

Microsoft®

Updated for R2

Windows Server® 2008

2
SECOND
EDITION



William R. Stanek
Author and Series Editor

Administrator's Pocket Consultant

Sample Chapters

Copyright © 2010 by William R. Stanek

All rights reserved.

To learn more about this book visit Microsoft Learning at:

<http://www.microsoft.com/learning/en/us/Book.aspx?ID=13931>

Contents

	<i>Acknowledgments</i>	<i>xix</i>
	<i>Introduction</i>	<i>xxi</i>
Chapter 1	Windows Server 2008 R2 Administration Overview	1
	Windows Server 2008 R2 and Windows 7	2
	Getting to Know Windows Server 2008 R2	3
	Power Management Options	7
	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	17
	Name-Resolution Services	18
	Using Domain Name System	18
	Using Windows Internet Name Service	21
	Using Link-Local Multicast Name Resolution	23
	Frequently Used Tools	24
	Windows PowerShell 2.0	25
	Windows Remote Management	26
Chapter 2	Deploying Windows Server 2008 R2	31
	Server Roles, Role Services, and Features for Windows Server 2008 R2	32
	Full-Server and Core-Server Installations of Windows Server 2008 R2	39
	Installing Windows Server 2008 R2	43
	Performing a Clean Installation	44
	Performing an Upgrade Installation	47
	Performing Additional Administration Tasks During Installation	49

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

	Managing Roles, Role Services, and Features	56
	Viewing Configured Roles and Role Services	57
	Adding or Removing Roles on Servers	58
	Viewing and Modifying Role Services on Servers	60
	Adding or Removing Features in Windows Server 2008 R2	61
	Deploying Virtualized Servers	62
Chapter 3	Managing Servers Running Windows Server 2008 R2	67
	Performing Initial Configuration Tasks	68
	Managing Your Servers	70
	Managing Your Servers Remotely	74
	Managing System Properties	76
	The Computer Name Tab	77
	The Hardware Tab	78
	The Advanced Tab	79
	The Remote Tab	88
	Managing Dynamic-Link Libraries	88
Chapter 4	Monitoring Processes, Services, and Events	89
	Managing Applications, Processes, and Performance	89
	Task Manager	90
	Managing Applications	90
	Administering Processes	91
	Viewing System Services	94
	Viewing and Managing System Performance	95
	Viewing and Managing Networking Performance	97
	Viewing and Managing Remote User Sessions	98
	Managing System Services	99
	Starting, Stopping, and Pausing Services	101
	Configuring Service Startup	102
	Configuring Service Logon	103
	Configuring Service Recovery	104
	Disabling Unnecessary Services	106
	Event Logging and Viewing	107
	Accessing and Using the Event Logs	108
	Filtering Event Logs	110
	Setting Event Log Options	112
	Clearing Event Logs	114
	Archiving Event Logs	114

Monitoring Server Performance and Activity	116
Why Monitor Your Server?	116
Getting Ready to Monitor	116
Using the Monitoring Consoles	117
Choosing Counters to Monitor	120
Performance Logging	122
Viewing Data Collector Reports	126
Configuring Performance Counter Alerts	127
Tuning System Performance	128
Monitoring and Tuning Memory Usage	128
Monitoring and Tuning Processor Usage	130
Monitoring and Tuning Disk I/O	131
Monitoring and Tuning Network Bandwidth and Connectivity	132
Chapter 5 Automating Administrative Tasks, Policies, and Procedures	133
Understanding Group Policies	136
Group Policy Essentials	136
In What Order Are Multiple Policies Applied?	137
When Are Group Policies Applied?	138
Group Policy Requirements and Version Compatibility	139
Navigating Group Policy Changes	139
Managing Local Group Policies	142
Local Group Policy Objects	142
Accessing the Top-Level Local Policy Settings	143
Local Group Policy Object Settings	144
Accessing Administrator, Non-Administrator, and User-Specific Local Group Policy	144
Managing Site, Domain, and Organizational Unit Policies	145
Understanding Domain and Default Policies	146
Using the Group Policy Management Console	147
Getting to Know the Policy Editor	148
Using Administrative Templates to Set Policies	149
Creating and Linking GPOs	151
Creating and Using Starter GPOs	152
Delegating Privileges for Group Policy Management	152
Blocking, Overriding, and Disabling Policies	154
Maintaining and Troubleshooting Group Policy	157
Refreshing Group Policy	157

Configuring the Refresh Interval	158
Modeling Group Policy for Planning Purposes	160
Copying, Pasting, and Importing Policy Objects	163
Backing Up and Restoring Policy Objects	164
Determining Current Group Policy Settings and Refresh Status	165
Disabling an Unused Part of Group Policy	165
Changing Policy Processing Preferences	166
Configuring Slow-Link Detection	166
Removing Links and Deleting GPOs	169
Troubleshooting Group Policy	170
Fixing Default Group Policy Objects	171
Managing Users and Computers with Group Policy	172
Centrally Managing Special Folders	172
User and Computer Script Management	176
Deploying Software Through Group Policy	180
Automatically Enrolling Computer and User Certificates	185
Managing Automatic Updates in Group Policy	186
Chapter 6 Enhancing Computer Security	189
Using Security Templates	189
Using the Security Templates and Security Configuration And Analysis Snap-Ins	191
Reviewing and Changing Template Settings	192
Analyzing, Reviewing, and Applying Security Templates	199
Deploying Security Templates to Multiple Computers	202
Using the Security Configuration Wizard	203
Creating Security Policies	204
Editing Security Policies	208
Applying Security Policies	209
Rolling Back the Last-Applied Security Policy	209
Deploying a Security Policy to Multiple Computers	210
Chapter 7 Using Active Directory	211
Introducing Active Directory	211
Active Directory and DNS	211
Read-Only Domain Controller Deployment	212
New Active Directory Features	213
Working with Domain Structures	215
Understanding Domains	215

Understanding Domain Forests and Domain Trees	216
Understanding Organizational Units	218
Understanding Sites and Subnets	220
Working with Active Directory Domains	221
Using Windows 2000 and Later Computers with Active Directory	221
Working with Domain Functional Levels	222
Raising Domain and Forest Functionality	225
Understanding the Directory Structure	227
Exploring the Data Store	227
Exploring Global Catalogs	228
Universal Group Membership Caching	229
Replication and Active Directory	230
Active Directory and LDAP	231
Understanding Operations Master Roles	232
Using the Active Directory Recycle Bin	233
Preparing Schema for the Recycle Bin	234
Recovering Deleted Objects	234
Chapter 8 Core Active Directory Administration	237
Tools for Managing Active Directory	237
Active Directory Administration Tools	238
Active Directory Command-Line Tools	238
Active Directory Support Tools	239
Active Directory Administrative Center and Windows PowerShell	240
Using Active Directory Users And Computers	242
Getting Started with Active Directory Users And Computers	242
Connecting to a Domain Controller	243
Connecting to a Domain	244
Searching for Accounts and Shared Resources	244
Managing Computer Accounts	246
Creating Computer Accounts on a Workstation or Server	246
Creating Computer Accounts in Active Directory Users And Computers	247
Viewing and Editing Computer Account Properties	248
Deleting, Disabling, and Enabling Computer Accounts	248
Resetting Locked Computer Accounts	249
Moving Computer Accounts	250

Managing Computers	250
Joining a Computer to a Domain or Workgroup	250
Using Offline Domain Join	252
Managing Domain Controllers, Roles, and Catalogs	253
Installing and Demoting Domain Controllers	253
Viewing and Transferring Domainwide Roles	255
Viewing and Transferring the Domain Naming Master Role	257
Viewing and Transferring Schema Master Roles	257
Transferring Roles Using the Command Line	258
Seizing Roles Using the Command Line	258
Configuring Global Catalogs	260
Configuring Universal Group Membership Caching	261
Managing Organizational Units	261
Creating Organizational Units	261
Viewing and Editing Organizational Unit Properties	262
Renaming and Deleting Organizational Units	262
Moving Organizational Units	262
Managing Sites	262
Creating Sites	263
Creating Subnets	264
Associating Domain Controllers with Sites	265
Configuring Site Links	265
Configuring Site Link Bridges	268
Maintaining Active Directory	269
Using ADSI Edit	269
Examining Intersite Topology	271
Troubleshooting Active Directory	272
Chapter 9 Understanding User and Group Accounts	275
The Windows Server Security Model	275
Authentication Protocols	276
Access Controls	277
Differences Between User and Group Accounts	277
User Accounts	278
Group Accounts	279
Default User Accounts and Groups	283
Built-in User Accounts	283
Predefined User Accounts	284

Built-in and Predefined Groups	285
Implicit Groups and Special Identities	286
Account Capabilities	286
Privileges	287
Logon Rights	290
Built-in Capabilities for Groups in Active Directory	291
Using Default Group Accounts	293
Groups Used by Administrators	293
Implicit Groups and Identities	294
Chapter 10 Creating User and Group Accounts	297
User Account Setup and Organization	297
Account Naming Policies	298
Password and Account Policies	299
Configuring Account Policies	302
Configuring Password Policies	302
Configuring Account Lockout Policies	304
Configuring Kerberos Policies	305
Configuring User Rights Policies	307
Configuring Global User Rights	307
Configuring Local User Rights	309
Adding a User Account	310
Creating Domain User Accounts	310
Creating Local User Accounts	312
Adding a Group Account	313
Creating a Global Group	314
Creating a Local Group and Assigning Members	315
Handling Global Group Membership	316
Managing Individual Membership	316
Managing Multiple Memberships in a Group	317
Setting the Primary Group for Users and Computers	317
Implementing Managed Accounts	318
Creating and Using Managed Service Accounts	319
Configuring Services to Use Managed Service Accounts	320
Removing Managed Service Accounts	321
Moving Managed Service Accounts	321
Using Virtual Accounts	322

Chapter 11	Managing User and Group Accounts	323
	Managing User Contact Information	323
	Setting Contact Information	323
	Searching for Users and Groups in Active Directory	325
	Configuring the User's Environment Settings	326
	System Environment Variables	327
	Logon Scripts	328
	Assigning Home Directories	329
	Setting Account Options and Restrictions	330
	Managing Logon Hours	330
	Setting Permitted Logon Workstations	332
	Setting Dial-In and VPN Privileges	333
	Setting Account Security Options	335
	Managing User Profiles	336
	Local, Roaming, and Mandatory Profiles	337
	Using the System Utility to Manage Local Profiles	339
	Updating User and Group Accounts	343
	Renaming User and Group Accounts	344
	Copying Domain User Accounts	345
	Importing and Exporting Accounts	346
	Changing and Resetting Passwords	347
	Enabling User Accounts	348
	Managing Multiple User Accounts	349
	Setting Profiles for Multiple Accounts	350
	Setting Logon Hours for Multiple Accounts	351
	Setting Permitted Logon Workstations for Multiple Accounts	351
	Setting Logon, Password, and Expiration Properties for Multiple Accounts	351
	Troubleshooting Logon Problems	352
	Viewing and Setting Active Directory Permissions	353
Chapter 12	Managing File Systems and Drives	357
	Managing the File Services Role	357
	Adding Hard Disk Drives	363
	Physical Drives	363
	Preparing a Physical Drive for Use	365
	Using Disk Management	365
	Removable Storage Devices	368

Installing and Checking for a New Drive	370
Understanding Drive Status	370
Working with Basic, Dynamic, and Virtual Disks	372
Using Basic and Dynamic Disks	372
Special Considerations for Basic and Dynamic Disks	373
Changing Drive Types	374
Reactivating Dynamic Disks	376
Rescanning Disks	376
Moving a Dynamic Disk to a New System	376
Managing Virtual Hard Disks	377
Using Basic Disks and Partitions	378
Partitioning Basics	378
Creating Partitions and Simple Volumes	379
Formatting Partitions	382
Managing Existing Partitions and Drives	384
Assigning Drive Letters and Paths	384
Changing or Deleting the Volume Label	385
Deleting Partitions and Drives	386
Converting a Volume to NTFS	386
Resizing Partitions and Volumes	388
Repairing Disk Errors and Inconsistencies	390
Defragmenting Disks	393
Compressing Drives and Data	395
Encrypting Drives and Data	397
Understanding Encryption and the Encrypting File System	398
Working with Encrypted Files and Folders	400
Configuring Recovery Policy	401
Chapter 13 Administering Volume Sets and RAID Arrays	403
Using Volumes and Volume Sets	404
Understanding Volume Basics	404
Understanding Volume Sets	405
Creating Volumes and Volume Sets	407
Deleting Volumes and Volume Sets	410
Managing Volumes	410
Improving Performance and Fault Tolerance with RAID.	411
Implementing RAID on Windows Server 2008 R2	412
Implementing RAID 0: Disk Striping	412

Implementing RAID 1: Disk Mirroring	413
Implementing RAID 5: Disk Striping with Parity	415
Managing RAID and Recovering from Failures	416
Breaking a Mirrored Set	416
Resynchronizing and Repairing a Mirrored Set	417
Repairing a Mirrored System Volume to Enable Boot	418
Removing a Mirrored Set	418
Repairing a Striped Set Without Parity	419
Regenerating a Striped Set with Parity	419
Managing LUNs on SANs	420
Configuring Fibre Channel SAN Connections	421
Configuring iSCSI SAN Connections	422
Adding and Removing Targets	423
Creating, Extending, Assigning, and Deleting LUNs	423
Defining a Server Cluster in Storage Manager For SANs	424
Chapter 14 Managing File Screening and Storage Reporting	425
Understanding File Screening and Storage Reporting	425
Managing File Screening and Storage Reporting	429
Managing Global File Resource Settings	430
Managing the File Groups to Which Screens Are Applied	433
Managing File Screen Templates	435
Creating File Screens	438
Defining File Screening Exceptions	438
Scheduling and Generating Storage Reports	439
Chapter 15 Data Sharing, Security, and Auditing	441
Using and Enabling File Sharing	442
Configuring Standard File Sharing	445
Viewing Existing Shares	445
Creating Shared Folders	447
Creating Additional Shares on an Existing Share	450
Managing Share Permissions	450
The Different Share Permissions	451
Viewing Share Permissions	451
Configuring Share Permissions	452
Modifying Existing Share Permissions	453
Removing Share Permissions for Users and Groups	454

Managing Existing Shares.	454
Understanding Special Shares	454
Connecting to Special Shares	455
Viewing User and Computer Sessions	456
Stopping File and Folder Sharing	459
Configuring NFS Sharing.	459
Using Shadow Copies	461
Understanding Shadow Copies	461
Creating Shadow Copies	462
Restoring a Shadow Copy	462
Reverting an Entire Volume to a Previous Shadow Copy	463
Deleting Shadow Copies	463
Disabling Shadow Copies	464
Connecting to Network Drives	464
Mapping a Network Drive	464
Disconnecting a Network Drive	465
Object Management, Ownership, and Inheritance	465
Objects and Object Managers	466
Object Ownership and Transfer	466
Object Inheritance	467
File and Folder Permissions	468
Understanding File and Folder Permissions	468
Setting File and Folder Permissions	471
Auditing System Resources.	473
Setting Auditing Policies	473
Auditing Files and Folders	475
Auditing the Registry	477
Auditing Active Directory Objects	478
Using, Configuring, and Managing NTFS Disk Quotas	478
Understanding NTFS Disk Quotas and How NTFS Quotas Are Used	479
Setting NTFS Disk Quota Policies	481
Enabling NTFS Disk Quotas on NTFS Volumes	484
Viewing Disk Quota Entries	485
Creating Disk Quota Entries	486
Deleting Disk Quota Entries	487
Exporting and Importing NTFS Disk Quota Settings	488
Disabling NTFS Disk Quotas	489

Using, Configuring, and Managing Resource Manager	
Disk Quotas	489
Understanding Resource Manager Disk Quotas	489
Managing Disk Quota Templates	491
Creating Resource Manager Disk Quotas	493
Chapter 16 Data Backup and Recovery	495
Creating a Backup and Recovery Plan	495
Figuring Out a Backup Plan	495
The Basic Types of Backup	497
Differential and Incremental Backups	498
Selecting Backup Devices and Media	499
Common Backup Solutions	499
Buying and Using Backup Media	500
Selecting a Backup Utility	501
Backing Up Your Data: The Essentials	502
Installing the Windows Backup and Recovery Utilities	502
Getting Started with Windows Server Backup	503
Getting Started with the Backup Command-Line Utility	506
Working with Wbadmin Commands	508
Using General-Purpose Commands	508
Using Backup Management Commands	508
Using Recovery Management Commands	509
Performing Server Backups	510
Configuring Scheduled Backups	511
Modifying or Stopping Scheduled Backups	514
Creating and Scheduling Backups with Wbadmin	515
Running Manual Backups	517
Recovering Your Server from Hardware or Startup Failure	518
Starting a Server in Safe Mode	520
Resuming After a Failed Start	522
Backing Up and Restoring the System State	522
Restoring Active Directory	523
Restoring the Operating System and the Full System	524
Restoring Applications, Nonsystem Volumes, and Files and Folders	526

Managing Encryption Recovery Policy	528
Understanding Encryption Certificates and Recovery Policy	528
Configuring the EFS Recovery Policy	530
Backing Up and Restoring Encrypted Data and Certificates	531
Backing Up Encryption Certificates	531
Restoring Encryption Certificates	532
Chapter 17 Managing TCP/IP Networking	533
Navigating Networking in Windows Server 2008 R2	533
Managing Networking in Windows 7 and Windows Server 2008 R2	537
Installing TCP/IP Networking	539
Configuring TCP/IP Networking	541
Configuring Static IP Addresses	541
Configuring Dynamic IP Addresses and Alternate IP Addressing	543
Configuring Multiple Gateways	544
Configuring Networking for Hyper-V	546
Managing Network Connections	547
Checking the Status, Speed, and Activity for Local Area Connections	547
Enabling and Disabling Local Area Connections	547
Renaming Local Area Connections	548
Chapter 18 Administering Network Printers and Print Services	549
Managing the Print and Document Services Role	549
Using Print Devices	550
Printing Essentials	550
Configuring Print Servers	552
Enabling and Disabling File and Printer Sharing	553
Getting Started with Print Management	554
Installing Printers	556
Using the Autoinstall Feature of Print Management	556
Installing and Configuring Physically Attached Print Devices	556
Installing Network-Attached Print Devices	561
Connecting to Printers Created on the Network	563
Deploying Printer Connections	565

Configuring Point and Print Restrictions	566
Moving Printers to a New Print Server	569
Monitoring Printers and Printer Queues Automatically	570
Solving Spooling Problems	571
Configuring Printer Properties	572
Adding Comments and Location Information	572
Listing Printers in Active Directory	572
Managing Printer Drivers	573
Setting a Separator Page and Changing Print Device Mode	574
Changing the Printer Port	574
Scheduling and Prioritizing Print Jobs	575
Starting and Stopping Printer Sharing	576
Setting Printer Access Permissions	577
Auditing Print Jobs	578
Setting Document Defaults	579
Configuring Print Server Properties	579
Locating the Spool Folder and Enabling Printing on NTFS	579
Managing High-Volume Printing	580
Logging Printer Events	580
Enabling Print Job Error Notification	580
Managing Print Jobs on Local and Remote Printers	580
Viewing Printer Queues and Print Jobs	581
Pausing the Printer and Resuming Printing	581
Emptying the Print Queue	582
Pausing, Resuming, and Restarting Individual Document Printing	582
Removing a Document and Canceling a Print Job	582
Checking the Properties of Documents in the Printer	582
Setting the Priority of Individual Documents	582
Scheduling the Printing of Individual Documents	583
Chapter 19 Running DHCP Clients and Servers	585
Understanding DHCP	585
Using Dynamic IPv4 Addressing and Configuration	586
Using Dynamic IPv6 Addressing and Configuration	586
Checking IP Address Assignment	589
Understanding Scopes	590

Installing a DHCP Server	591
Installing DHCP Components	591
Starting and Using the DHCP Console	594
Connecting to Remote DHCP Servers	595
Starting and Stopping a DHCP Server	596
Authorizing a DHCP Server in Active Directory	596
Configuring DHCP Servers	596
Binding a DHCP Server with Multiple Network Interface Cards to a Specific IP Address	597
Updating DHCP Statistics	597
DHCP Auditing and Troubleshooting	597
Integrating DHCP and DNS	599
Integrating DHCP and NAP	601
Avoiding IP Address Conflicts	604
Saving and Restoring the DHCP Configuration	604
Managing DHCP Scopes	605
Creating and Managing Superscopes	605
Creating and Managing Scopes	606
Creating and Managing Split Scopes	616
Managing the Address Pool, Leases, and Reservations.	620
Viewing Scope Statistics	620
Enabling and Configuring MAC Address Filtering	621
Setting a New Exclusion Range	622
Deleting an Exclusion Range	623
Reserving DHCP Addresses	623
Modifying Reservation Properties	625
Deleting Leases and Reservations	625
Backing Up and Restoring the DHCP Database	625
Backing Up the DHCP Database	626
Restoring the DHCP Database from Backup	626
Using Backup and Restore to Move the DHCP Database to a New Server	627
Forcing the DHCP Server Service to Regenerate the DHCP Database	627
Reconciling Leases and Reservations	628
Chapter 20 Optimizing DNS	629
Understanding DNS	629
Integrating Active Directory and DNS	630
Enabling DNS on the Network	631

Understanding DNSSEC	633
Configuring Name Resolution on DNS Clients	635
Installing DNS Servers	637
Installing and Configuring the DNS Server Service	637
Configuring a Primary DNS Server	639
Configuring a Secondary DNS Server	642
Configuring Reverse Lookups	642
Configuring Global Names	644
Managing DNS Servers	645
Adding Remote Servers to the DNS Manager Console	646
Removing a Server from the DNS Manager Console	646
Starting and Stopping a DNS Server	647
Creating Child Domains Within Zones	647
Creating Child Domains in Separate Zones	647
Deleting a Domain or Subnet	648
Managing DNS Records	649
Adding Address and Pointer Records	650
Adding DNS Aliases with CNAME	651
Adding Mail Exchange Servers	652
Adding Name Servers	654
Viewing and Updating DNS Records	655
Updating Zone Properties and the SOA Record	655
Modifying the SOA Record	655
Allowing and Restricting Zone Transfers	657
Notifying Secondaries of Changes	658
Setting the Zone Type	659
Enabling and Disabling Dynamic Updates	660
Managing DNS Server Configuration and Security	660
Enabling and Disabling IP Addresses for a DNS Server	660
Controlling Access to DNS Servers Outside the Organization	661
Enabling and Disabling Event Logging	663
Using Debug Logging to Track DNS Activity	663
Monitoring a DNS Server	664
 <i>Index</i>	 667

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

Using Active Directory

- Introducing Active Directory **211**
- Working with Domain Structures **215**
- Working with Active Directory Domains **221**
- Understanding the Directory Structure **227**
- Using the Active Directory Recycle Bin **233**

Active Directory Domain Services (AD DS) is an extensible and scalable directory service that 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. To help manage this complex technology, I'll start with an overview of Active Directory and then explore its components.

Introducing Active Directory

Since Windows 2000, Active Directory has been the heart of Windows-based domains. Just about every administrative task you perform affects Active Directory in some way. Active Directory technology is based on standard Internet protocols and is designed to help you clearly define your network's structure.

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, such as `zeta.microsoft.com`, to numeric TCP/IP addresses, such as `192.168.19.2`. 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 20, "Optimizing DNS."

With Windows Server 2008 R2, you install Active Directory in a two-part process. First you use the Add Roles Wizard to add the Active Directory Domain Services role to the server. Then you run the Active Directory Installation Wizard (click Start, type **dcpromo** in the Search field, and then press Enter). 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, you can:

- 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.

Read-Only Domain Controller Deployment

As discussed in Chapter 1, "Windows Server 2008 R2 Administration Overview," domain controllers running Windows Server 2008 R2 can be configured as read-only domain controllers (RODCs). 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 domain controller installed in a forest or domain cannot be an RODC. However, you can configure subsequent domain controllers as read-only. For planning purposes, keep the following in mind:

- Prior to adding AD DS to a server that is running Windows Server 2008 R2 in a Windows Server 2003 or Windows 2000 Server forest, you must update the schema on the schema operations master in the forest by running `adprep /forestprep`.
- Prior to adding AD DS to a server that is running Windows Server 2008 R2 in a Windows Server 2003 or Windows 2000 Server domain, you must update the infrastructure master in the domain by running `adprep /domainprep /gpprep`.
- Prior to installing AD DS to create your first RODC in a forest, you must prepare the forest by running `adprep /rodcprep`.

New Active Directory Features

Active Directory Domain Service in Windows Server 2008 R2 has many new features that give administrators additional options for implementing and managing Active Directory. When you are using Windows Server 2008 R2 and have deployed the operating system 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 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 10.
- **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 10.

REAL WORLD Technically, you can use managed service accounts and managed virtual accounts in a mixed-mode domain environment. However, you must update the Active Directory schema for Windows Server 2008 R2 and you have to manually manage SPNs for managed service accounts.

- **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 2008 R2. These include:

- **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. A related option is on the Administrative Tools menu.
- **Active Directory Administrative Center** Provides a task-orientated interface for managing Active Directory. A related option is on the Administrative Tools menu.
- **Active Directory Web Services** Introduces a Web service interface for Active Directory domains.

These features are discussed in more detail in Chapter 8, “Core Active Directory Administration.” Also keep in mind that you must prepare Active Directory schema for the Active Directory Recycle Bin. The preparation procedures are the same as those discussed for RODCs in the previous section.

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).

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 7-1. 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 7-1 This network diagram depicts a wide area network (WAN) with multiple sites and subnets.

NOTE User and group accounts are discussed in Chapter 9, “Understanding User and Group Accounts.” Computer accounts and the various types of computers used in Windows Server 2008 R2 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, Windows Server 2008, and Windows Server 2008 R2.
- **Windows Server 2008** Supports domain controllers running Windows Server 2008 and Windows Server 2008 R2.
- **Windows Server 2008 R2** Supports domain controllers running Windows Server 2008 R2.

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 discontinuous 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 7-2 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 7-2 Domains in the same tree share a contiguous naming structure.

If the domains in a forest have discontinuous DNS names, they form separate domain trees within the forest. As shown in Figure 7-3, 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 7-3 Multiple trees in a forest have discontinuous naming structures.

You can access domain structures by using Active Directory Domains And Trusts, shown in Figure 7-4. Active Directory Domains And Trusts is a snap-in for the Microsoft Management Console (MMC). You can also start it from the Administrative Tools menu. You'll find separate entries for each root domain. In Figure 7-4, the active domain is `cpandl.com`.




FIGURE 7-4 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:

- **Windows Server 2003** Supports domain controllers running Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2.
- **Windows Server 2008** Supports domain controllers running Windows Server 2008 and Windows Server 2008 R2.
- **Windows Server 2008 R2** Supports domain controllers running Windows Server 2008 and Windows Server 2008 R2.

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 new features. These features include the Active Directory Recycle Bin, managed service accounts, and Authentication Mechanism Assurance.

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 only come 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:

- 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 7-5. This utility is a snap-in for the MMC, and you can also start it from the Administrative Tools menu.




FIGURE 7-5 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 2008 R2 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 7-6. 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 Administrative Tools menu.




FIGURE 7-6 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 2008 R2 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 2008 R2 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 20.

Using Windows 2000 and Later Computers with Active Directory

User computers running professional or business editions of Windows 2000, Windows XP, Windows Vista, or Windows 7 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 running Windows 2000 Server, Windows Server 2003, and Windows Server 2008 or later 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 Remove Role 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 several domain functional levels, including:

- **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. Windows NT and Windows 2000 domain controllers are no longer supported. 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 NT, Windows 2000, and Windows Server 2003 domain controllers are no longer supported. The

good news is that 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 NT, Windows 2000, Windows Server 2003, and Windows Server 2008 domain controllers are no longer supported. The good news is that a domain operating in Windows Server 2008 R2 mode can use all the latest Active Directory features, including the Active Directory Recycle Bin.

Using Windows Server 2003 Functional Level

After you upgrade the Windows NT structures in your organization, you can begin upgrading to Windows Server 2003 domain structures by upgrading Windows 2000 domain controllers to Windows Server 2003 or Windows Server 2008 domain controllers. Then, if you want to, you can change the functional level to Windows Server 2003 mode operations. Note that since Windows Server 2008 R2 runs only on 64-bit hardware, you'll likely need to install Windows Server 2008 R2 on new hardware rather than hardware designed for Windows NT, Windows 2000, or Windows Server 2003.

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 Windows Server 2003 installation media, you'll find `Adprep` in the `i386` subfolder.

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`.

After upgrading your servers, 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 2003, Windows Server 2008, and Windows Server 2008 R2 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 NT and Windows 2000 structures in your organization, you can begin upgrading to Windows Server 2008 domain structures by upgrading Windows Server 2003 domain controllers to Windows Server 2008 or Windows Server 2008 R2 domain controllers. Then, if you want to, 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 /gpprep`**. 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, 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 runs only on 64-bit hardware. You'll likely need to install Windows Server 2008 R2 on new hardware rather than on hardware designed for Windows NT, Windows 2000, or Windows Server 2003.

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 /gpprep`**. 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 and you can't go back to any other mode. 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.

Raising Domain and Forest Functionality

Domains operating in 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 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. Click Start, point to Administrative Tools, and then click Active Directory Domains And Trusts.
2. In the console tree, right-click the domain you want to work with, and then click Raise Domain Functional Level.

The current domain name and functional level are displayed in the Raise Domain Functional Level dialog box.

3. To change the domain functionality, select the new domain functional level from the list provided, and then click Raise. You can't reverse this action. Consider the implications carefully before you do this.
4. Click OK. The new domain functional level is replicated to each domain controller in the domain. This operation can take some time in a large organization.

You can raise the forest level functionality by following these steps:

1. Click Start, point to Administrative Tools, and then click Active Directory Domains And Trusts.
2. In the console tree, right-click the Active Directory Domains And Trusts node, and then click Raise Forest Functional Level.

The current forest name and functional level are displayed in the Raise Forest Functional Level dialog box.

3. To change the forest functionality, select the new forest functional level by using the list provided, and then click Raise. You can't reverse this action. Consider the implications carefully before you do this.

4. 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.

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 that 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 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 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 Windows Server 2008 R2, 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 Ntds.dit. This file's location is set when Active Directory is installed, and it should be on an NTFS file system

drive formatted for use with Windows Server 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 2008 R2 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 that 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 that 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 8. 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 Windows Server 2003 or higher functional mode. To learn how to configure caching, see "Configuring Universal Group Membership Caching" in Chapter 8.

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, configuration information for all domains in the forest, 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.

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 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 master 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 would make one domain controller the operations master, and you would designate the second as your standby operations master. The standby operations master could 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 that 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 relative ID 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 would probably want to place the relative ID 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 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, as discussed earlier in “Using Windows Server 2008 R2 Functional Level.” 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 and then raise the domain and forest functional levels to the Windows Server 2008 R2 level.

After these operations, you can access the recycle bin. From now on, when an Active Directory object is deleted, the object is put in a state referred to as *logically deleted*, moved to the Deleted Objects container, and its distinguished name is altered. A deleted object remains in the Deleted Objects container for the period of time set in the delete object lifetime value, which is 180 days by default.

REAL WORLD The `msDS-deletedObjectLifetime` attribute replaces the `tombstoneLifetime` attribute. However, when `msDS-deletedObjectLifetime` is set to `$null`, the lifetime value comes from the `tombstoneLifetime`. If the `tombstoneLifetime` is also set to `$null`, the default value is 180 days.

Recovering Deleted Objects

You can recover deleted objects from the Deleted Objects container by using an authoritative restore. The procedure has not changed from previous releases of Windows Server. What has changed, however, is the fact 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. 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. Click Start, type **Ldp.exe** in the Search box, and then press Enter.
2. On the Options menu, click Controls. In the Controls dialog box, select Return Deleted Objects in the Load Predefined list, and then 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 click OK.
4. On the View menu, 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 click OK.
5. In the console tree, double-click the root distinguished name and locate the CN=Deleted Objects container.
6. Locate and right-click the Active Directory object that you want to restore, and then 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, click Delete, and then 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, click Replace. Select the Extended check box, click Enter, and then click Run.

Using Windows PowerShell for Basic and Advanced Recovery

You can also use the Active Directory cmdlets for Windows PowerShell to recover deleted objects. 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. You need to import the module before you can use the cmdlets it provides. For more information, see “Active Directory Administrative Center and Windows PowerShell” in Chapter 8.

To use the Active Directory cmdlets for recovery, you need to open an elevated, administrator PowerShell prompt by right-clicking the Windows PowerShell entry on the menu and 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 Haddock"}  
-IncludeDeletedObjects | Restore-ADObject
```

```
Get-ADObject -Filter {SamAccountName -eq "richh"} -IncludeDeletedObjects  
| Restore-ADObject
```

It's important to 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
```

Managing File Systems and Drives

- Managing the File Services Role **357**
- Adding Hard Disk Drives **363**
- Working with Basic, Dynamic, and Virtual Disks **372**
- Using Basic Disks and Partitions **378**
- Managing Existing Partitions and Drives **384**

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 Windows 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 13, "Administering Volume Sets and RAID Arrays," looks at 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 Windows

Server operating system, all servers were installed with basic file services. With Windows Server 2008 R2, 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 12-1 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 backup utility included with Windows Server 2008 R2.
- **Storage Manager for SANs** Allows you to provision storage for storage area networks (SANs).
- **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.

TABLE 12-1 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.
Distributed File System (DFS)	Provides tools and services for DFS Namespaces and DFS Replication. DFS Replication is a newer and preferred replication technology. When a domain is running in Windows 2008 domain functional level, domain controllers use DFS Replication to provide more robust and granular replication of the SYSVOL directory.
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.
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.

ROLE SERVICE	DESCRIPTION
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.
Indexing Service	Allows indexing of files and folders for faster searching. Using the related query language, users can find files quickly. You cannot install Indexing Service and Windows Search Service on the same computer.
Services for Network File System	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 2008 R2 and UNIX operating systems by using the NFS protocol.
Windows Search Service	Enables fast file searches of resources on the server from clients that are compatible with Windows Search Service. This feature is designed primarily for desktop and small office implementations.
Windows Server 2003 File Services	Provides file services that are compatible with Windows Server 2003. This allows you to use a server running Windows Server 2008 R2 with servers running Windows Server 2003.

You can add the File Services role to a server by following these steps:


1. In Server Manager, select the Roles node in the left pane, and then click Add Roles. This starts the Add Roles Wizard. If the wizard displays the Before You Begin page, read the Welcome text, and then click Next.

NOTE During the setup process, shared files are created on the server. If you encounter a problem that causes the setup process to fail, you need to resume the setup process using the Add Role Services Wizard. After you restart Server Manager, select the File Services node under Roles. In the main pane, scroll down and then click Add Role Services. You can continue with the installation starting with step 3. If you were in the process of configuring domain-based DFS, you need to provide administrator credentials.
2. On the Select Server Roles page, select File Services, and then click Next twice.
3. On the Select Role Services page, select one or more role services to install. A summary of each role service is provided in Table 12-1. To allow for interoperability with UNIX, be sure to add Services for Network File System. Click Next.

4. A DFS namespace is a virtual view of shared folders located on different servers. To install DFS Namespaces, you work with several additional configuration pages:
- On the Create A DFS Namespace page, set the root name for the first namespace or elect to create a namespace later, as shown in the following screen. The namespace root name should be something that is easy for users to remember, such as CorpData. In a large enterprise, you may need to create separate namespaces for each major division.




- On the Select Namespace Type page, specify whether you want to create a domain-based namespace or a stand-alone namespace, as shown in the following screen. Domain-based namespaces can be replicated with multiple namespace servers to provide high availability, but they can have only up to 5,000 DFS folders. Stand-alone namespaces can have up to 50,000 DFS folders, but they are replicated only when you use failover server clusters and configure replication.




- If you are creating a domain-based namespace, on the Provide Credentials To Create A Namespace page, click Select, and then specify the user name and password for an account that is a member of the Domain Admins groups. This account is used to create the namespace.
- On the Configure Namespace page, you can add shared folders to the namespace as well as namespaces that are associated with a DFS folder, as shown in the following screen. Click Add. In the Add Folder To Namespace dialog box, click Browse. In the Browse For Shared Folders dialog box, select the shared folder to add, and then click OK. Type a

name for the folder to add, and then click OK. Next, type a name for the folder in the namespace. This name can be the same as the original folder name or a new name that will be associated with the original folder in the namespace. After you type a name, click OK to add the folder and complete the process.



NOTE You do not have to configure DFS Namespaces at this time. Once you install DFS Namespaces, DFS Replication, or both, you can use the DFS Management console to manage the related features. This console is installed and available on the Administrative Tools menu. See Chapter 15, “Data Sharing, Security, and Auditing,” for more information.

- With File Server Resource Manager, you can monitor the amount of space used on disk volumes and create storage reports. To install File Server Resource Manager, you work with two additional configuration pages:
 - On the Configure Storage Usage Monitoring page, select disk volumes for monitoring as shown in the following screen. When you select a volume and then click Options, you can set the volume usage threshold and choose the reports to generate when the volume reaches the threshold value. By default, the usage threshold is 85 percent.



- On the Set Report Options page, you can select a save location for usage reports, as shown in the following screen. One usage report of each type you select is generated each time a volume reaches its threshold. Old reports are not automatically deleted. The default save location is %SystemDrive%\StorageReports. To change the default location, click Browse, and then select the new save location in the Browse For Folder dialog box. You can also elect to receive reports by e-mail. To do this, you must specify the recipient e-mail addresses and the SMTP server to use.

Save reports at this location:

Receive reports by e-mail
 Reports can be sent to one or more e-mail addresses. Type each e-mail address where you want to receive the reports. Use semicolons (;) to separate multiple addresses.

E-mail addresses:

 Format: account@domain

An SMTP server must be used for sending the reports by e-mail. Select the SMTP server to use.

SMTP server:

NOTE You do not have to configure monitoring and reporting at this time. After you install FSRM, you can use the File Server Resource Manager console to manage the related features. This console is installed and available on the Administrative Tools menu.

6. To install Windows Search Service, you work with an additional configuration page that allows you to select the volumes to index. Indexing a volume makes it possible for users to search a volume quickly. However, indexing entire volumes can affect service performance, especially if you index the system volume. Therefore, you may want to index only specific shared folders on volumes, which you can do later on a per-folder basis.

NOTE You do not have to configure indexing at this time. After you install Windows Search Service, you can use the Indexing Options utility in Control Panel to manage the related features.

7. After you complete the optional pages, click Next. You'll see the Confirm Installation Options page. Click Install to begin the installation process. When Setup finishes installing the server with the features you selected, you'll see the Installation Results page. Review the installation details to ensure that all phases of the installation were completed successfully.

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. In Server Manager, expand the Roles node, and then select the File

Services node. In the main pane, the window is divided into several panels. Scroll down until you see the Role Services panel, and then click Add Role Services. You can then follow the previous procedure starting with step 3 to add role services.

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 2008 R2, 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 stand-alone 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 2008 R2 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. Typical drives have capacities of 500 gigabytes (GB) to 2 terabytes (TB). Many drive types are available for use with Windows Server 2008 R2, including Small Computer System Interface (SCSI), Parallel ATA (PATA), and Serial ATA (SATA).

The terms SCSI, PATA, 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, PATA drives use PATA controllers, and so on. 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 Ultra320 SCSI or SATA II, 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 7,200 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 7,200 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, 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).

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 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 that 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.

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 and the Server Manager console. You can also add it to custom MMCs. In Computer Management and in Server Manager, you can access Disk Management by expanding the Storage node and then selecting Disk Management.

Regardless of whether you are using Computer Management or Server Manager, 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 12-1, 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 12-1 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 2008 R2 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 can be used with previous versions of Windows.
- **Dynamic** An enhanced fixed disk type for Windows Server 2008 R2 that you can update without having to restart the system (in most cases). Dynamic disks are divided into volumes and can be used only with Windows 2000 and later releases of Windows.

- **Removable** The standard disk type associated with removable storage devices. Removable storage devices can be formatted with exFAT, FAT, FAT32, or NTFS.
- **Virtual** The virtual hard disk (VHD) disk type associated with virtualization can be used when a computer is running Windows 7, Windows Server 2008 R2, or later releases. Computers can use VHDs just like they use regular fixed disks and can even be configured to boot from a VHD.

REAL WORLD Windows Vista with SP1 or later, Windows 7, and Windows Server 2008 or later all support exFAT with removable storage devices. The exFAT file system is the next generation file system in the FAT (FAT12/16, FAT32) family. While retaining the ease-of-use advantages of FAT32, exFAT overcomes the 4-GB file size limit on FAT32 and its 32-GB partition size limit on Windows systems. The exFAT file system also supports allocation unit sizes of up to 32,768 KB.

The exFAT file system is designed so that it can be used with any compliant operating system or device. This means you can remove an exFAT storage device from a compliant camera and insert it into a compliant phone or vice versa without having to do any reformatting. It also means that you can remove an exFAT storage device from a computer running Mac OS or Linux and insert it into a computer running Windows.

From the Disk Management window, you can get more detailed information on a drive section by right-clicking it and then selecting Properties. When you do this, you see a dialog box. With fixed disks, the dialog box is much like the one shown on the left in Figure 12-2. With removable disks, the dialog box is much like the one shown on the right in Figure 12-2. This is the same dialog box that you can open from Windows Explorer (by selecting the top-level folder for the drive and then choosing Properties from the File menu).




FIGURE 12-2 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 3, "Managing Servers Running Windows Server 2008 R2," 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:

- 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, including USB 1.0, USB 1.1, USB 2.0, FireWire 400, and FireWire 800.

USB 2.0 is the industry standard, and it supports data transfers at a maximum rate of 480 Mbps, with sustained data transfer rates usually from 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. 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 LCD 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 currently are used, including FireWire 400 and FireWire 800. FireWire 400 (IEEE 1394a) has maximum sustained transfer rates of up to 400 Mbps. FireWire 800 (IEEE 1394b) has maximum sustained transfer rates of up to 800 Mbps. As with USB, if you connect a FireWire 800 device to a FireWire 400 port or vice versa, the device operates at the significantly reduced FireWire 400 transfer speed.

FireWire 400 and FireWire 800 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 cables always have bus power and have nine pins and nine connectors.

Another option is External Serial ATA (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 may be able to get a device with a dual interface that supports USB 2.0 and FireWire 400, or a triple interface that supports USB 2.0, FireWire 400, and FireWire 800. A device with dual or triple interfaces gives you more options. There also are devices with quadruple interfaces.

Working with removable disks is similar to working with fixed disks. You can do the following:

- Right-click a removable disk and select Open or Explore to examine the disk’s contents in Windows Explorer.
- 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.
- 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” later in this chapter.

When you work with removable disks, you can customize disk and folder views. To do this, right-click the disk or folder, select Properties, and then click the Customize tab. You can then specify the default folder type to control the default details displayed. For example, you can set the default folder type as Documents or Pictures And Videos. You can also set folder pictures and folder icons.

Removable disks support network file and folder sharing. You configure sharing on removable disks in the same way that 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 devices without shutting off the computer. Typically, hot-swappable drives are installed and removed from the front of the computer. If your computer supports hot swapping of 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 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 that you installed.
2. Disks can use either the MBR or GPT partition style. Select the partition style you want to use for the disk or disks you are initializing.
3. Click OK. If you elected to initialize disks, Windows writes a disk signature to the disks and initializes the disks with the basic disk type.

If you don't want to use the Initialize Disk dialog box, you can close it and use Disk Management instead to view and work with the 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 right-click the disk's icon and select Online. Right-click the disk's icon again, and select Initialize Disk. You can then initialize the disk as discussed previously.

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 12-2 summarizes the most common status values.

TABLE 12-2 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 right-clicking the disk and selecting Reactivate Disk. If this doesn't work, the disk might have physical damage or you might need to run a thorough check of the disk.
Offline	The disk isn't accessible and might be corrupted or temporarily unavailable. If the disk name changes to Missing, the disk can no longer be located or identified on the system.	Check for problems with the drive, its controller, and cables. Make sure that the drive has power and is connected properly. Use the Reactivate Disk command to bring the disk back online (if possible).
Foreign	The disk has been moved to your computer but hasn't been imported for use. A failed drive brought back online might sometimes be listed as Foreign.	Right-click the disk, and then click Import Foreign Disks to add the disk to the system.
Unreadable	The disk isn't accessible currently, which can occur when disks are being rescanned. Both dynamic and basic disks display this status.	With FireWire and USB card readers, you might 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. Right-click the disk, and then click Rescan Disk (on the Action menu) to try to correct the problem. You might also want to reboot the system.

STATUS	DESCRIPTION	RESOLUTION
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 2008 R2, right-click the disk, and then click Initialize Disk.
No Media	No media has been inserted into the CD-ROM or removable drive, or the media has been removed. Only CD-ROM and removable disk types display this status.	Insert a CD-ROM, a floppy disk, 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 2008 R2 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 2008 R2 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 may 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, right-click the primary partition you want to mark as active, and then click Mark Partition As Active. You can't mark dynamic disk volumes as active. When you convert a basic disk containing the active partition to a dynamic disk, this partition becomes a simple volume that's active automatically.

Changing Drive Types

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 2008 R2 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 2008 R2.

To convert a basic disk to a dynamic disk, follow these steps:

1. In Disk Management, right-click a basic disk that you want to convert, either in the Disk List view or in the left pane of the Graphical View. Then click Convert To Dynamic Disk.
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. Click OK to continue.

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, 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. 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, 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, right-click the dynamic disk you want to reactivate, and then click Reactivate Disk. Confirm the action when prompted.
2. If the drive status doesn't change, you might need to reboot the computer. If this still doesn't resolve the problem, check for problems with the drive, its controller, and the cables. Also make sure that the drive has power and is connected properly.

Rescanning Disks

Rescanning all drives on a system updates the drive configuration information on the computer. Rescanning can sometimes resolve a problem with drives that show a status of Unreadable. You rescan disks on a computer by choosing Rescan Disks from the Action menu in Disk Management.

Moving a Dynamic Disk to a New System

An important advantage of dynamic disks over basic disks is that you can easily move 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 2008 R2 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 boot might fail.
3. Check whether any dynamic disks that you want to move are part of a spanned, extended, or striped set. If they are, you should make a note of which disks are part of which set and plan on moving all disks in a set

together. If you are moving only part of a disk set, you should be aware of the consequences. For spanned, extended, or striped volumes, moving only part of the set will make the related volumes unusable on the current computer and on the computer to which you are planning to move the disks.

When you are ready to move the disks, follow these steps:

1. On the original computer, start Computer Management. Then, in the left pane, select Device Manager. In the Device list, expand Disk Drives. This shows a list of the physical disk drives on the computer. Right-click each disk that you want to move, and then click Uninstall. If you are unsure which disks to uninstall, right-click each disk and click Properties. In the Properties dialog box, 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 that you want to move are still listed, right-click each disk, and then click Remove Disk.
3. After you perform these procedures, you can move the dynamic disks. If the disks are hot-swappable disks and this feature is supported on both computers, remove the disks from the original computer and then install them on the destination computer. Otherwise, turn off both computers, remove the drives from the original computer, and then install them on the destination computer. When you have finished, restart the computers.
4. On the destination computer, access Disk Management, and then choose Rescan Disks from the Action menu. When Disk Management finishes scanning the disks, right-click any disk marked Foreign, and then click Import. You should now be able to access the disks and their volumes on the destination computer.

NOTE In most cases, the volumes on the dynamic disks should retain the drive letters that 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, 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 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 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, right-click the disk entry in Graphical View, and then click Initialize Disk. In the Initialize Disk dialog box, the disk is selected for initialization. Specify the disk type as MBR or GPT, and then click OK.

After initializing the disk, 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 right-clicking the disk entry in Graphical View and selecting Offline or Online, respectively. If you no longer want to use a VHD, you can detach it by right-clicking the disk entry in Graphical View, selecting Detach VHD, and then clicking OK in the Detach Virtual Hard Disk dialog box.

You can use VHDs created with other programs as well. If you created a VHD using another program or have a detached VHD that you want to attach, you can work with the VHD by completing the following steps:

1. In Disk Management, click the Attach VHD option on the Action menu.
2. In the Attach Virtual Hard Disk dialog box, click Browse. Use the Browse Virtual Disk Files dialog box to select the .vhd file for the virtual hard disk, and then click Open.
3. If you want to attach the VHD in read-only mode, select Read-Only. Click OK to attach the VHD.

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 2008 R2, 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 a high-level formatting that creates the file system structure rather than a low-level formatting that sets up the drive for initial use. You're probably very familiar with the C drive used by Windows Server 2008 R2. 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 2008 R2 lets you specify drive letters. Generally, the drive letters C through Z are available for your use.

NOTE The drive letter A is usually assigned to a system's floppy disk drive. If the system has a second floppy disk drive, the letter B is assigned to it, so you can use only the letters C through Z. Don't forget that CD-ROMs, Zip 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 2008 R2 simplifies the Disk Management user interface by using one set of dialog boxes and wizards for both partitions and volumes. The first three volumes on a basic drive are created automatically as primary partitions. If you try to create a fourth volume on a basic drive, the remaining free space on the drive is converted automatically to an extended partition with a logical drive of the size you designate by using the new volume feature in the extended partition. Any subsequent volumes are created in the extended partitions as logical drives automatically.

In Disk Management, you create partitions, logical drives, and simple volumes by following these steps:

1. In Disk Management's Graphical View, right-click an unallocated or free area, and then click New Simple Volume. This starts the New Simple Volume Wizard. Read the Welcome page, and then click Next.
2. The Specify Volume Size page, shown in Figure 12-3, 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 field, and then click Next.




FIGURE 12-3 Set the size of the volume on the Specify Volume Size page.

3. On the Assign Drive Letter Or Path page, shown in Figure 12-4, specify whether you want to assign a drive letter or path, and then click Next. The following options are available:
 - **Assign The Following Drive Letter** Choose this option to assign a drive letter. Then select an available drive letter in the list provided. By default, Windows Server 2008 R2 selects the lowest available drive letter and excludes reserved drive letters as well as those assigned to local disks or network drives.
 - **Mount In The Following Empty NTFS Folder** Choose this option to mount the partition in an empty NTFS folder. You must then type the path to an existing folder or click Browse to search for or create a folder to use.
 - **Do Not Assign A Drive Letter Or Drive Path** Choose this option if you want to create the partition without assigning a drive letter or path. If you

later want the partition to be available for storage, you can assign a drive letter or path at that time.

NOTE You don't have to assign volumes a drive letter or a path. A volume with no designators is considered to be unmounted and is for the most part unusable. An unmounted volume can be mounted by assigning a drive letter or a path at a later date. See "Assigning Drive Letters and Paths" later in this chapter.




FIGURE 12-4 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 12-5, 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 FAT32 or NTFS. NTFS is selected by default in most cases. 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 1,024 bytes. With these settings, small files use less disk space.
 - **Volume Label** Sets a text label for the partition. This label is the partition's volume name and is set to New Volume by default. You can

change the volume label at any time by right-clicking the volume in Windows Explorer, clicking Properties, and typing a new value in the Label field provided on the General tab.

- **Perform A Quick Format** Tells Windows Server 2008 R2 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. 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 12-5 Set the formatting options for the partition on the Format Partition page.

5. Click Next, confirm your options, and then click Finish.

Formatting Partitions

Formatting creates a file system on a partition and permanently deletes any existing data. This is a high-level formatting that creates the file system structure rather than a low-level formatting that initializes a drive for use. To format a partition, right-click the partition and then click Format. This opens the Format dialog box, shown in Figure 12-6.




FIGURE 12-6 Format a partition in the Format dialog box by specifying its file system type and volume label.

You use the formatting fields 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 FAT32 or NTFS. NTFS is the native file system type for Windows NT and later releases of Windows.
- **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 1,024 bytes. With these settings, small files use less disk space.
- **Perform A Quick Format** Tells Windows Server 2008 R2 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. Built-in compression is available only for NTFS. 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.

When you're ready to proceed, click OK. Because formatting a partition destroys any existing data, Disk Management gives you one last chance to cancel the procedure. Click OK to start formatting the partition. Disk Management changes the drive's status to reflect the formatting and the percentage of completion. When formatting is complete, the drive status changes to reflect this.

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 2008 R2 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, Windows 7, Windows Server 2008, and 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/en-us). 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, right-click the drive you want to configure in Disk Management, and then click Change Drive Letter And Paths. This opens the dialog box shown in Figure 12-7. You can now do the following:

- **Add a drive path** Click Add, select Mount In The Following Empty NTFS Folder, and then type the path to an existing folder, or click Browse to search for or create a folder.
- **Remove a drive path** Select the drive path to remove, click Remove, and then click Yes.
- **Assign a drive letter** Click Add, select Assign The Following Drive Letter, and then choose an available letter to assign to the drive.
- **Change the drive letter** Select the current drive letter, and then click Change. Select Assign The Following Drive Letter, and then choose a different letter to assign to the drive.
- **Remove a drive letter** Select the current drive letter, click Remove, and then click Yes.

NOTE If you try to change the letter of a drive that's in use, Windows Server 2008 R2 displays a warning. You need to exit programs that are using the drive and try again or allow Disk Management to force the change by clicking Yes when prompted.




FIGURE 12-7 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 FAT32, 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 FAT32 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 2008 R2 utilities, including Windows Explorer, it can provide information about a drive's contents. You can change or delete a volume label using Disk Management or Windows Explorer.

Using Disk Management, you can change or delete a label by following these steps:

1. Right-click the partition, and then click Properties.
2. On the General tab of the Properties dialog box, type a new label for the volume in the Label text box or delete the existing label. Click OK.

Using Windows Explorer, you can change or delete a label by following these steps:

1. Right-click the drive icon, and then click Properties.
2. On the General tab of the Properties dialog box, type a new label for the volume in the Label text box or delete the existing label. 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 2008 R2 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, a volume, or a logical drive by following these steps:

1. In Disk Management, right-click the partition, volume, or drive you want to delete, and then click Explore. Using Windows Explorer, move all the data to another volume or verify an existing backup to ensure that the data was properly saved.
2. In Disk Management, right-click the partition, volume, or drive again, and then click Delete Partition, Delete Volume, or Delete Logical Drive as appropriate.
3. Confirm that you want to delete the selected item by clicking Yes.

The steps for deleting an extended partition differ slightly from those for deleting a primary partition or a logical drive. To delete an extended partition, follow these steps:

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 2008 R2 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 2008 R2 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
```

where *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:

<i>volume</i>	Sets the volume to work with
/FS:NTFS	Converts to NTFS
/V	Sets verbose mode
/X	Forces the volume to dismount before the conversion (if necessary)
/CvtArea: <i>filename</i>	Sets the name of a contiguous file in the root directory to be a placeholder for NTFS system files
/NoSecurity	Removes all security attributes and makes all files and directories accessible to the group Everyone

The following sample statement uses Convert:

```
convert C: /FS:NTFS /V
```

Using the Convert Utility

Before you use the Convert utility, determine whether the partition is being used as the active boot partition or a system partition containing the operating system. You can convert the active boot partition to NTFS. Doing so requires that the system gain exclusive access to this partition, which can be obtained only during startup. Thus, if you try to convert the active boot partition to NTFS, Windows Server 2008 R2 displays a prompt asking if you want to schedule the drive to be converted the next time the system starts. If you click Yes, you can restart the system to begin the conversion process.

TIP Often, you will need to restart a system several times to completely convert the active boot partition. Don't panic. Let the system proceed with the conversion.

Before the Convert utility actually converts a drive to NTFS, the utility checks whether the drive has enough free space to perform the conversion. Generally, Convert needs a block of free space that's roughly equal to 25 percent of the total space used on the drive. For example, if the drive stores 200 GB of data, Convert needs about 50 GB of free space. If the drive doesn't have enough free space, Convert

aborts and tells you that you need to free up some space. On the other hand, if the drive has enough free space, Convert initiates the conversion. Be patient. The conversion process takes several minutes (longer for large drives). Don't access files or applications on the drive while the conversion is in progress.

You can use the `/CvtArea` option to improve performance on the volume so that space for the master file table (MFT) is reserved. This option helps to prevent fragmentation of the MFT. How? Over time, the MFT might grow larger than the space allocated to it. The operating system must then expand the MFT into other areas of the disk. Although the Disk Defragmenter 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 2008 R2 doesn't use Ntldr and `Boot.ini` to load the operating system. Instead, Windows Server 2008 R2 has a preboot environment in which Windows Boot Manager is used to control startup and load the boot application you've 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 2008 R2, 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 FAT32. Additionally, you can't extend a system or boot volume, regardless of its configuration.

You can shrink a simple volume or a spanned volume by following these steps:

1. In Disk Management, right-click the volume that you want to shrink, and then click Shrink Volume. This option is available only if the volume meets the previously discussed criteria.
2. In the field provided in the Shrink dialog box, shown in Figure 12-8, enter the amount of space to shrink.




FIGURE 12-8 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. Click Shrink to shrink the volume.

You can extend a simple volume or a spanned volume by following these steps:

1. In Disk Management, right-click the volume that you want to extend, and then click Extend Volume. This option is available only if the volume meets the previously discussed criteria and free space is available on one or more of the system's dynamic disks.
2. In the Extend Volume Wizard, read the introductory message, and then click Next.
3. On the Select Disks page, select the disk or disks from which you want to allocate free space. Any disks currently being used by the volume are automatically selected. By default, all remaining free space on those disks is selected for use.
4. With dynamic disks, you can specify the additional space that you want to use on other disks by performing the following tasks:
 - Click the disk, and then click Add to add the disk to the Selected list.
 - Select each disk in the Selected list, and then, in the Select The Amount Of Space In MB list, specify the amount of unallocated space to use on the selected disk.
5. Click Next, confirm your options, and then click Finish.

Repairing Disk Errors and Inconsistencies

Windows Server 2008 R2 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.

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 have 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 2008 R2 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 2008 R2 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 may want to (or may 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, 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.

Running Check Disk from the Command Line

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 E:
```

To find and repair errors that are on the E drive, use the following command:

```
chkdsk /f E:
```

NOTE Check Disk can't repair volumes that are in use. If a volume is in use, Check Disk displays a prompt that asks if you want to schedule the volume to be checked the next time you start the system. Click Yes to schedule this.

The complete syntax for Check Disk is shown here:

```
chkdsk [volume[[path]filename]] [/F] [/V] [/R] [/X] [/I] [/C] [/L[:size]]
```

The options and switches for Check Disk are used as follows:

<i>volume</i>	Sets the volume to work with.
<i>[path]filename</i>	FAT/FAT32 only: Specifies files to check for fragmentation.
<i>/F</i>	Fixes errors on the disk.
<i>/V</i>	On FAT/FAT32: Displays the full path and name of every file on the disk. On NTFS: Displays cleanup messages, if any.
<i>/R</i>	Locates bad sectors and recovers readable information (implies <i>/F</i>).
<i>/X</i>	Forces the volume to dismount first if necessary (implies <i>/F</i>).
<i>/I</i>	NTFS only: Performs a minimum check of index entries.
<i>/C</i>	NTFS only: Skips checking of cycles within the folder structure.
<i>/L:size</i>	NTFS only: Changes the log file size.

Running Check Disk Interactively

You can run Check Disk interactively by using Windows Explorer or Disk Management. Follow these steps:

1. Right-click the drive, and then click Properties.
2. On the Tools tab of the Properties dialog box, click Check Now.
3. As shown in Figure 12-9, you can now do the following:
 - Check for errors without repairing them. Click Start without selecting either of the check boxes.
 - Check for errors and fix them. Make the appropriate selections in the check boxes to fix file system errors, recover bad sectors, or both, and then click Start.



FIGURE 12-9 Use Check Disk to check a disk for errors and repair them.

Defragmenting 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 2008 R2 can manually or automatically defragment disks using Disk Defragmenter. The more frequently data is updated on drives, the more often you should run this tool.

You can manually defragment a disk by following these steps:

1. In Server Manager, select the Storage node and then the Disk Management node. Right-click a drive, and then click Properties.
2. On the Tools tab, click Defragment Now. In the Disk Defragmenter dialog box, select a disk, and then click Analyze Disk. Disk Defragmenter then analyzes the disk to determine whether it needs to be defragmented. If so, it recommends that you defragment now.
3. In the Disk Defragmenter dialog box, select a disk, and then click Defragment Disk.

NOTE Depending on the size of the disk, defragmentation can take several hours. You can click Stop Operation at any time to stop defragmentation.

When you enable automatic defragmentation, Windows Server 2008 R2 runs Disk Defragmenter automatically on a specific schedule, such as at 1:00 A.M. every Wednesday. As long as the computer is powered on at the scheduled run time, automatic defragmentation occurs. You can configure and manage automated defragmentation by following these steps:

1. In Server Manager, select the Storage node and then the Disk Management node. Right-click a drive, and then click Properties.
2. On the Tools tab, click Defragment Now. This displays the Disk Defragmenter dialog box, shown in Figure 12-10.




FIGURE 12-10 Disk Defragmenter analyzes and defragments disks efficiently.

3. To cancel automated defragmentation, click **Configure Schedule**, clear **Run On A Schedule**, and then click **OK**. Click **Close**, and skip the remaining steps.
4. To enable automated defragmentation, click **Turn On Schedule**. In the **Modify Schedule** dialog box, shown in Figure 12-11, select **Run On A Schedule**, and then set the run schedule. In the **Frequency** list, you can choose **Daily**, **Weekly**, or **Monthly**. If you choose a weekly or monthly run schedule, you need to select the run day of the week or month from the **Day** list. Finally, the **Time** list lets you set the time of the day that automated defragmentation should occur.
5. If you want to modify the run schedule, click **Configure Schedule**. In the **Modify Schedule** dialog box, shown in Figure 12-11, set the run schedule as discussed in the previous step.
6. If you want to manage which disks are defragmented, click **Select Disks**. In the **Select Disks For Schedule** dialog box, select which disks should be defragmented. By default, all disks installed within or connected to the computer are defragmented, and any new disks are defragmented automatically as well. In the **Disks To Include In Schedule** list, select the check boxes for disks that should be defragmented automatically and clear the check boxes for disks that should not be defragmented automatically. Click **OK**.
7. Click **OK**, and then click **Close** to save your settings.

NOTE Windows Vista with SP1 or later, Windows 7, and Windows Server 2008 or later releases of Windows automatically perform cyclic pickup defragmentation. With this feature, when a scheduled defragmentation pass is stopped and rerun, the computer automatically picks up the next unfinished volume in line to be defragmented.




FIGURE 12-11 Set the run schedule for automated defragmentation.

Compressing Drives and Data

When you format a drive for NTFS, Windows Server 2008 R2 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.

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 2008 R2 automatically decrypts the data and then compresses it. Likewise, if you try to encrypt compressed data, Windows Server 2008 R2 uncompresses the data and then encrypts it.

Compressing Drives

To compress a drive and all its contents, follow these steps:

1. In Windows Explorer or Disk Management, right-click the drive that you want to compress, and then click Properties.
2. On the General tab, select Compress Drive To Save Disk Space, and then click OK.
3. In the Confirm Attribute Changes dialog box, select whether to apply the changes to subfolders and files, and then click OK.

Compressing Directories and Files

If you decide not to compress a drive, Windows Server 2008 R2 lets you selectively compress directories and files. To compress a file or directory, follow these steps:

1. In Windows Explorer, right-click the file or directory that you want to compress, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced. In the Advanced Attributes dialog box, select the Compress Contents To Save Disk Space check box, as shown in Figure 12-12. Click OK twice.




FIGURE 12-12 With NTFS, you can compress a file or directory by selecting the Compress Contents To Save Disk Space check box in the Advanced Attributes dialog box.

For an individual file, Windows Server 2008 R2 marks the file as compressed and then compresses it. For a directory, Windows Server 2008 R2 marks the directory as compressed and then compresses all the files in it. If the directory contains subfolders, Windows Server 2008 R2 displays a dialog box that allows you to compress all the subfolders associated with the directory. Simply select Apply Changes To This Folder, Subfolders, And Files, and then click OK. 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

You can remove compression from a drive by following these steps:

1. In Windows Explorer or Disk Management, right-click the drive that contains the data you want to expand, and then click Properties.

2. Clear the Compress Drive To Save Disk Space check box, and then click OK.
3. In the Confirm Attribute Changes dialog box, select whether to apply the change to subfolders and files, and then click OK.

TIP Windows always checks the available disk space before expanding compressed data. You should too. If less free space is available than used space, you might not be able to complete the expansion. For example, if a compressed drive uses 150 GB of space and has 70 GB of free space available, you won't have enough free space to expand the data.

Expanding Compressed Directories and Files

If you decide that you want to expand a compressed file or directory, follow these steps:

1. Right-click the file or directory in Windows Explorer, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced. Clear the Compress Contents To Save Disk Space check box. Click OK twice.

With files, Windows Server 2008 R2 removes compression and expands the file. With directories, Windows Server 2008 R2 expands all the files within the directory. If the directory contains subfolders, you also have the opportunity to remove compression from the subfolders. To do this, select Apply Changes To This Folder, Subfolders, And Files when prompted, and then click OK.

TIP Windows Server 2008 R2 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 that you can use with Windows Server 2008 R2. 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 he or 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 16. 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 in the event that 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 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 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 2008 R2 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 2008 R2 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 2008 R2 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.

To encrypt a file or directory, follow these steps:

1. Right-click the file or directory that you want to encrypt, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced, and then select the Encrypt Contents To Secure Data check box. Click OK twice.

NOTE You can't encrypt compressed files, system files, or read-only files. If you try to encrypt compressed files, the files are automatically uncompressed and then encrypted. If you try to encrypt system files, you get an error.

For an individual file, Windows Server 2008 R2 marks the file as encrypted and then encrypts it. For a directory, Windows Server 2008 R2 marks the directory as encrypted and then encrypts all the files in it. If the directory contains subfolders, Windows Server 2008 R2 displays a dialog box that allows you to encrypt all the

subfolders associated with the directory. Simply select Apply Changes To This Folder, Subfolders, And Files, and then click OK.

NOTE On NTFS volumes, files remain encrypted even when they're moved, copied, or renamed. If you copy or move an encrypted file to a FAT or FAT32 drive, the file is automatically decrypted before being copied or moved. Thus, you must have proper permissions to copy or move the file.

You can grant special access to an encrypted file or folder by right-clicking the file or folder in Windows Explorer and then selecting Properties. On the General tab of the Properties dialog box, click Advanced. In the Advanced Attributes dialog box, click Details. In the Encryption Details For dialog box, users who have access to the encrypted file are listed by name. To allow another user access to the file, click Add. If a user certificate is available for the user, select the user's name in the list provided, and then click OK. Otherwise, click Find User to locate the certificate for the user.

Working with Encrypted Files and Folders

Previously, I said that 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 or FAT32 volume, the files are decrypted before transfer and then transferred as standard files. FAT and FAT32 don'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 or FAT32 volume on another computer. FAT and FAT32 don't support encryption.

After you transfer a sensitive file that has been encrypted, you might want to confirm that the encryption is still applied. Right-click the file and then select Properties. On the General tab of the Properties dialog box, click Advanced. The Encrypt Contents To Secure Data option should be selected.

Configuring Recovery 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 5.
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 data, purpose, and more.
4. To designate an additional recovery agent, right-click Encrypting File System, and then click Add Data Recovery Agent. This starts the Add Recovery Agent Wizard, which you can use to select a previously generated certificate that has been assigned to a user and mark it as a designated recovery certificate. Click Next.
5. On the Select Recovery Agents page, you can select certificates published in Active Directory or use certificate files. If you want to use a published certificate, click Browse Directory, and then, in the Find Users, Contacts, And Groups dialog box, select the user 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, click Browse Folders. In the Open dialog box, use the options provided to select and open the certificate file you want to use.

SECURITY ALERT Before you designate additional recovery agents, you should consider setting up a root certificate authority (CA) in the domain. Then you can use the Certificates snap-in to generate a personal certificate that uses the EFS Recovery Agent template. The root CA must then approve the certificate request so that the certificate can be used.

6. To delete a recovery agent, select the recovery agent's certificate in the right pane, and then press Delete. When prompted to confirm the action, click Yes to permanently and irrevocably delete the certificate. If the recovery policy is empty (meaning that it has no other designated recovery agents), EFS will be turned off so that files can no longer be encrypted.

Decrypting Files and Directories

If you want to decrypt a file or directory, follow these steps:

1. In Windows Explorer, right-click the file or directory, and then click Properties.
2. On the General tab of the Properties dialog box, click Advanced. Clear the Encrypt Contents To Secure Data check box. Click OK twice.

With files, Windows Server 2008 R2 decrypts the file and restores it to its original format. With directories, Windows Server 2008 R2 decrypts all the files within the directory. If the directory contains subfolders, you also have the option to remove encryption from the subfolders. To do this, select Apply Changes To This Folder, Subfolders, And Files when prompted, and then click OK.

TIP Windows Server 2008 R2 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.