20
  - failover service, 574–575
  - hot standby mode, 574–575
  - IP address conflicts, avoiding, 591
  - IPv4 addressing, 573–575
  - IPv6 configuration, 575–577
  - leases, 574, 577–578, 591, 594–598, 601, 609
  - load balance mode, 574
  - name protection, 586
  - NAP with, 588–591
  - number of servers needed, 574
  - relay agents, 575, 577, 579
  - reserving addresses, 608–609
  - responsibilities of, 573
  - router advertisement messages, 576–577
  - scopes with, 578–579, 582
  - server. *See* DHCP server
  - Server Core default, 46
  - stateful mode, 575–577
  - stateless mode, 575–577

- WINS with, 21
- DHCP server
  - activation and deactivation of scopes of, 601
  - auditing, 585–586
  - authorizing in Active Directory, 583
  - backing up the database, 610–611
  - bindings, configuring, 584
  - credentials specification, 580–581
  - database files for, 610
  - deleting leases and reservations, 610
  - deleting scopes, 602
  - DHCP console, 581–582
  - DNS integration with, 586–587
  - exclusion ranges, setting, 607–608
  - failover scopes, 602–605
  - forcing database regeneration, 612
  - groups delegated for administration, 580
  - installing, 579–581
  - IP address conflicts, avoiding, 591
  - MAC address filtering, 606–607
  - modifying reservation properties, 609
  - moving the database, 611–612
  - multicast scopes, 598
  - multiple scopes, 602
  - Network Policy Server, setting up as, 588–591
  - normal scopes, 593–598
  - reconciliation, 612–613
  - remote connections, 582–583
  - reserving addresses, 608–610
  - restoring the database, 611
  - role description, 33
  - saving configurations, 591–592
  - scope indicators, 582
  - scope options, 599–601
  - service, 583
  - source files for, 580
  - starting and stopping, 583
  - statistics configuration, 584
  - status indicators, 582
  - superscopes, 579, 592–593
  - viewing scope statistics, 605
- dial-up users, 317, 358–360, 359
- differential backups, 521–522, 525
- direct reports, setting for users, 350
- DirectAccess, 13, 34
- directories, domain. *See* Active Directory directories, file, 49, 412–413, 415–416. *See also* folders
- Directory Service log, 109
- directory services. *See* Active Directory
- Directory Services Access Control Lists Utility, 252
- Directory Services Restore Mode, 16–17
- Disable Automatic Restart On System Failure, 545
- disable- cmdlets, 25
- discovery, network, 12, 560–562
- Disk Cleanup, 5
- disk controllers, duplexing, 429
- disk drives. *See* hard disk drives
- disk I/O performance tuning, 133
- Disk Management
  - active partitions, marking, 402
  - assigning drive letters, 443–444
  - converting basic disks to dynamic disks, 402–403
  - creating partitions, 407–410
  - creating shadow copies, 482
  - creating volumes, 407–410, 424–425
  - deleting partitions, 445
  - disk mirroring with, 429–431
  - disk striping with parity, 431
  - drive path management, 443
  - drive status values, table of, 399–400
  - extending volumes, 449
  - hot swapping disks, 398
  - initializing disks, 399
  - moving dynamic disks to new systems, 404–405
  - Optimize Drives, 453–455
  - partition color coding scheme, 407
  - properties, viewing, 395–396
  - remote management limitations, 394, 396
  - removing a mirror, 434
  - Rescan Disks command, 398, 405
  - resizing partitions, 447–449
  - shrinking volumes, 448–449
  - striped set creation, 428
  - VHDs (virtual hard disks), managing, 405–406
  - views, 394
  - volume status issues, table of, 422–423
  - volumes, viewing properties of, 420–421
- disk mirroring, 426–434
- disk quotas
  - application issues, 504
  - NTFS. *See* NTFS disk quotas
  - Resource Manager, 503–504, 514–518
- disk striping, 426–428, 431, 434–435, 440
- DISKPART command, 49

- Diskraid.exe, 42
- display names, 319, 331, 333–334
- displays, 7, 76–77
- Distributed File System (DFS), 34, 109, 142, 151, 252, 386–387
- distribution groups, 303
- DLLs (dynamic-link libraries), 39, 51
- DNS (Domain Name System)
  - Active Directory, integrating with, 616–617, 625–626, 628
  - Active Directory, use by, 217–218
  - aliases, 638
  - assigning for IP address configuration, 568–569
  - child domains, 18, 281, 616
  - clients, 617–621
  - configuring in Advanced TCP/IP Settings, 620–621
  - DHCP integration with, 19, 586–587, 616, 618
  - DNSSEC, 19–20, 632–634
  - domains vs. Active Directory domains, 230
  - dynamic update options, 626, 645
  - enabling on a network, 617–619
  - FQDNs, 218–219, 616
  - full integration with Active Directory, 19, 616
  - GlobalNames zone, 619, 629–630
  - host name structure, 615
  - intranets, relation to, 615
  - IPv6 with, 617–618
  - LLMNR with, 618
  - log, DNS Server, 109
  - Manager console. *See* DNS Manager console
  - name protection, 586
  - organization of, 18–20
  - parent domains, 18, 218, 615–616
  - partial integration with Active Directory, 18–19, 616
  - port used by, 293
  - record types for, 636–637
  - replication across networks, 616–617
  - RODCs with, 19, 219, 619
  - security issues, 626–627, 632–634
  - servers. *See* DNS servers
  - single-label name resolution, 619
  - SOA (start of authority) records, 637, 641–642
  - structure of, 615–616
  - zones, 616–617, 625–630
- DNS Manager console
  - adding or removing servers, 631–632
  - address and pointer records, adding, 636–638
  - child domain creation, 634–636
  - CNAME records, adding, 638
  - configuring servers, 623–627
  - debugging, logging, 648–649
  - deleting domains or subnets, 636
  - DNSSEC configuration, 632–634
  - dynamic update settings, 645
  - editing records, 640
  - event logging settings, 648
  - expiration parameter, 642
  - forwarding options, setting, 646–648
  - IP addresses, disabling, 646
  - monitoring options, 649–650
  - MX (mail exchange) records, 637–639
  - name servers, adding, 639–640
  - primary server designation, 642
  - record management with, 636–640
  - refresh intervals, 642
  - responsible persons, 642
  - serial numbers, 642
  - SOA record modification, 641–642
  - starting or stopping DNS servers, 632
  - structure of, 631
  - TTL parameters, 642
  - viewing records, 640
  - zone properties, setting, 641–645
- DNS servers
  - Active Directory–integrated primary servers, 621
  - configuring during setup, 623–624
  - configuring primary servers, 624–627, 642
  - configuring secondary servers, 627, 643–644
  - dynamic update options, 626, 645
  - event logging, 648
  - forwarding options, 646–648
  - forward-only servers, 622
  - GlobalNames zone configuration, 619, 629–630
  - installing services for, 622–624
  - IP addresses, disabling, 646
  - Manager console. *See* DNS Manager console
  - new zone configuration, 625–626
  - primary servers, 621, 624–627, 642
  - restart times, 617
  - reverse lookups, 624, 628–629



- role, DNS Server, 33
- root hints, 624
- secondary servers, 622, 627, 643–644
- server addresses, order of use, 620
- service, 616–617, 622–623, 632
- starting or stopping from Manager console, 632
- startup tasks of, 617
- troubleshooting, 252, 649–650
- well-known site-local addresses of servers, 617
- DNS64 support, 13
- Dnscmd.exe, 252
- DNSSEC (DNS Security Extensions), 19–20, 632–634
- domain accounts
  - creating, 330–334
  - display names of, 319
  - expiration dates, changing, 375
  - listing, 368–369
- Domain Admins group, 309, 315–317
- Domain Computers group, 339
- domain controllers
  - AD DS (Active Directory Domain Services), 15
  - bridgehead servers, 291–293
  - claims-based policy for, 497
  - creation with Active Directory Domain Services, 230
  - default GPOs of, 147
  - demoting, 272–273
  - Directory Services Restore Mode, 16–17
  - DNS installation, 622
  - Domain Controllers group, 339
  - functional levels of, 225–227
  - global catalogs, designation of, 239–240
  - grouping, 58
  - infrastructure masters, 239, 243, 273–280
  - installing, 270–273
  - multimaster replication model, 14
  - names of, changing, 75
  - operations masters, 230–231
  - ports used by, 292–293
  - purpose of, 6, 14
  - read-only, 16, 19, 219, 619
  - refreshing Group Policy, 158–161, 166
  - replication of directory data, 230–231
  - restoring system state, 547
  - sites, associating with, 285
  - states of Active Directory, 17
  - Synchronize Directory Service Data privilege, 312
  - verifying installations of, 218–219
  - virtual domain controller cloning, 223
- domain joins, off-premises feature, 222
- domain local groups, 303–306
- Domain Name System (DNS). *See* DNS (Domain Name System)
- domain naming master role, 242, 244, 275–276, 276–280
- domain networks, 560
- domain user accounts
  - copying to create new, 371
  - creating, 330–334
  - definition of, 301
  - deleting, 373
  - logon troubleshooting, 379
  - properties, configuring, 375–378
- Domain Users group, 338
- domains
  - AD DS. *See* Active Directory; AD DS (Active Directory Domain Services)
  - CAs (certificate authorities) for, 14–15, 33, 75
  - child domains, 18, 218, 634–636
  - controllers. *See* domain controllers
  - default GPOs, 147
  - definition of, 6
  - DNS vs. Active Directory, 230
  - Domain network type, 12
  - forests of, 225–227
  - FQDNs (fully qualified domain names), 101, 218–219, 301–302, 616
  - functional levels of, 225–227
  - GPMC Domains node, 149
  - joining computers to, 75, 267–270
  - local GPO issues in, 143
  - member servers of, 14
  - name system. *See* DNS (Domain Name System)
  - parents, 218
  - permissions to manage Group Policy, 155
  - remote management within, 65
  - root domains, 18
  - structural overview of, 223
  - subdomains, 18
  - time synchronization, 137
  - top-level (TLDs), 218
  - trees in, 225–227
  - Windows Domain Manager, 252

- Domains And Trusts snap-in, 226–227
- drive letters
  - assigning, 407, 408, 425, 443–444
  - changing on standards-based systems, 436
  - moved disks, of, 405
  - standard volumes, assigning to, 442
- drive paths, 407, 443
- Driveletter\$ share, 475
- drivers, device, 42, 53, 75–76, 311, 545
- drives
  - basic, 395, 401–404, 421
  - dynamic. *See* dynamic drives
  - hard disk. *See* hard disk drives
  - network drives, 484–485
  - partitioning. *See* partitions
  - tape drives, 523–525
  - volumes on. *See* volumes
- Ds commands, list of, 251
- Dsacls.exe, 252
- Dsquery commands, 242–243, 251
- dual boot systems, 421
- dump files, 84–85
- DVDs for backups, 525, 535, 541
- dynamic drives
  - capabilities of, 401
  - compared to other types, 395, 421
  - converting to and from basic disks, 402–404
  - drive section types, 402
  - moving to new systems, 404–405
  - RAID capabilities of, 420
  - reactivating, 404
  - UNIX compatibility of, 402
- Dynamic Host Configuration Protocol. *See* DHCP (Dynamic Host Configuration Protocol)
- dynamic IP addressing, 568–569
- dynamic updates, 645. *See also* DHCP (Dynamic Host Configuration Protocol)
- dynamic-link libraries (DLLs), 39, 51

## E

- ECHO command, 49
- editions of Windows Server 2012, 6
- EFS (Encrypting File System)
  - backing up certificates, 554–555
  - copying encrypted files, 416
  - data recovery system of, 414–415, 417
  - decrypting files and directories, 418
  - granting special access to encrypted files, 416
  - mechanics of, 414
  - NTFS with, 413
  - ownership of files, 414
  - recovery policy, configuring, 417, 551–553
  - restoring certificates, 555
  - roaming profiles with, 362
  - steps for encrypting files, 415–416
- elevation prompts, 5, 6, 93
- email
  - distribution groups, 303
  - setting for user accounts, 348
  - SMTP Server, 37
- enable- cmdlets, 25
- encryption
  - AES (Advanced Encryption Standard), 361
  - backing up certificates, 554–555
  - BitLocker Drive Encryption, 35
  - compression not allowed with, 413
  - copying encrypted files, 416
  - data recovery system for files, 414–415, 417
  - decrypting files and directories, 418
  - drive hardware encryption, 391
  - EFS. *See* EFS (Encrypting File System)
  - mechanics of, 414
  - NTFS, with, 413, 416
  - recovery policy for, 551–553
  - restoring certificates, 555
  - SMB, 457, 459
  - steps for encrypting files, 415–416
  - user profiles for certificates, 414
- End Task command, 89–90, 94
- enhanced Active Directory recycle bin, 247–248
- Enhanced Fine-Grained Password Policy, 222
- Enhanced Storage, 36, 386
- Enterprise Admins group, 315–317
- environment variables
  - configuring, 81–83
  - SET commands, 52
  - %UserName%, 352–353, 355, 363–364, 376–377
- errors. *See also* events
  - event levels of, 112
  - improvements in recovery mechanisms for, 541–542
  - Microsoft Online Crash Analysis, 543–544

- Stop errors, 84–85
- eSATA, 397
- event logs. *See also* logs
  - archives, of, 116–118
  - clearing, 116
  - custom views, creating, 112–114
  - data collector sets, 124–130
  - Data column, 112
  - DNS servers, of, 648
  - Event Viewer, 111–115
  - filtering, 112–115
  - Forwarded Events log, 108
  - levels of events, 111–112
  - list of available, 108–109
  - Log Properties dialog box, 115–116
  - options for, setting, 115–116
  - overwriting modes, 116
  - properties, columns of, 110, 112
  - remote management of, 67
  - security templates for, 193–195
  - Server Manager for viewing, 109–115
  - size, setting maximum, 115
  - task categories of, 112
  - viewing archives of, 118
  - Windows Event Log service, 108–109
- Event Viewer, 4, 111–115
- events
  - Audit Success or Failure, 112
  - command-line commands for configuring, 43
  - Critical, 112
  - Error, 112
  - Event Viewer, 111–115
  - Events, Server Manager, 60
  - IDs of, 110, 112
  - informational, 111
  - levels of, 111–112
  - properties of, 110
  - Server Manager for viewing, 109–115
  - task categories of, 112
  - traces, data collector sets, 125
  - Warning, 112
- Everyone group
  - Full Control share permission default, 468
  - Guest account membership in, 308
- Everyone identity, 318
- EXIT command, 49
- expired accounts, 375
- exporting accounts, 372
- extended partitions, 407

## F

- Failover Clustering, 36
- failover scopes, DNS, 579, 602–605
- failures, system, recovery mechanisms for, 541–544
- fan speeds, 8
- FAT file systems, 392–393, 445–447, 451–452, 458
- fault tolerance, RAID for, 426–427
- Fax Server, 4, 34
- FAX\$ share, 474
- features
  - Add Roles And Features Wizard, 58, 70–72
  - adding with ServerManager, 63
  - add-windowsfeature command, 63
  - associations with types of installations, 55–56
  - configuring, tool for, 32
  - get-windowsfeature command, 63
  - Remove Roles And Features Wizard, 71, 72–73
  - subordinates, adding all, 63
  - table of, 35–39
- federation services, 15, 33, 297
- File And Storage Services, 34, 394, 468–469
- file compression, 49, 410–412, 425
- File Explorer
  - Control Panel access from, 7
  - decrypting files and directories, 418
  - file encryption with, 415–416
  - network drives, 484–485
  - object ownership, 486–487
  - permissions, setting, 487–488, 491–495
  - restoring shadow copies, 482–483
  - special shares, connecting to, 475–476
  - viewing permissions, 488–489
- file paths, security templates for, 198–201
- File Replication Service log, 109
- File Server Resource Manager (FSRM), 387, 513–518
- File Server role service, 34, 387
- File Server VSS Agent Service, 387
- file servers, 385–389, 488
- file sharing
  - administrative shares, 474–476
  - advanced management properties, 468
  - Advanced Sharing Settings, 460–461
  - central access policies, 459
  - claims-based access controls, 459

## file systems

### file sharing, *continued*

- clients currently accessing, viewing
  - number of, 462
- closing open shares, 478
- creating shared folders, 463–468
- Enable Access-Based Enumeration option, 467
- File And Printer Sharing, enabling or disabling, 461
- file systems supporting, 458
- get-smbshare cmdlet, 461
- hidden shares, 464, 474–476
- Hyper-V issues, 468
- modifying settings for folders, 468–469
- multiple shares on single folders, 465–466
- naming shares, 464, 466
- net session command, 476–477
- net share command, 461
- NFS for, 34, 36, 463, 466, 479–481
- NTFS permissions, viewing, 488
- offline setting options, 464–465, 467
- permissions for. *See* share permissions
- public folder sharing, 458–460
- publishing shares, 469
- quota templates, 468
- SMB protocol for, 209, 457, 459, 461, 466, 468
- special shares, 474–476
- standard, 457–468
- stopping sharing on folders, 478–479
- viewing existing shares, 461–463
- viewing sessions of, 476–478

### file systems

- allocation unit size, setting, 409, 411
- changing on standards-based systems, 436
- determining type of, 421
- FAT, 392–393, 445–447, 451–452, 458
- list of supported, 392–393
- MFT (master file table), 393
- NTFS. *See* NTFS file system
- organization of, 385
- partitions, relation to, 392
- ReFS (Resilient File System), 393, 449, 459, 504
- repairing errors, 437
- selecting for partitions, 409, 411
- standard volumes, formatting with, 442

### files

- basic permissions for, 491–493

- compressing, 412
  - encrypting, 415–416, 551–553. *See also* EFS (Encrypting File System)
  - expanding compressed, 413
  - inheritance options for, 495
  - restoring, 550–551
  - searching for, 7
  - special permissions for, 490, 493–495
- ### FIND command, 49
- fingerprint-based authentication, 38
- ### FireWire, 396–397
- ### firmware, ACPI version compliance, 10
- ### folders
- encrypting, 415–416
  - inheritance options for, 495
  - NTFS permissions, viewing, 488
  - parent-child relationships, 487
  - redirection, 173–177
  - restoring nonsystem, 550–551
  - security templates for, 200–201
  - setting permissions for, 491–495
  - sharing. *See* file sharing; shared folders
  - special permissions for, 491
- ### FOR command, 49
- ### Force Shutdown From A Remote System privilege, 311
- ### foreground processes, 87
- ### forests, domain, 148, 225–227, 317
- ### FORMAT command, 49
- ### formatting partitions, 407, 409–411, 425
- ### forward lookups, 624
- ### Forwarded Events log, 108
- ### FQDN (fully qualified domain name)
- group account names using, 302
  - logon names using, 301–302
  - services, identifying for, 101
  - structure of, 218–219, 616
- ### FSMO (flexible single-master operation)
- role, 71, 273, 276. *See also* operations masters
- ### FSRM (File Server Resource Manager), 387
- ### Fsutil, 389
- ### FTP command, 49
- ### full backups, 521, 525, 528–529
- ### Full Control permission, 489
- ### Full Control share permission, 470
- ### full-server installations
- features included in, 4, 40
  - minimal-interface installations, converting to and from, 56

- functional levels
  - 2003 mode, 231, 232–233, 235–236
  - 2008 mode, 231, 233–234
  - 2012 mode, 231–232, 235
  - 2008 R2 mode, 231–232, 234–235
  - Adprep for upgrading, 232–235
  - lists and descriptions of, 225–227, 231–232
  - minimal level requirement, 232
  - raising, 236–237

## G

- G states, 11
- Gadgets, Windows, 5
- gateways, 568–570
- Generate Security Audits privilege, 311
- get- cmdlets, 25
- get-help cmdlet, 25
- get-service command, 26
- get-smbshare command, 461
- global catalogs
  - authentication role of, 231
  - configuring, 239, 280
  - domain controller assignment of, 239–240
  - ports used by, 292
  - universal membership caching, 281
- global groups
  - accounts, creating, 336–337
  - best use of, 305–306
  - definition of, 303
  - scope capabilities, 304
- global user rights configuration, 328–329
- GlobalNames zone, 619
- gpedit.msc command, 144
- GPMC (Group Policy Management Console)
  - Active Directory, use of, 149
  - Add Forest command, 148
  - administrative templates, 150–152
  - ADMX format, 141–142
  - auditing, configuring, 498–500
  - autoenrollment, 186–187
  - Automatic Update management, 187–190
  - backing up GPOs, 165
  - blocking inheritance, 156–158
  - central store, 142
  - Computer Configuration node, 149–150
  - copying GPOs, 164
  - creating GPOs, 152–153
  - creating local policies, 146
  - delegation of privileges to manage, 154–155
  - deleting GPOs, 170–171
  - deploying software, 183
  - disabling GPOs, 144, 166
  - Domains node, 149
  - editing a GPO, 144–145, 149
  - Editor choices, 141
  - enforcing inheritance, 157–158
  - gpmc.msc command, 148
  - Group Policy Modeling node, 149
  - Group Policy Object Editor, 141
  - Group Policy Results node, 149
  - importing GPOs, 164–165
  - inheritance, managing, 155–158
  - installing, 141
  - link order, changing, 156
  - link removal, 170–171, 213
  - linking GPOs to containers, 152–153
  - links vs. accessing actual GPOs, 149
  - local GPOs, 143–146
  - Local Object Editor, 141
  - loopback processing settings, 167
  - Modeling Wizard, 161–163
  - node structure, 148
  - opening, 148
  - opening local policies, 146
  - OS version issues, 141–143
  - permissions to manage GPOs, granting, 155
  - Policies node, 150
  - precedence preferences, setting, 156, 167
  - Preferences node, 150
  - redirecting special folders, 174–177
  - refreshing Group Policy, 158–161, 166
  - restoring GPOs, 165–166
  - Results node, 166, 171
  - RSoP (Resultant Set of Policy), 166, 171–172
  - script assignment with, 178–181
  - security template deployments, 205
  - shortcuts, 150
  - Sites node, 149
  - slow-link detection, configuring, 167–170
  - Software Settings node, 150
  - starter GPOs, 141, 153
  - troubleshooting, 171–172
  - User Configuration node, 150
  - Windows Settings node, 150

- GPOE (Group Policy Object Editor). *See* GPMC (Group Policy Management Console)
- GPOs (Group Policy Objects)
  - account policies, setting, 321–322
  - administrative templates, 151–152
  - Administrators and Non-Administrators local layer, 143–144, 146
  - auditing, configuring, 498–503
  - backing up, 165
  - blocking inheritance, 156–158
  - central stores for, 142
  - claims-based policy deployment, 299–300
  - conflict resolution, 144
  - containers, 138–139
  - copying, 164
  - creating, 152–153
  - default domain policies, 147, 172–173
  - definition of, 138
  - delegation of privileges to manage, 154–155
  - deleting, 170–171
  - disabling local, 144
  - disabling unused parts of, 166
  - editing. *See* GPMC (Group Policy Management Console)
  - event log, 143
  - folder for domain policies, 148
  - Folder Redirection, editing, 173–177
  - gpedit.msc command, 144
  - importing, 164–165
  - inheritance, 138–139, 155–158
  - link order, changing, 155–156
  - link removal, 170–171
  - linking to containers, 152–153
  - local, 143–146
  - local Group Policy Object, 143–146
  - logon/logoff script assignment with, 180–181
  - loopback processing settings, 167
  - Machine folder, 145–146, 148
  - precedence of, 152, 155–158, 167
  - priority of local GPOs, 144
  - refreshing, 158–161, 166
  - replication of, 142
  - restoring, 165–166
  - RSoP (Resultant Set of Policy), 171–172
  - security policies, including in, 212
  - security template deployments, 204–205
  - shutdown script assignment with, 178–179
  - starter GPOs, 153
  - startup script assignment with, 178–179
  - User folder, 145–146, 148
  - user rights assignments to, 328
  - User-specific local Group Policy layer, 144, 146
- Gpresult command, 171–172
- GPT (GUID partition table), 392, 399, 403, 407
- Gpupdate, 161
- Graphical Management Tools And Infrastructure, 57
- Graphical Shell, 4, 56
- group accounts
  - basic operation of, 302–303
  - built-in capabilities for, 313–315
  - creating, 335, 336–338
  - default groups used by administrators, 315–317
  - deleting, 373
  - global, creating, 336–337
  - implicit groups, 317–318
  - local, creating, 337–338
  - logon rights for, 312–313
  - names of, 302
  - privileges available, table of, 310–312
  - SIDs of, 304–305
  - tools for creating, 318
  - user accounts compared to, 300–301
- Group Managed Service Accounts, 222
- Group Policy
  - Active Directory, relation to, 138
  - administrative templates, 151–152
  - auditing, configuring, 498–503
  - autoenrollment, 186–187
  - Automatic Updates, managing, 187–190
  - blocking inheritance, 156–158
  - capabilities of, 138
  - claims-based policy, 298
  - compatibility with OS versions, 140
  - console. *See* GPMC (Group Policy Management Console)
  - containers, 138–139
  - delegation of privileges to manage, 154–155
  - deleting, 170–171
  - deploying software using, 181–188
  - disabling unused parts, 166
  - disk quota issues with, 504



hard disk drives, *continued*

- drive interface architectures supported, 390
- drive paths, 407
- drive section types, 402
- drivers for, loading during installations, 53
- dynamic. *See* dynamic drives
- expanding compressed drives, 412–413
- fault tolerance, 426–427
- file systems of. *See* file systems
- Foreign status, 400
- FORMAT command, 49
- Free Space, 394
- GPT (GUID partition table), 392, 399, 403, 407
- hardware encryption, 391
- hot swapping, 398
- initializing, 398–399
- installing, 398–399
- listing for a computer, 405
- logical drives, 392, 401, 404, 407
- manager for. *See* Disk Management
- maximum sustained data transfer rate, 391
- MBR (master boot record), 392, 399, 403, 406–407
- mean time to failure, 391
- moving to new systems, 404–405
- network drives, 484–485
- No Media status, 400
- Not Initialized status, 400
- Offline status, 399, 404
- Online status, 399, 404
- optimizing, 453–455
- paging, configuring, 77–80
- partitions. *See* partitions
- physical drive formats, 389–392
- physical sector sizes, 389
- planning considerations, 389
- preparing for use, 392–394
- properties, viewing, 395–396
- reactivating, 404
- removable disks, 395
- removable storage devices, 396–398
- removing, 405
- rescanning all drives, 404
- resetting, 437
- Resource Monitor Disk usage statistics, 120
- rotational speed, 390–391

- selecting for inclusion in volumes, 424–425
- self-healing NTFS, 450–451
- Standard Format drives, 389
- standard volumes, creating, 441–443
- status values, table of, 399–400
- system volumes, 402
- temperature characteristics, 391
- Transactional NTFS, 449–450
- types of disk configurations, 395
- uninstalling, 405
- unmounted drives, 443
- Unreadable status, 400, 404
- Unrecognized status, 400
- VHD (virtual hard disk), 395
- virtual. *See* standards-based storage; virtual disks
- virtual hard disks, adding or removing roles, 71–72
- volume status issues, table of, 422–423

hardware

- backups, options for, 522–524
- built-in diagnostics for, 543
- configuration, 6
- CPUs. *See* processors
- disks. *See* hard disk drives
- Hardware Events log, 109
- Hardware tab access to configuration, 75–76
- independence, modularization for, 5
- hibernate state, absence of, 5
- hidden shares, 474–476
- home directories, user accounts, 351, 353–355, 371, 373, 377
- Home network type, 12
- hot swapping, 398
- hotfixes, 43
- Hyper-V
  - file sharing issues, 468
  - network configuration for, 570–571
  - service description, 34

I

- Identity Federation Support service, 33
- idle processor states, 10–11
- IE ESC (Internet Explorer Enhanced Security Configuration), 61
- IIS (Internet Information Services)
  - IIS (Web Server) role, 35
  - service accounts for, 341



- WinRM Extension for, 39
- images, Windows
  - DISM command, 49
  - features, installing from PowerShell, 63–64
- Impersonate A Client privilege, 311
- implicit accounts, 306
- implicit groups
  - Everyone group, 308
  - list of, 317–318
  - role of, 309
- importing accounts, 372
- Increase privileges, 311
- incremental backups, 521–522, 525, 528–529
- indexes, 221
- InetOrgPerson object, 301
- infrastructure masters
  - duties of, 243
  - global catalogs, relation to, 239
  - seizing, 276–280
  - transferring, 276
  - viewing, 273–274
- inheritance
  - ACEs, of, 297
  - GPOs, of, 138–139, 155–158
  - object permissions, of, 487–488
  - options for, 495
- Ink and Handwriting Services, 36
- installations
  - 64-bit processor requirement, 43–44
  - changing the type, 55–56
  - clean, steps for, 44–47
  - clean vs. upgrades, 43
  - command line use during, 48–49
  - disk space required for, 44
  - full-server installations, 4, 40, 56
  - hardware requirements, 43–44
  - keyboard layouts, 44–45
  - licensing terms, 45
  - minimal-interface installations, 4, 40, 56
  - partitions, managing during, 52–55
  - product keys, 45, 47
  - Server Core installations, 4, 40–47, 56–58, 526
  - Server With A GUI installations, 4. *See also* full-server installations
  - Server With Minimal Interface installation, 4, 40, 56
  - type of, choosing, 45, 48
  - types changeable after installation, 4
  - upgrades, steps for performing, 47–48
  - Where Do You Want to Install Windows, 45–46
- Interactive identity, 318
- Internet Explorer, 61
- Internet Printing, 4, 34, 36
- Internet Storage Naming (iSNS) Server Service, 36
- intersite replication topology, 291–292
- intranets, 15, 18
- IP addresses
  - assignment methods for, 566
  - conflicts, avoiding, 591
  - dial-in connections, assigning static, 359–360
  - disabling for DNS, 646
  - DNS records mapping to, 636–637
  - dynamic configuration of. *See* DHCP (Dynamic Host Configuration Protocol)
  - exclusion ranges, setting, 607–608
  - IP Address Management Server, 36
  - IPAM servers for managing, 137
  - multicast scopes, 598
  - name resolution for. *See* name resolution
  - normal scopes for, 593–598
  - reserving DHCP, 607–610
  - scopes of, 578–579, 582, 599–601
  - static addresses, 566–568
  - superscopes, 579, 592–593
- IPC\$ share, 474
- ipconfig
  - all parameter, 42
  - autoconfiguration addresses, 578
  - MAC addresses, 606–607
  - Mini Windows PC version, 50
  - releasing reserved addresses, 609–610
- IPv4 addresses
  - BOOTP (Bootstrap Protocol), 601
  - configuring, 564–569, 573–575, 591
  - conflicting addresses, avoiding, 591
  - decimal expressions of, 13–14
  - dynamic addressing, 573–575
  - failover scopes, 602–605
  - MAC address filtering, 606–607
  - normal scopes for addresses, 593–597
  - reserving addresses, 608
- IPv6 addresses
  - configuring, 564–569
  - decimal expressions of, 13–14
  - dynamic addressing with, 575–577

- IPv6 addresses, *continued*
  - enabling for client computers, 617–618
  - LLMNR for name resolution, 22–23
  - normal scopes for, 579, 596–598
  - reserving addresses, 608–609

## iSCSI

- iSCSI Target role services, 387
- iSNS Server Service, 36
- new virtual disk, creating, 436
- role services for, 436
- Target Server service, 34

- ISTG (Inter-Site Replication Topology Generator), 291–292

**J**

- joining computers to domains or workgroups, 267–270

**K**

## Kerberos

- account security options for, 361
- with Armoring, 222, 297–298, 496
- authentication with, 296
- claims-based access controls, 496–498
- configuring policies for, 321–322, 326–327
- constrained delegation across domains, 222
- preauthentication, 361
- tickets, 327

- Kerberos with Armoring, 222, 297–298, 496

## kernel, OS

- Kernel Transaction Manager (KTM), 450
- memory statistics, 98
- usage, tracking, 96

- key masters, 633–634

- keyboard layouts, 44–45

- KTM (Kernel Transaction Manager), 450

**L**

- LABEL command, 50

- LAN Manager authentication, 209–210

- language independence, modularization for, 5

- laptops running Windows Server, 8

- Last Known Good Configuration, 545

- latency, network, 134

- LDAP (Lightweight Directory Access Protocol)

- Active Directory Administration Tool, 252
- Active Directory use of, 242
- importing files with LDAP attributes, 372
- InetOrgPerson object for, 301
- ports used by, 292
- Security Configuration Wizard setting for, 209

## Ldp.exe, 246

- level 0, RAID, 426–428, 434

- level 1, RAID, 426–434

- level 5, RAID, 426–427, 431, 434–435

- licensing, 34, 42

- Linear Tape Open (LTO), 523

- link-layer filtering, 606–607

- Link-Local Multicast Name Resolution (LLMNR), 22–23

- List Folder Contents permission, 489

- List Folder/Read Data special permission, 490–491

- LLMNR (Link-Local Multicast Name Resolution), 22–23, 37, 618

## load balancing

- of failover scopes, 602–605
- Network Load Balancing (NLB), 37, 63

## local groups

- adding users to, 338
- creating, 337–338
- local GPOs, 143–146
- Local Group Policy Editor, 313–315
- local policies, 327–330
- security templates for local policies, 193–195
- types of, 303

## local profiles

- changing types of, 368
- creating, 363, 366
- managing, 366–368
- names of, 365
- purpose of, 362

## local user accounts

- assigning to local groups, 337–338
- creating new, 334–335
- definition of, 301

## Local Users And Groups

- local group creation, 337–338
- passwords, resetting, 373

- LocalService account, 307, 504

- LocalSystem account, 307

- locked out accounts, 373–374, 379

- lockout policies, 325–326

- logging. *See* auditing; logs

- logical drives, 392, 401, 404, 407–408, 445
- logical processors, 6, 10–11
- logical structures, domain, 223
- logoff scripts, 180–181
- logon names
  - creating for an account, 331, 333–334
  - naming schemes for, 320
  - numeric code assignment, 320
  - parts of, 301–302
  - pre-Windows 2000, issues with, 331, 333
  - rules for, 319
- logon rights, 312–313
- logon scripts
  - assigning, 180–181
  - environment variable for path, 352
  - paths, setting, 351
  - renaming, 371
  - specifying, 353–354
  - %UserName% in, 376–377
- logons
  - Anonymous Logon identity, 317
  - authentication protocols for, 296–297
  - error messages from, 379–380
  - failures, logging, 379
  - forcibly disconnecting users when logon
    - hours expire, 357
  - hours allowed for, managing, 355–357, 377
  - lockout policies, 325–326
  - names for. *See* logon names
  - options for controlling, 378
  - passwords for. *See* passwords
  - permitting by workstation, 357–358, 378, 380
  - rights, 312–313
  - scripts. *See* logon scripts
  - security tokens generated by, 305
  - Server Core, 40–41
  - troubleshooting, 378–380
  - user profile issues, 361. *See also* user profiles
  - validation of. *See* authentication
- logs
  - archives, of, 116–118
  - auditing with. *See* auditing
  - clearing, 116
  - data collector sets, 124–130
  - DHCP logs, 585–586
  - Event Viewer, 111–115
  - Hardware Events log, 109
  - levels of events, 111–112
  - list of available, 108–109

- Log Properties dialog box, 115–116
- overwriting modes, 116
- properties, columns of, 110, 112
- Server Manager for viewing, 109–115
- size, setting maximum, 115
- Stop error logging, 84
- Windows Event Log service, 108–109
- loopback processing, Group Policy, 167
- Lpd.exe, 252
- LPR Port Monitor, 36
- LTO (Linear Tape Open), 523
- LUNs. *See* virtual disks

## M

- MAC address filtering, 606–607
- MAC addresses, 13–14
- Manage Auditing And Security Log privilege, 312
- managed service accounts, 341–344
- mandatory user profiles, 362, 364, 379–380
- mapping network drives, 475–476, 484–485
- masters. *See* operations masters
- maximum processor state, 8–11
- MBR (master boot record)
  - basic to dynamic disk conversions, 403
  - GPT compared to, 392
  - number of allowed partitions, 406–407
  - selecting during initialization, 399
- Media Foundation feature, 36
- member servers, relation to domains, 14
- memory
  - allocating between OS and applications, 130
  - available, guideline for, 131
  - built-in diagnostics for, 543
  - Cached statistic, 98
  - counters, table of, 131–132
  - Data Execution Prevention, 80–81
  - dump files, 84–85
  - Lock Pages In Memory privilege, 312
  - NP (nonpaged) pool, 92
  - page faults, 92, 131–132
  - peak working set, process, 92
  - Performance tab, Task Manager for
    - statistics, 97–98
  - performance tuning, 130–132
  - processes, reserved for, 90
  - Resource Monitor usage statistics, 120
  - virtual, configuring, 77–80
  - Windows Memory Diagnostic Tools, 548
- Message Queuing, 36

MFT (master file table), 393  
 Microsoft Exchange Server  
   service accounts for, 341  
   Windows Server Backup with, 527  
 Microsoft Management Console, 4  
 Microsoft Online Backup Service, 525, 527  
 Microsoft Online Crash Analysis, 543–544  
 Microsoft Update, 35  
 MicrosoftWindows log, 109  
 migration, 47  
 mini Windows PC commands, 48–52  
 minimal-interface installations, 4, 40, 56  
 minimum processor state, 8–11  
 MINWINPC commands, 48–52  
 mirrored sets, 426–434  
 mobile device networking, 168  
 Modeling Wizard, Group Policy, 161–163  
 Modify Firmware Environment privilege, 312  
 Modify permission, 489  
 modularization, 5  
 monitoring servers, 118–119. *See also*  
   performance  
 monitors, 7, 76–77  
 MPIO (Multipath I/O), 36, 386  
 MS-DOS, RAID noncompliance of, 427  
 MSI (Microsoft Installer), 43, 182–183  
 multicast scopes, 579, 598  
 multimaster replication model, 14  
 Multipath I/O (MPIO), 36, 386  
 multiple operating systems, 83–84  
 multiple processor cores support, 6  
 multiple processor systems, 96, 122

## N

name protection, 586  
 name resolution  
   appending suffixes automatically, 620  
   client computer settings for, 620–621  
   Computer Name tab, 74–75  
   DNS for. *See* DNS (Domain Name System)  
   DNSSEC (DNS Security Extensions), 19–20, 632–634  
   forward lookups, 624  
   LLMNR for, 22–23  
   reverse lookups, 624, 628–629  
   services supported, 17  
   single-label name resolution, 619  
   WINS for, 20–21  
   name servers, adding, 639–640  
   names of computers, 46–47, 75  
   naming contexts, 289–290  
   NAP (Network Access Protection), 588–591  
   NAT64/DNS64 support, 13  
   NDF (Network Diagnostic Framework), 564  
 net commands  
   list of, 50–51  
   net help command, 23  
   net session command, 476–477  
   net share command, 461  
 .NET Framework, 37  
 NetBIOS, 20–21, 50, 618  
 Net-DMA (network direct memory access), 13  
 Netdom commands, 42, 252, 265–266, 279–280  
 NETLOGON share, 474  
 Netmon, 564  
 Netsh command, 42, 46, 51, 564, 591  
 Network Access Protection (NAP), 588–591  
 network adapters  
   adding, assessing need for, 99  
   DHCP server bindings, configuring, 584  
   discovery settings, 560–561  
   installing TCP/IP, 565–566  
   MAC addresses, 13–14  
   network connections of, 571–572  
   NIC Teaming, 61  
   performance issues, 134  
   utilization statistics, 98  
   virtual, for Hyper-V, 570–571  
 Network And Sharing Center  
   Advanced Sharing Settings, 460–461  
   alternate IP addressing, 568–569  
   DNS settings, 620–621  
   IP addressing, 567–569  
   multiple gateway configurations, 569–570  
   network connections, 571–572  
   Network Connections page, 562  
   settings, basic, 12  
   status indicators of, 561–562  
   troubleshooting options, 562  
 network authentication, 296–297  
 Network Awareness, 560  
 network connections, 571–572, 584  
 network device configuration, 6  
 Network Device Enrollment Service, 33  
 Network Diagnostic Framework (NDF), 564

- Network Diagnostics, 559, 564
- network direct memory access (Net-DMA), 13
- Network Discovery, 12, 560–562
- network drives, 484–485
- Network Explorer, 559–561
- Network File System. *See* NFS (Network File System)
- Network Load Balancing (NLB), 37, 63
- Network Location Awareness, 143
- Network Monitoring, 564
- Network Policy and Access Services (NPAS), 34, 588–591
- Network Policy Server, 588–591
- Network Unlock, 391–392
- networking
  - addresses, network, 13–14. *See also* IP addresses
  - administrative template policies for, 151
  - costed networks, 168
  - DirectAccess, 13
  - discovery, 12, 560–562
  - DNS64 support, 13
  - domain names. *See* DNS (Domain Name System)
  - dynamic address assignment. *See* DHCP (Dynamic Host Configuration Protocol)
  - Ethernet properties configuration, Server Manager, 60–61
  - Home network type, 12
  - Hyper-V with, 570–571
  - IPv4. *See* IPv4 addresses
  - IPv6. *See* IPv6 addresses
  - latency issues, 134
  - mobile device policies, 168
  - name servers, without, 22–23
  - NAT64/DNS64 support, 13
  - Netsh command, 42, 46, 51, 564, 591
  - Network identity, 318
  - performance tuning, 134
  - ping command, 567
  - private networks, 12, 626–627
  - profile management, 562
  - public networks, 12, 626–627
  - Resource Monitor usage statistics, 120
  - security features, list of, 563
  - security policy configuration, 209
  - sharing settings, 12
  - TCP Chimney Offload, 13
  - TCP/IP. *See* TCP/IP protocol
  - Teredo, 564
  - tool suite for, 11–12
  - tools for, list of, 559–560
  - types of networks, 12, 560
  - utilization statistics, 98
  - Work network type, 12
- NetworkService account, 307, 504
- new- cmdlets, 25
- new features of Windows Server 2012, 6–7
- Next Generation TCP/IP stack, 14
- NFS (Network File System)
  - Client for NFS, 36
  - file sharing with, 463, 479–481
  - role service for, 34
  - Server for NFS role service, 387
  - share profile options, 466
- NLB (Network Load Balancing), 37, 63
- No Access share permission, 469
- normal backups, 521, 525, 528–529
- normal scopes, 579, 593–598
- Notepad, 41
- NP (nonpaged) pool, 92, 98
- NPAS (Network Policy and Access Services), 34
- NPS (Network Policy Server), 358
- ntdsutil command, 251, 276
- NTFS disk quotas
  - enabling on per-volume basis, 508–510
  - Group Policy for, 505–508
  - importing and exporting, 512–513
  - limits, 504
  - managing, 510–513
  - vs. Resource Manager disk quotas, 503
- NTFS file system
  - compression, enabling, 410
  - converting volumes to, 445–447
  - encryption with, 413, 416
  - file sharing, advantages for, 463
  - permissions, 458, 488
  - quotas. *See* NTFS disk quotas
  - self-healing NTFS, 450–451
  - structure of volumes, 393
  - Transactional NTFS, 449–450
- NTLM authentication, 296
- NX (no-execute) processor feature, 80

**O**

## objects

- editing permissions for, 487
- group. *See* GPOs (Group Policy Objects)
- inheritance with, 487–488
- managers and management tools for, 485–486
- ownership of, 486–487
- types of, 485–486

OCsetup.exe, 42

OCSP (Online Certificate Status Protocol), 563

offline domain join, 268–270

Online Certificate Status Protocol (OCSP), 563

Online Responder, 33

operating system recovery. *See* system recovery

## operations masters

- assignment of, 243
- definition of, 231
- domain naming master role, 242, 244
- infrastructure master role. *See* infrastructure masters
- PDC emulator role, 243–244
- relative ID master role, 242, 244
- required roles for Active Directory, 242–243
- schema master role, 242, 244
- seizing roles, 276–280
- single domain controller scenario, 243
- standby operations masters, 243
- transferring roles, 274–276
- viewing, 273–274

Optimize Drives, 453–455

Organization panel, 349

organizational units. *See* OUs (organizational units)

OS X file sharing with NFS, 463, 479–481

## OUs (organizational units)

- advantages of creating, 227–228
- default domain policies, 147
- definition of, 138
- managing, 228, 281–282
- permissions to manage Group Policy, 155

ownership of objects, 486–487

**P**

packages for deploying software with Group Policy, 182

page faults, 92, 131

paged pool, 92, 98, 132

paging, 77–80, 402

parent domains, 18, 218, 615

parent objects, 487

parity checking, 426

## partitions

active, 402, 445

basic to dynamic disk conversions, 402

boot, 402, 433–434, 443

clean installations, options for, 44

compression, enabling, 410–412

crash dump, 402

creating, 54, 407–410

deleting, 52–55, 445

DiskPart utility, 49, 52–53

extended, 55, 407

file system format, selecting, 409, 411

formatting, 54, 407, 409–411, 425

GPT (GUID partition table), 392, 399, 403, 407

installations, changing during, 53–55

labeling, 409, 425

logical drives in, 392, 401, 404, 407

MBR (master boot record), 392, 399, 403, 406–407

mounting in empty NTFS folders, 408

number allowed per drive, 407

preparing, 392–394

primary, 407

quick formats, 410–411

resizing, 447–449

structure of, 406–407

system partitions, 402

passive cooling mode, 8–9

## passwords

Account Lockout policies, 325–326

Administrator accounts, for, 308

computer accounts, for, 264–266

managed service accounts, for, 343–345

new, creating for accounts, 332–333

options for controlling, 378

options for, setting at creation, 332–333

policies governing, 321–323

PSOs (password-setting objects), 222

resetting, 373

secure channel passwords, 264–266

secure password strategies, 320–321

- security account options for, 360–361
  - security templates for, 193–195
  - single sign-on, 296–297
  - user accounts, relationship to, 302
- payloads, 24, 57, 63–65, 72, 218
- PDC emulator role, 243–244, 273–274, 276–280
- Peer Name Resolution Protocol, 37
- peer-to-peer networks with LLMNR, 22–23
- performance
  - alerts, 60, 129–130
  - applications options, 77
  - built-in diagnostics for, 543
  - CPU usage statistics, 96–97
  - c-states, 10–11
  - data collector sets, 124–130
  - Data Execution Prevention options, 80–81
  - disk striping for I/O speed, 426–428
  - graphics for, 76–77, 95–96
  - memory, tuning, 130–132
  - monitoring, goals for, 118
  - monitoring tool. *See* Performance Monitor
  - network tuning, 134
  - paging, configuring, 77–80
  - power option effects on, 8–11
  - processors, tuning, 132–133
  - Profile privileges, 312
  - p-states, 10–11
  - Server Manager Performance panel, 60
  - Task Manager Performance tab, 95–99
  - virtual memory, configuring, 77–80
- Performance Monitor
  - alerts, 129–130
  - counters, 121, 122–124
  - credentials needed for, 123
  - Data Collector Sets node, 124–130
  - graphic display, 121
  - instances, 122–123
  - objects, performance, 123
  - opening, 121
  - purpose of, 119
  - Reports, data collector, 128–129
  - views, 122
- permissions
  - Active Directory, setting, 380–381
  - backups, for, 528
  - basic, 488–493
  - claims-based access controls, setting for, 496–498
  - editing for objects, 487
  - file and folder permissions, table of, 489
  - file sharing. *See* share permissions
  - folder, security templates for, 200–201
  - Group Policy management permissions, 154–155
  - inheritance of, 487–488, 495
  - logon rights, granted by, 312–313
  - network drives, 484
  - NFS, 479–481
  - privileges, relation to, 310
  - security tokens for, 305
  - setting for files and folders, 491–493, 493–495
  - special, 489–491, 493–495
  - system services, setting, 197
- physical drives. *See* hard disk drives
- physical processors, 6
- physical structures, domain, 223
- ping command, 51, 567
- planning with Modeling Wizard, 161–163
- Pnputil.exe, 42
- PNRP (Peer Name Resolution Protocol), 37
- policy editors. *See* GPMC (Group Policy Management Console)
- policy objects. *See* GPOs (Group Policy Objects)
- ports used by Active Directory, 292–293
- power options, 5–11
- PowerShell. *See* Windows PowerShell 3.0
- preboot environment, 5
- predefined accounts, 306
- preinstallation environment, 5
- Previous Versions, 482–483
- primary computers
  - for folder redirection, 173
  - for restricting roaming, 363
- primary groups, 340
- primary partitions, 407
- printing, 34, 51, 151, 475
- private networks, 12, 560, 626–627
- privileges
  - processes, viewing for, 92
  - table of, for users and groups, 310–312
  - user accounts, granting to, 309–310
- problems, viewing current, 542
- process working set privilege, 311
- processes
  - dependencies of, 93–94
  - End Task command, 89–90, 94
  - executables, finding related, 90
  - foreground vs. background, 87
  - memory reserved for, 90, 91

## processors

- processes, *continued*
  - PIDs of, 90, 93–95
  - properties of, columns for, 91–93
  - publishers of, 90
  - session IDs, 93
  - status, determining, 89, 91
  - thread counts, 93
  - trees of, 94
  - types of, 90
  - waiting for locked resources, 93
- processors
  - affinity settings, 10
  - counters for, 132–133
  - CPU usage statistics, 96–97
  - c-states, 10–11
  - logical processor idling, 10
  - logical, viewing number of, 96
  - multiple processor cores support, 6
  - NX (no-execute) processor feature, 80
  - performance tuning, 132–133
  - power states of, 8–11
  - processor architecture environment variable, 353
  - p-states, 10–11
  - scheduling options, 77
  - thread queuing issue, 132
  - throttling, 8–11
  - usage, monitoring, 133
- product identifiers, 61
- product keys, 45, 46, 47, 73–74
- profiles, user. *See* user profiles
- programs
  - Apps, 7, 137, 363
  - Debug Programs privilege, 311
  - DEP exceptions for, 80–81
  - deploying, 181–186
  - directory data stores specific to, 15
  - opening, options for, 7
- PROMPT command, 51
- properties
  - resource, 298–299
  - server, viewing in Server Manager, 60–62
- Proxy identity, 318
- PSOs (password-setting objects), 222
- p-states, 10–11
- Public folder, 458–460
- public networks, 12, 560–561, 626–627
- publishing
  - data stores, 238
  - process publishers, 90

- shares, 469
- software, 182
- trusted publisher list, 564

## Q

- Quality Windows Audio Video Experience, 37

## R

- RACTask, 122
- RAID
  - backup systems using, 523–524
  - deleting volumes, 426
  - disk mirroring, 426–434
  - disk striping, 426–428, 431, 434–435
  - Diskraid.exe, 42
  - level 0, 426–428, 434
  - level 1, 426–433
  - level 5, 426–427, 431, 434–435
  - parity checking, 426
  - queue length counter, 133
  - repairing, 432–435
  - supported levels of, 420
- RAS Connection Manager, 37
- Read Attributes special permission, 490–491
- Read permission, 489
- Read Permissions special permission, 490–491
- Read share permission, 470
- read-only domain controllers (RODCs), 16, 19, 219, 619
- reads, disk performance for, 133
- recovery. *See also* restoring
  - Active Directory, 245–248, 547
  - EFS based, 414–415, 417
  - encryption certificates, of, 551–553
  - improvements in Windows Server 2012 for, 541–542
  - nonsystem volumes with Windows Server Backup, 550–551
  - policy, configuring, 417
  - RECOVER command, 51
  - Recovery screen options, 544
  - safe mode startups, 544–546
  - services, configuring for, 106–107
  - Startup And Recovery dialog box, 83–85
  - system. *See* system recovery
  - system image recovery, 44, 544
  - tools for, installing, 526–527
  - Wbadm backup recovery commands, 533



- recovery point objective (RPO), 520–521
- recovery time objective (RTO), 520–521
- Recycle Bin, Active Directory, 222, 244–248
- redirected folders, 173–177, 363
- refreshing Group Policy, 158–161, 166
- ReFS (Resilient File System), 393, 449, 459, 504
- Regedit, 41
- register cache utilization, 96–97
- registry
  - administrative templates, 151–152
  - auditing, 502
  - data collectors for, 127–128
  - REG commands, list of, 51
  - security policy configuration, 209–210
  - security templates for, 198–200
- Regsvr32, 39
- relative ID master role
  - managing, 273–274
  - placement of, 244
  - purpose of, 242
  - seizing, 276–280
  - SID generation task, 302
  - transferring, 276
- Reliability Monitor, 119–122
- Remote Access service, 34, 358–360
- remote assistance, 6, 37, 136
- Remote Desktop
  - clients, viewing, 100
  - disconnecting, 100
  - managing, 136
  - RDP file signing, 564
  - Remote Desktop Services User identity, 318
  - remote management, independence of, 67
  - Services, 34, 61
  - Session Host, 4
  - Task Manager for managing, 99–100
  - utilization statistics, 100
- Remote Differential Compression, 37
- remote management
  - adding servers to Server Manager, 67–68
  - blocking for the local server, 66
  - credentials required for, 66–67
  - Disk Management capabilities, 394, 396
  - groups, 68–70
  - PowerShell for, 70
  - PowerShell Web Access for, 24
  - properties, 61
  - Remote Desktop independence of, 67
  - RSAT for, 37, 65, 135, 256
  - tasks available through, 65–67
  - tools available for, 65
  - web client WinRM IIS Extension, 26
  - Windows Firewall, applications enabled for, 66–67
  - WinRM, 25–29
- Remote Server Administration Tools (RSAT), 37, 65, 135, 256
- Remote Service Management, 67
- removable storage devices, 395–398, 443
- remove- cmdlets, 25
- Remove Computer From Docking Station privilege, 312
- Remove Roles And Features Wizard, 71, 72–73, 272–273
- Repadmin.exe, 252, 277, 293–294
- Repair Your Computer, 525
- Replace A Process Level Token privilege, 312
- replication
  - application directory partitions with, 238
  - bridgehead servers, 291–292, 293
  - cmdlets for viewing, 278–279
  - DFS Replication role service, 387
  - of directory data by domain controllers, 230–231
  - DNS configuration for, 625–626
  - domain controller multimaster model, 14
  - File Replication Service log, 109
  - of GPOs using DFS, 142
  - intersite replication topology, 291–292
  - listing data about, 293–294
  - recovering from failures of, 291
  - Repadmin.exe, 277–278, 293–294
  - Replication Diagnostics tool, 252
  - service dependencies of, 292
  - site links, 285–286
  - troubleshooting, 292–294
  - types of data replicated, 238–239, 241–242
  - USNs (update sequence numbers), 277, 293
- reports, data collector, 128–129
- rescanning all drives, 404
- Resource Manager disk quotas, 503–504, 514–518
- Resource Manager, Windows System, 10, 38
- Resource Monitor, 119–120

## resources

- resources
  - properties, 298–299
  - sessions, viewing, 476–478
- Restart Manager, 542
- Restartable Active Directory Domain Services, 16–17
- restarting servers, 7–8
- restoring. *See also* recovery
  - Active Directory, 244–248, 547
  - default Group Policy Objects, 172–173
  - DHCP configurations, 591–592
  - encryption certificates, 554–555
  - GPOs (Group Policy Objects), 165–166
  - mirrored drive sets, 432–433
  - nonsystem volumes, 550–551
  - Restore Files And Directories privilege, 312
  - security settings from GPOs, 213
  - security settings with rollback templates, 204
  - shadow copies, 482–483
  - system state, 546
- Restricted Groups settings, 195–196
- Restricted identity, 318
- Resultant Set of Policy (RSOP), 154, 166
- resume functionality, absence of, 5
- reverse lookups, 624, 628–629
- RID (Relative ID)
  - master role. *See* relative ID master role
  - soft ceilings and warnings, 223
- rights, 15, 327–330. *See also* permissions
- roaming
  - encryption certificates in profiles, 414, 552
  - roaming profiles, 362–364, 368, 552
- RODCs (read-only domain controllers), 16, 19, 219, 619
- role services
  - adding with ServerManager module, 63
  - Ocsetup.exe for configuring, 42
  - server roles, relation to, 32
- roles, server. *See* server roles
- root domains, 18
- RPCs (Remote Procedure Calls)
  - RPC over HTTP Proxy, 37
  - RPC over IP, site links with, 286
  - Windows Firewall with, 67
- RPO (recovery point objective), 520–521
- RSAT (Remote Server Administration Tools), 37, 65, 135, 256

- RSOP (Resultant Set of Policy), 154, 166, 171–172
- RSS (receive-side scaling), 13
- RTO (recovery time objective), 520–521

## S

- safe mode, 544–546
- SATA (Serial ATA), 390, 397
- scanning, 34
- scheduled tasks, 67
- schemas
  - Active Directory Recycle Bin, preparing for, 244–245
  - default, 239
  - functional levels required for feature support, 220–221
  - master role, 242, 244, 275–280
  - Schema Admins group, 315–317
- Sconfig (Server Configuration), 41, 46, 58
- scope, group, 336–337
- scopes of IP addresses
  - classes and types of, 578–579
  - exclusion ranges, 607–608
  - failover scopes, 602–605
  - icon indicators for, 582
  - managing, 601–602
  - normal scopes, 593–598
  - scope options, 599–601
  - statistics, viewing, 605
  - superscopes, 579, 592–593
- scripts, 178–179
- SCSI (Small Computer System Interface), 390
- Scwcmd command, 206, 212
- Search box, Start options panel, 7
- searching
  - accounts, listing, 368–369
  - Active Directory for users and groups, 350–351
  - for Active Directory objects, 254–256, 258
  - Apps, installed, 7
  - Search box focus, 137
  - Windows TIFF IFilter, 39
- Secure Socket Tunneling Protocol (SSTP), 563
- Secure Sockets Layer (SSL), 563
- Secured Boot, 391–392
- security
  - account options, 360–361
  - Administrator accounts, steps to secure, 307–308

- Configuration Wizard. *See* Security Configuration Wizard
- Data Execution Prevention, 80–81
- Default Domain Policy GPO, 147
- DNS issues, 626–627, 632–634
- identifiers. *See* SIDs (security identifiers)
- networking features for, 563
- permissions. *See* permissions policies for. *See* security policies
- refreshing Group Policy, 158–161, 166, 168–169
- Security Configuration And Analysis snap-in, 192–193, 201–204
- Security log, 108
- service accounts, considerations, 105
- services, disabling unnecessary, 107
- slow-link detection effects, 168–169
- templates for. *See* security templates
- tokens, 305
- user verification. *See* authentication
- Security Configuration And Analysis, 192–193, 201–204
- Security Configuration Wizard
  - capabilities of, 206
  - configuration sections, 206–207
  - network configuration, 209
  - registry settings configuration, 209–210
  - security policies, managing, 207, 210–213
  - server roles, services, and features, 208–209
  - templates, adding, 210
  - viewing security configuration data-bases, 207–208
- security descriptors, 297, 303
- security groups, 303, 305
- security log, auditing, 498–503
- security policies
  - applying, 211
  - audit policy configuration, 210
  - creating, 207
  - editing, 211
  - folders for, default, 210
  - multiple computers, deploying to, 212
  - network configuration, 209
  - purpose of, 206
  - registry settings configuration, 209–210
  - rolling back, 211–213
  - Save Security Policy options, 210
  - Scwcmd command, 206, 212
  - security templates, adding, 210
  - server roles configuration, 208
  - services and features, configuring, 208–209
  - viewing security configuration data-bases, 207–208
- security prompts, 5, 6, 93
- security templates
  - applicability of, 191
  - configuring, 201–204
  - creating, 192, 193
  - file path settings, 198–201
  - folders for, 193
  - multiple computer deployments, 204–205
  - policy settings, 193–195
  - registry settings, 198–200
  - Restricted Groups settings, 195–196
  - rollback templates, 203–204
  - Security Configuration And Analysis, 192–193, 201–204
  - security policies, incorporating in, 206, 210
  - steps for using, general, 192
  - system services policy settings, 196–197
- security zones, 61
- seizing server roles, 276–280
- Self identity, 318
- self-healing NTFS, 450–451
- Server Core installations
  - backup tool options, 526
  - command prompts, opening new, 41
  - converting other installation types to or from, 56–57
  - DHCP default, 46
  - features that can be installed, 42
  - limited functionality of, 4
  - roles supported by, 40
  - Sconfig with, 41–42, 58
  - setup commands, 46–47
  - user interface of, 40–41
  - Windows Logon, 40–41
- Server Graphical Shell, 4, 56
- Server Manager
  - Active Directory integration, 223
  - Add Other Servers To Manage, 58
  - Add Roles And Features Wizard, 24, 58, 70–72, 218
  - adding servers to, 67–68
  - administrative wizard access from, 23

Server Manager, *continued*

- All Servers view, 68
- alternate credentials for servers, 58
- alternate credentials for servers, entering, 58
- Best Practices Analyzer, 60
- capabilities of, 32, 57–58
- command-line version, 63–64
- console tree options, 59
- Create Server Group, 58
- default view, 58–59
- demoting domain controllers, 272–273
- dependencies, notifications of, 32
- DHCP console, 581–582
- Disks node, 437–438
- Events panel, 60, 109–115
- File And Storage Services node, 436–439
- file shares, viewing, 462–463
- Group Policy tool. *See* GPMC (Group Policy Management Console)
- grouping servers, 58
- initial configuration with, 58–62
- Local Server properties, 59
- opening, 62
- Performance panel, 60
- permissions, setting, 493, 495
- Properties panel, 60–62
- remote management requirements, 65–66
- Remove Roles And Features Wizard, 71–73
- role-based group management, 70–73
- Roles And Features panel, 60
- Services panel, 60, 100–101. *See also* services
- share permission configuration, 472–474
- shared folder management, 463–469
- startup options, 58
- stopping sharing on folders, 478–479
- storage pool creation, 438–440
- task capabilities of, 31
- Volumes node, 436–437

server roles

- Add Roles And Features Wizard, 58, 70–72
- binaries for, 72
- definition of, 32
- hardware requirement considerations, 32
- managing with Server Manager, 70–73
- managing with ServerManager module, 63

- Ocsetup.exe for configuring, 42
- Remove Roles and Features Wizard, 71–73
- Roles And Features panel, Server Manager, 60
- security policy configuration for, 208
- table of available, 33–35

Server With A GUI installations, 4. *See also* full-server installations

Server With Minimal Interface installations, 4, 40, 56

ServerManager module for PowerShell, 63–64

Service identity, 318

services. *See also specific services*

- Computer Management, configuration with, 102–107
- disk quotas with accounts, 504
- group organization of, 101
- logon accounts of, viewing, 102
- logon configuration, 104–105
- managed accounts for, 341–344
- managing, 103, 107
- names of, 101–102
- recovery configuration, 106–107
- Remote Service Management, 67
- restrictions, running under, 94–95
- security considerations for, 105
- security policy configuration for, 208–209
- Server Manager Services panel, 60, 100–101
- start types of, viewing, 101
- startup configuration, 103–104
- status of, viewing, 101–102
- stopping, 95, 103
- system, 94–95, 196–197
- Task Manager Services tab, 94–95

session IDs, 93, 99

sessions

- managing, 476–478
- type, viewing, 99

set- cmdlets, 25–26

SET commands, 52

Setup

- command-line commands, table of, 48–52
- installation steps, 44–48
- log, 108

shadow copies, 481–484, 520

- share permissions
  - configuring in Computer Management, 470–472
  - configuring in Server Manager, 472–474
  - file sharing role of, 458
  - list of, 469–470
  - options for, 465, 467–468
- shared folders
  - administrative template policies for, 151
  - claims-based permissions for, 498
  - public folder sharing, 458
  - removable disks with, 398
- sharing
  - administrative shares, 474–476
  - files. *See* file sharing
  - hidden shares, 464, 474–476
  - network options for, 12
  - permissions for. *See* share permissions
  - public folder sharing, 458
  - shadow copies of shared folders, 481–484
  - special shares, 474–476
  - standard file model, 457–458
- shutting down
  - methods for, 7–8
  - scripts for, 178–179
  - Shut Down The System privilege, 312
  - shutdown command, 47
- Sidebar, Windows, 5
- SIDs (security identifiers)
  - group accounts with, 304–305
  - NTFS disk quotas, use with, 506
  - renaming user accounts, 369–371
  - structure of, 302
- Simple Network Management Protocol (SNMP), 38
- simple volumes, 420. *See also* volumes
- single-label name resolution, 619
- single sign-on, 296–297
- sites, Active Directory
  - Active Directory Sites And Services for managing, 229–230
  - bridgehead servers, 291–293
  - creating, 282–283
  - domain controllers, associating with, 285
  - GPMC Sites node, 149
  - intersite replication topology, 291–292
  - links between, configuring, 285–289
  - permissions to manage Group Policy, 155
  - structural relations of, 138, 229
  - subnets, associating with, 284
  - well connected ideal for, 229
- 64-bit systems, 6, 43–44
- Sleep state, server, 5
- sleep states, processor, 10–11
- Slmgr commands, 42
- slow-link detection, 159, 167–170
- smart cards, 302, 320, 360
- SMB (server message block)
  - advantages for file sharing, 459
  - encryption of shares option, 468
  - enhancements in version 3.0, 457
  - get-smbshare cmdlet, 461
  - port used by, 293
  - security policy configuration for, 209
  - share profile options, 466
- SMTP (Simple Mail Transfer Protocol), 37, 286
- Snipping Tool, 5–6
- SNMP (Simple Network Management Protocol), 38
- SOA (start of authority) records, 637, 641–642
- Sound Recorder, 5–6
- spanned volumes
  - managing, 424–426, 447–449
  - status issues, table of, 422–423
  - vs. simple volumes, 420
- special folders, redirecting, 173–177
- special identities, 306
- special permissions, 489–491
- special share symbol (\$), 474–475
- SQL Server, 341
- SSL (Secure Sockets Layer), 563
- SSTP (Secure Socket Tunneling Protocol), 563
- standalone servers, 14
- standard file sharing, 457–468
- standard volumes, creating, 441–443
- standards-based storage
  - abstraction, 435
  - deduplication, 436
  - Disks node options, Server Manager, 437–438
  - layers of, 436
  - shares, creating, 437
  - storage pools, 435, 438–440
  - storage spaces, 435
  - subsystems, 436, 439

## Start screen

- standards-based storage, *continued*
  - traditional storage compared to, 419–420
  - virtual disk creation in storage spaces, 440–441
  - volume management, 436, 437, 441–443
- Windows Standards-Based Storage Management, 38, 436
- Start screen, 7, 151
- starter GPOs, 153
- startup
  - debugging mode, 545
  - DHCP server mechanism for, 574
  - restore options from, 545
  - safe mode, 544–546
  - scripts for, 178–179
  - settings options, Recovery screen, 544
  - Startup And Recovery dialog box, 83–85
  - Startup Recovery Options, 548
  - Startup Repair (StR), 543, 545, 548
  - Windows Preboot Environment, 5
- static IP addresses, 566–568
- Stop errors, 84–85
- storage. *See also* hard disk drives
  - attached, 419
  - BitLocker Drive Encryption, 35, 136, 404
  - Enhanced Storage, 36
  - I/O performance tuning, 133
  - iSNS Server Service, 36
  - PhysicalDisk counters, 133
  - removable media, 443
  - removable storage devices, 396–398
  - Resource Monitor Disk usage statistics, 120
  - standards-based techniques. *See* standards-based storage
  - Storage Services role service, 387
  - subsystems, 436, 439
  - traditional vs. standards-based techniques, 419
  - Windows Standards-Based Storage Management, 38
- storage pools
  - creating, 438–440
  - physical disks, handling by, 439
  - primordial pools, 439
  - standards-based storage, role in, 435
  - virtual disk creation in, 440–441
- Storage Services, 34
- StR (Startup Repair), 543, 545, 548
- striped volumes, 424–428, 431, 434–435
- SUA (Subsystem for UNIX-Based Applications), 38
- subdomains, 18
- subnet masks, 566–567, 574
- subnets
  - creating and associating with sites, 284
  - deleting from DNS servers, 636
  - name resolution with LLMNR, 22–23
  - place in overall domain structures, 223–224
  - reverse lookup zones for, 628–629
  - sites, relation to, 138
- Subsystem for UNIX-Based Applications, 38
- superscopes, 579, 592–593
- Support Dynamic Access Control And Kerberos Armoring policy, 496–497
- Svchost.exe, 95
- Sync Center, 5
- system account disk quotas, 504
- System console, 73, 75–85
- system environment variables
  - common, list of, 352–353
  - configuring, 81–83
- system files
  - backing up, 528, 534
  - Startup Repair, 543, 545
- System identity, 318
- System Idle Process, 93
- System Image Recovery feature, 44
- System log, 108, 143
- system partitions, 402
- System Properties dialog box, 75–85
- system recovery
  - full system recovery issues, 548
  - restoring system state, 546
  - Server Core recovery, 557
  - tools for, 547–549
- system services, 94–95, 196–197
- system settings, template, 151
- System utility
  - console, 73, 75–85
  - local profile management, 365–368
  - task capabilities of, 31
- system volume drive letters, 443
- SystemInfo command, 43
- SYSVOL share, 475

## T

- Take Ownership Of Files Or Other Objects privilege, 312
- Take Ownership special permission, 490–491
- tape drives, 523–525
- Task Manager
  - Details tab, 91–94
  - End Task command, 89–90, 94
  - options for, 88, 90
  - Performance tab, 95–99
  - Processes tab, 88–94
  - Remote Desktop, managing, 99–100
  - Services tab, 94–95
  - Status column, processes, 89, 91
  - Users tab, 99–100
- Task Scheduler, 136, 539–540
- taskbars, template policies for, 151
- TCP Chimney Offload, 13, 564
- TCP/IP protocol
  - addresses for. *See* IP addresses
  - dual layer architecture of, 12–14
  - dynamic configuration of. *See* DHCP (Dynamic Host Configuration Protocol)
  - Group Policy settings for, 559
  - installing, 565–566
  - name resolution for. *See* DNS (Domain Name System)
  - Simple TCP/IP Services, 37
- telephone numbers, user contact, 348–350
- telnet, 38
- Teredo, 564
- themes, desktop, 5
- 32-bit processes, 92–93
- threads, 93, 97, 132
- throttling, processor, 8–11
- tickets, Kerberos, 327
- time
  - TIME command, 52
  - Time Zone property, 62
  - Windows Time feature, 137
- TLDs (top-level domains), 218
- tokens, 305
- trace data, collecting, 125, 127
- Transactional NTFS, 449–450
- Transport Server role service, 35
- Traverse Folder/Execute File special permission, 490–491
- trees, domain, 225–227
- troubleshooting

- Active Directory, 292–294
- DNS servers, 649–650
- improvements in recovery mechanisms, 541–542
- logon problems, 378–380
- problems, viewing current, 542
- safe mode for, 544–546
- Startup Repair Wizard, 548
- trust
  - Rights Management Services, 15
  - transitive trust relationships, 230
- tuning performance. *See* performance types of installations
  - changing, 55–58
  - choosing, 45, 48
  - list of, 4

## U

- UAC (User Account Control), 5, 6, 93
- universal groups, 231, 239–241, 281, 304–306
- UNIX, 38, 463, 479–481
- unmounted drives, 443
- Up Time system property, 97
- updates
  - applications, for, 541–542
  - using Group Policy, 185–186
  - Windows Server Update Services, 35
  - Wusa.exe command for, 43
- upgrade installations, 47–48
- upgrading software using Group Policy, 185–186
- USB interface, 396–397
- User Account Control (UAC), 5–6
- user accounts
  - built-in, 307
  - capabilities, types of, 309–310
  - contact information, setting, 347–350
  - credentials, preventing from delegating, 361
  - deletion effects, 302
  - disabled status, 360, 378–380
  - domain. *See* domain user accounts
  - enabling disabled accounts, 374
  - environment variables for, 352–353
  - global user rights configuration, 328–329
  - group accounts compared to, 300–301

- group membership, managing, 339–340, 375
- home folders, setting, 351, 371
- local. *See* local user accounts
- local user rights configuration, 330
- locked out accounts, 373–374, 379
- logon names identified with, 301–302
- logon rights for, 312–313
- logon script name, renaming, 371
- logon script paths, setting, 351
- managing, 373, 375–380
- naming policies, 319–320. *See also* user names
- passwords for, 302
- permissions, setting, 380–381
- predefined, 307–308
- privileges of, 309–312
- profile paths, setting, 351, 371
- properties, configuring, 375–378
- public certificates for, 302
- renaming, 369–371
- rights management for, 327–330
- security options, 360–361
- SIDs (security identifiers), 302
- smart cards for, 302
- tools for creating, 318
- types of, 301
- updating, 368–369
- user claims, 496–497
- user environment variables, 81–83
- User Interfaces And Infrastructure, 38–39
- user names
  - changing, role of SIDs, 302
  - environment variable for, 352
  - logon names, construction from, 301–302
  - naming policies, 319–320
- user objects, 301
- user profiles
  - app deployment issues, 363
  - caching of, 362
  - EFS with roaming profiles, 362
  - encryption certificates in, 414
  - local profiles, 362–368
  - managing, 361, 363–368, 373
  - mandatory profiles, 362, 364, 379–380
  - multiple accounts, setting for, 376–377
  - names of, 365
  - paths for, setting, 351
  - primary computers with, 363
  - roaming profiles, 362–364

- user publishing method for software deployment, 182
- %UserName% variable, 352–353, 355, 363–364, 376–377
- USNs (update sequence numbers), 277, 293
- USNs folder binaries, 64–65
- UUIDs (universally unique identifiers), 262

## V

- VHD (virtual hard disk), 395, 405–406, 526
- video
  - Quality Windows Audio Video Experience, 37
  - Video for Windows, 5
- virtual accounts, 346
- virtual disks
  - creating in storage pools, 439–441
  - iSCSI, 436
  - layout options, 439–440
  - place in standards-based storage, 435–436
  - roles, adding or removing, 70–73
  - Server Manager display of, 437
  - standard volumes, creating, 441–443
  - storage pool creation, 438–440
- virtual domains, 223
- virtual memory
  - Committed Bytes counter, 131
  - Committed statistic, 97
  - configuring, 77–80
  - Resource Exhaustion Detection And Recovery, 542
- Visual Effects options, 76–77
- Volume Activation Services, 34
- volume sets. *See also* spanned volumes
  - advantages of, 422
  - creating, 424–425
  - definition of, 420
  - deleting, 426
  - drive letter assignment, 425
  - sizing segments by disk, 424–425
  - status issues, table of, 422–423
- Volume Shadow Copy Service (VSS), 525
- volumes
  - active, 402
  - allocation unit size, setting, 425
  - boot, 402, 433–434, 443
  - capabilities of, 421
  - CHKNTFS command, 49
  - compression, enabling, 410–412



- converting to NTFS, 445–447
  - creating, 407–410, 424–425, 437
  - deleting, 426, 436
  - drive letter assignment, 408, 425
  - drive paths for, 407
  - extending, 436, 447–449
  - Failed status, 422
  - File System property of, 421
  - Free Space property of, 421
  - Healthy status, 423
  - labeling, 409, 411, 425, 444–445
  - Layout property of, 420
  - logical drives as, 407–408
  - mirrored, 426–434
  - MOUNTVOL command, 50
  - NTFS disk quotas, enabling on, 508–510
  - number of active, limitation of, 407
  - paging, configuring, 77–80, 443
  - Perform Volume Maintenance Tasks
    - privilege, 312
    - remote management of, 67
    - resizing, 447–449
    - restoring nonsystem, 550–551
    - reverting entire to shadow copies, 483
    - simple, 420
    - size of, specifying, 408
    - standard, creating, 441–443
    - standards-based. *See* standards-based storage
    - status issues, 421–423
    - striped, 424–428, 431, 434–435
    - system volumes, 402, 443
    - Type property of, 420
    - VOL command, 52
    - Volumes node, Server Manager, 436–437
  - VPNs (Virtual Private Networks), 34, 358–360, 563
  - VSS (Volume Shadow Copy Service), 525, 535
- W**
- wait chains, viewing, 93
  - WANs (wide area networks), 224, 565
  - Wbadm backup command
    - commands available in, 530–533, 537–539
    - compared to other backup utilities, 525
    - critical volumes only option, 534
    - deleting system state backups, 532
    - disabling daily backups, 532
    - enabling daily backups, 532
    - help, 530
    - modifying scheduled backups, 538
    - parameters, 531
    - running, 529
    - scheduling automatic backups, 538–540
    - storage location specification, 534–535
    - system state, 546
    - Task Scheduler with, 539–540
    - volume options for, 534
  - WDS (Windows Deployment Services), 4, 35, 262–263
  - web applications, internal, 15
  - Web edition, 6
  - web page, setting for users, 348, 350
  - web servers
    - domain controller, unable to run as, 6
    - Web Server (IIS) role, 35
  - Web Services, Active Directory, 221
  - Wecutil command, 43
  - WER (Windows Error Reporting), 42, 62
  - Wevutil command, 43
  - Wi-Fi. *See* wireless networks
  - WIM (Windows Imaging Format), 5, 57, 64–65
  - Windows 8, 3, 5, 65, 67
  - Windows Aero, 5
  - Windows Biometric Framework, 38
  - Windows Boot Manager
    - configuring, 83–84
    - resizing possible with, 447
  - Windows Defender, 5
  - Windows Deployment Services, 4, 35, 262–263
  - Windows Domain Manager, 252
  - Windows Error Recovery mode, 544
  - Windows Error Reporting (WER), 42, 62
  - Windows Event Log service, 108–109
  - Windows File Protection, 39
  - Windows Firewall
    - accessing, 137
    - Active Directory issues from, 250
    - exceptions, 66–67
    - graphical shell requirement, 4
    - Properties panel, 62
  - Windows Gadgets, 5
  - Windows Imaging Format (WIM), 5
  - Windows Internal Database, 38

## Windows Internet Name Service

- Windows Internet Name Service (WINS), 20–21
- Windows key, 7
- Windows Logon, 40–42
- Windows logs, 108
- Windows Management Instrumentation (WMI), 66
- Windows Media Foundation, 36
- Windows Media Player, 5
- Windows Memory Diagnostics, 543
- Windows NT PDC emulation, 243
- Windows PE (Preinstallation Environment) 4.0, 5
- Windows PowerShell 3.0
  - Active Directory module for, 220–221, 259
  - aliases, 25
  - backward compatibility of, 24
  - cmdlets, 24–25
  - disabling remote management for local server, 66
  - event log for, 109
  - execution order issues, 25
  - features, 24
  - listing cmdlets, 25
  - remote management, 25–29, 65–66, 70
  - script support, 178
  - Search box command execution, 137
  - Server Manager function, 63–64
  - starting, 24
- Windows Preboot Environment, 5
- Windows Process Activation Service, 38
- Windows Remote Management (WinRM).
  - See WinRM (Windows Remote Management)
- Windows Server Backup
  - advantages of, 525–526
  - application data, 528, 534
  - command line option to. See Wbadmin backup command
  - compared to other backup utilities, 525
  - configuring, 527–529
  - critical volumes only option, 534
  - destination type options, 537
  - excluding selected locations or file types, 535
  - full system recovery issues, 548
  - installing, 386, 525–527
  - manual backups, 540–541
  - Microsoft Exchange Server with, 527
  - modifying scheduled backups, 538
  - permissions for, 528
  - recovering nonsystem volumes, 550–551
  - remote shared folders for, 534, 539, 541
  - scheduling automatic backups, 535–538
  - starting, 527
  - stopping scheduled backups, 538
  - storage location specification, 534–535
  - system state data, 528, 534
  - volume options for, 534
  - VSS Settings, 535
- Windows Server Migration tools, 47
- Windows Server Update Services (WSUS), 189–190
- Windows Sidebar, 5
- Windows Software Licensing Management tool, 42, 74
- Windows Standards-Based Storage Management, 38, 436
- Windows System Resource Manager, 10, 38
- Windows TIFF IFilter, 39
- Windows Time, 137
- Windows Token-Based Agent, 33
- Windows Update
  - managing with Group Policy, 187–190
  - payloads, Group Policy for, 64–65
  - Properties panel, 62
  - restoring payloads with, 218
- WinRM (Windows Remote Management)
  - authentication issues, 26–27, 28–29
  - configuring, 26–28
  - disabling for local server, 66
  - Group Policy affecting, 29
  - IIS Extension, 39
  - listeners, 28
  - requirements for, 25–26
  - web gateway with, 26
  - Windows 8, enabling remote management of, 67
  - Windows Firewall exceptions for, default, 66
- WINS (Windows Internet Name Service), 20–21
- wireless networks
  - policies for, 168, 562–563
  - Wireless LAN Service, 39, 135
- Wmic commands, 43
- Work network type, 12
- workgroups
  - Active Directory Lightweight Directory Services, 15
  - Administrator accounts for, 308

- CAs (certificate authorities), 15
  - definition of, 6
  - joining computers to, 75, 267–270
  - remote management within, 65
  - time synchronization, 137

- WoW64, 39

- Write Attributes special permission, 490–491

- Write Extended Attributes special permission, 490–491

- Write permission, 489

- writes, disk performance for, 133

- WSH (Windows Script Host), 178

- WS-Management, 43

- WSRM (Windows System Resource Manager), 10, 38

- WSUS (Windows Server Update Services), 35, 189–190

## X

- X.500 directory service migrations, 301
- XPS Viewer, 39

## Z

- ZAW (.zap) files, 182–183
- zones, DNS

- Active Directory integration, 616–617

- adding records to, 636–640

- child domains, creating in, 634–636

- configuring new, 625–628

- DNSSEC configuration, 20, 632–634

- integration modes, setting, 645

- properties, setting, 641–645

- secondary servers, notifications to, 644

- SOA (start of authority) records, 637, 641–642

- transfer restrictions, 643–644

- types, setting, 645



# About the Author

---



**WILLIAM R. STANEK** (<http://www.williamstanek.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. His current and forthcoming books include *Windows 8 Administration Pocket Consultant* and *Windows Server 2012 Inside Out*.

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 crew member 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!

Find William on Twitter at WilliamStanek and on Facebook at [www.facebook.com/William.Stanek.Author](http://www.facebook.com/William.Stanek.Author).

# What do you think of this book?

We want to hear from you!

To participate in a brief online survey, please visit:

[microsoft.com/learning/booksurvey](https://microsoft.com/learning/booksurvey)

Tell us how well this book meets your needs—what works effectively, and what we can do better. Your feedback will help us continually improve our books and learning resources for you.

Thank you in advance for your input!

**Microsoft**  
Press