

Microsoft® Exchange Server 2007 Administrator's Companion

*Walter Glenn, Scott Lowe,
and Joshua Maher*

To learn more about this book, visit Microsoft Learning at
<http://www.microsoft.com/MSPress/books/9545.aspx>

9780735623507

Microsoft®
Press

Resources for IT Professionals

Published and Forthcoming Titles from Microsoft Press

→ Windows Server

Microsoft® Windows Server® 2003
Resource Kit

Microsoft MVPs and Partners with
Microsoft Windows Server Team
978-0-7356-2232-6

Microsoft Windows Server 2003
Administrator's Companion
Second Edition

Charlie Russel, Sharon Crawford,
and Jason Gerend
978-0-7356-2047-6

Microsoft Windows Server 2003
Inside Out

William R. Stanek
978-0-7356-2048-3

Microsoft Windows Server 2003
Administrator's Pocket Consultant
Second Edition

William R. Stanek
978-0-7356-2245-6

→ Windows Client

Windows Vista™
Resource Kit

Tulloch, Northrup, Honeycutt,
Russel, and Wilson with the
Microsoft Windows Vista Team
978-0-7356-2283-8

Windows Vista
Administrator's Pocket Consultant

William R. Stanek
978-0-7356-2296-8

Microsoft Windows® XP
Professional
Resource Kit
Third Edition

The Microsoft Windows Team with
Charlie Russel and Sharon Crawford
978-0-7356-2167-1

Microsoft Windows XP
Professional
Administrator's Pocket Consultant
Second Edition

William R. Stanek
978-0-7356-2140-4

Microsoft Windows Command-Line
Administrator's Pocket Consultant

William R. Stanek
978-0-7356-2038-4

→ SQL Server 2005

Microsoft SQL Server™ 2005
Administrator's Pocket Consultant

William R. Stanek
978-0-7356-2107-7

Microsoft SQL Server 2005
Administrator's Companion

Whalen, Garcia, et al.
978-0-7356-2198-5

Inside Microsoft SQL Server 2005:
The Storage Engine

Kalen Delaney
978-0-7356-2105-3

Inside Microsoft SQL Server 2005:
T-SQL Programming

Itzik Ben-Gan, Dejan Sarka, and
Roger Wolter
978-0-7356-2197-8

→ Exchange Server 2007

Microsoft Exchange Server 2007
Administrator's Companion
Walter Glenn and Scott Lowe
978-0-7356-2350-7

Microsoft Exchange Server 2007
Administrator's Pocket Consultant

William R. Stanek
978-0-7356-2348-4

→ Scripting

Microsoft Windows PowerShell™
Step by Step

Ed Wilson
978-0-7356-2395-8

Microsoft VBScript
Step by Step

Ed Wilson
978-0-7356-2297-5

Microsoft Windows
Scripting with WMI:
Self-Paced Learning Guide

Ed Wilson
978-0-7356-2231-9

Advanced VBScript for Microsoft
Windows Administrators

Don Jones and Jeffery Hicks
978-0-7356-2244-9

RELATED TITLES



Microsoft Office
SharePoint® Server
2007 *Administrator's
Companion*
Bill English with the
Microsoft SharePoint
Community Experts
978-0-7356-2282-1



Microsoft Windows
Security
Resource Kit
Second Edition
Ben Smith and Brian
Komar with the
Microsoft Security
Team
978-0-7356-2174-9



Microsoft Windows
Small Business
Server 2003 R2
*Administrator's
Companion*
Charlie Russel and
Sharon Crawford
978-0-7356-2280-7



Microsoft Internet
Security and
Acceleration (ISA)
Server 2004
*Administrator's Pocket
Consultant*
Bud Ratliff and Jason
Ballard with the Microsoft
ISA Server Team
978-0-7356-2188-6


2007 Microsoft® Office System Resources for Developers and Administrators



Microsoft Office SharePoint® Server 2007 Administrator's Companion

Bill English with the Microsoft SharePoint Community Experts
ISBN 9780735622821

Get your mission-critical collaboration and information management systems up and running. This comprehensive, single-volume reference details features and capabilities of SharePoint Server 2007. It delivers easy-to-follow procedures, practical workarounds, and key troubleshooting tactics—for on-the-job results.



Microsoft Windows SharePoint Services Version 3.0 Inside Out

Errin O'Connor
ISBN 9780735623231


Conquer Microsoft Windows SharePoint Services—from the inside out! This ultimate, in-depth reference packs hundreds of time-saving solutions, troubleshooting tips, and workarounds. You're beyond the basics, so now learn how the experts tackle information sharing and team collaboration—and challenge yourself to new levels of mastery!



Microsoft SharePoint Products and Technologies Administrator's Pocket Consultant

Ben Curry
ISBN 9780735623828

Portable and precise, this pocket-sized guide delivers immediate answers for the day-to-day administration of SharePoint Products and Technologies. Featuring easy-to-scan tables, step-by-step instructions, and handy lists, this book offers the straightforward information you need to get the job done—whether you're at your desk or in the field!



Inside Microsoft Windows® SharePoint Services Version 3

Ted Pattison and Daniel Larson
ISBN 9780735623200

Get in-depth insights on Microsoft Windows SharePoint Services with this hands-on guide. You get a bottom-up view of the platform architecture, code samples, and task-oriented guidance for developing custom applications with Microsoft Visual Studio® 2005 and Collaborative Application Markup Language (CAML).

Inside Microsoft Office SharePoint Server 2007

Patrick Tisseghem
ISBN 9780735623682

Dig deep—and master the intricacies of Office SharePoint Server 2007. A bottom-up view of the platform architecture shows you how to manage and customize key components and how to integrate with Office programs—helping you create custom enterprise content management solutions.

Microsoft Office Communications Server 2007 Resource Kit

Microsoft Office Communications Server Team
ISBN 9780735624061

Your definitive reference to Office Communications Server 2007—direct from the experts who know the technology best. This comprehensive guide offers in-depth technical information and best practices for planning, designing, deploying, managing, and optimizing your systems. Includes a toolkit of valuable resources on CD.

Programming Applications for Microsoft Office Outlook® 2007

Randy Byrne and Ryan Gregg
ISBN 9780735622494

Microsoft Office Visio® 2007 Programming Step by Step

David A. Edson
ISBN 9780735623798

See more resources at microsoft.com/mspress and microsoft.com/learning

Microsoft Press® products are available worldwide wherever quality computer books are sold. For more information, contact your bookseller, computer retailer, software reseller, or local Microsoft Sales Office, or visit our Web site at microsoft.com/mspress. To locate a source near you, or to order directly, call 1-800-MSPRESS in the United States. (In Canada, call 1-800-268-2222.)

Microsoft®
Press

Windows Vista™ Resources for Administrators



Windows Vista Administrator's Pocket Consultant

William Stanek
ISBN 9780735622968

Portable and precise, this pocket-sized guide delivers immediate answers for the day-to-day administration of Windows Vista. Featuring easy-to-scan tables, step-by-step instructions, and handy lists, this book offers the straightforward information you need to solve problems and get the job done—whether you're at your desk or in the field!



Windows Vista Resource Kit

Mitch Tulloch, Tony Northrup, Jerry Honeycutt, Ed Wilson, Ralph Ramos, and the Windows Vista Team
ISBN 9780735622838

Get the definitive reference for deploying, configuring, and supporting Windows Vista—from the experts who know the technology best. This guide offers in-depth, comprehensive technical guidance on automating deployment; implementing security enhancements; administering group policy, files folders, and programs; and troubleshooting. Includes an essential toolkit of resources on DVD.



MCTS Self-Paced Training Kit (Exam 70-620): Configuring Windows Vista Client

Ian McLean and Orin Thomas
ISBN 9780735623903

Get in-depth preparation plus practice for Exam 70-620, the required exam for the new Microsoft Certified Technology Specialist (MCTS): Windows Vista Client certification. This 2-in-1 kit focuses on installing client software and configuring system settings, security features, network connectivity, media applications, and mobile devices. Ace your exam prep—and build real-world job skills—with lessons, practice tests, evaluation software, and more.

MCITP Self-Paced Training Kit (Exam 70-622): Installing, Maintaining, Supporting, and Troubleshooting Applications on the Windows Vista Client – Enterprise

Tony Northrup and J.C. Mackin
ISBN 9780735624085

Maximize your performance on Exam 70-622, the required exam for the new Microsoft® Certified IT Professional (MCITP): Enterprise Support Technician certification. Comprehensive and in-depth, this 2-in-1 kit covers managing security, configuring networking, and optimizing performance for Windows Vista clients in an enterprise environment. Ace your exam prep—and build real-world job skills—with lessons, practice tests, evaluation software, and more.

MCITP Self-Paced Training Kit (Exam 70-623): Installing, Maintaining, Supporting, and Troubleshooting Applications on the Windows Vista Client – Consumer

Anil Desai with Chris McCain of GrandMasters
ISBN 9780735624238

Get the 2-in-1 training kit for Exam 70-623, the required exam for the new Microsoft Certified IT Professional (MCITP): Consumer Support Technician certification. This comprehensive kit focuses on supporting Windows Vista clients for consumer PCs and devices, including configuring security settings, networking, troubleshooting, and removing malware. Ace your exam prep—and build real-world job skills—with lessons, practice tests, evaluation software, and more.

See more resources at microsoft.com/mspress and microsoft.com/learning

Microsoft Press® products are available worldwide wherever quality computer books are sold. For more information, contact your bookseller, computer retailer, software reseller, or local Microsoft Sales Office, or visit our Web site at microsoft.com/mspress. To locate a source near you, or to order directly, call 1-800-MSPRESS in the United States. (In Canada, call 1-800-268-2222.)

Microsoft®
Press

Table of Contents

<i>Introduction</i>	<i>xxi</i>
---------------------------	------------

Part I

Introduction

1 Overview of Microsoft Exchange Server 2007	3
What Is Exchange Server?	3
Editions of Exchange Server 2007	4
Exchange Server 2007 Standard Edition	4
Exchange Server 2007 Enterprise Edition	5
Understanding Basic Concepts	5
Messaging Systems	5
The Organization of an Exchange Environment	8
Exchange Server Storage	11
What's New in Exchange Server 2007	13
Active Directory Site Routing	14
Split Permissions Model	14
Exchange Server 2007 Setup Wizard	14
Exchange Management	14
Exchange Server Roles	15
Unified Messaging	15
Messaging Policy and Compliance	15
Anti-Spam and Antivirus	15
64-Bit Architecture	16
Outlook Web Access	16
Summary	17

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

Microsoft is interested in hearing your feedback about this publication so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit: www.microsoft.com/learning/booksurvey/

2 Active Directory for Exchange Administrators	19
Brief Overview of Active Directory	19
Directory Structure in Active Directory	20
Logical Structure of Active Directory	20
Groups	25
Other Active Directory Components	27
Naming Partitions	27
Sites	28
Location Service Providers	28
Global Catalog Servers	28
Client Authentication	29
Active Directory Names	30
Exchange Server 2007 and Active Directory	31
Exchange Server 2007 and Active Directory Site Topology	31
Storing Exchange Server 2007 Data in Active Directory	33
Exchange Server 2007 and Forest Boundaries	36
Configuration Partition and Directory Data	37
DNS Configuration	37
Summary	38
3 Exchange Server 2007 Architecture	39
The Role of Exchange Server 2007 Roles	39
Mailbox Server Role	40
Client Access Server Role	41
Hub Transport Server Role	42
Unified Messaging Server Role	43
Edge Transport Server Role	44
Storage Design Goals in Exchange Server 2007	45
Stores and Storage Groups	46
Increased User Support	48
Individual Backup and Restore	49
Database File Structure	50
On-Demand Content Conversion	50
Single-Instance Message Store	50

Data Recovery and Transaction Logs	51
The Extensible Storage Engine	51
Transaction Log Files	57
The Web Folder Client	62
Public Folders	63
Multiple Public Folder Trees	64
Indexing	64
Index Catalogs	66
Index Size	67
Exchange Server Storage Design	67
Supported Storage Technologies	67
Choosing a RAID Level	68
Planning for Disk Space	69
Logical Unit Number (LUN) Layout	70
Other Storage Notes	71
Testing Your Storage Architecture	72
Transport Architecture	73
SMTP Connectors	73
Creating SMTP Connectors	74
Message Routing	74
Message Transport Scenarios	78
Transport Protocols	81
Summary	83

Part II

Planning Your Deployment

4 Assessing Needs	87
Defining User Needs	88
Messaging	88
Public Folders	89
Connections to Other Systems	90
Remote Access	91
Custom Applications	91
Training and Support Services	91

Assessing Current Resources	92
Defining Your Geographic Profile	92
Defining Your Software Environment	92
Defining Your Network Topology	93
Defining Your Active Directory Model	96
Defining Administrative Needs	98
Summary	98
5 Planning for Deployment	99
Planning the Organization	99
Establishing a Naming Convention	99
Planning Public Folders	102
Planning Gateways	103
Planning Servers	104
Disk Considerations	104
Processor Considerations	106
Memory Considerations	108
Network Considerations	108
Ways to Add Fault Tolerance	109
Summary	109

Part III

Installation and Deployment

6 Installing Exchange Server 2007	113
Preparing for the Installation	114
Gathering Information	114
Verifying Hardware Requirements	116
Getting Service Packs	117
Defining the Role of Your Server	118
Optimizing Hardware through Configuration	119
Verifying System Requirements	120
Creating the Exchange Administrator's Account	124
Playing It Safe	125
Performing the Installation	125
Preparing the Active Directory Environment	126

Installing Exchange Server 2007 in a New Organization	128
The Role of Roles	130
Installing in an Existing Organization	135
Verifying Your Installation	136
Finalizing Exchange Server 2007 Deployment	138
Keeping Exchange Healthy	142
Summary	143
7 Coexisting with Previous Versions of Exchange Server	145
Chapter Background	146
Terminology	146
Exchange Server 2007 Coexistence Deployment Considerations	147
Exchange Server 2003 Native Mode	147
Automatic Coexistence Tasks	148
Global Settings	149
Installing Exchange Server 2007 into an Existing Exchange Server 2003 Organization	150
Coexistence Administration Issues	152
Creating Additional Routing Group Connectors	153
Coexistence Issue: Version-Specific Administration	155
SMTP Connectors and Internet E-Mail	156
Handling Internet E-Mail	157
Adding an SMTP Connector to Your Legacy Exchange Organization ..	158
Public Folders	162
Public Folder Replication	163
Handling Public Folder Referrals	164
Administering Public Folders	166
Recipient Update Service	167
Complete Coexistence Notes	168
Summary	181
8 Transitioning to Exchange Server 2007	183
The Example Scenario	184
Transition Options	185
Transition Limitations	185
Move Internet Mail to Exchange Server 2007	186

Allow Mail to Flow to the Internet	187
Allow Incoming Mail from the Internet	190
Moving Mailboxes to Exchange Server 2007	192
The Decommissioning Process	199
Re-Home Client Services	200
Remove SMTP Connectors from Your Legacy Exchange Organization	200
Re-Home Public Folders	201
Move the Offline Address Book to Exchange Server 2007	203
Move the Recipient Update Service to Exchange Server 2007	204
Remove Legacy Connectors	205
Uninstall Exchange from Legacy Exchange Servers	207
Remove Legacy Exchange Routing Groups	207
Summary	208
9 High Availability in Exchange Server 2007	209
Continuous Replication and Transaction Logs	210
Local Continuous Replication	213
Preparing for LCR	214
Enabling Local Continuous Replication	215
Cluster Continuous Replication	223
CCR Terminology	224
Preparing for CCR	226
Enabling Continuous Cluster Replication	227
Establishing the Cluster	229
Configure the MNS Quorum to Use the File Share Witness	233
Installing Exchange Server 2007 on Your Cluster	233
Verifying the Status of Your CCR	236
Verifying That a Server Can Handle a Failover	236
Configuring the Transport Dumpster	237
Closing Thoughts on CCR	238
Single Copy Clusters	239
Summary	242

Part IV

Management

10 Managing Exchange Server 2007	245
Microsoft Management Console	246
The MMC User Interface	246
How MMC Works	249
Using the Exchange Management Console	251
Major Areas of the Exchange Management Console	252
Examining the Exchange Hierarchy	254
Using the Exchange Management Shell	260
Understanding Cmdlets	262
Getting Help	263
Summary	265
11 Creating and Managing Recipients.....	267
Understanding Recipient Types	268
Users	269
Mailbox Users.....	269
Mail-Enabled Users	286
Mailbox Resources	288
Mail Contacts	289
Creating a Mail Contact	289
Configuring a Mail Contact	291
Distribution Groups	291
Creating a Distribution Group.....	292
Configuring a Group.....	293
Creating Dynamic Distribution Groups	296
Filtering Recipients	297
Templates	298
Address Lists.....	299
Summary	302

12 Using Public Folders	303
Understanding Public Folder Storage	304
Using Public Folders in Outlook 2007	305
Creating a Public Folder in Outlook	305
Managing Public Folders in Outlook	305
Managing Public Folder Databases in the Exchange Management Console	307
Creating a New Public Folder Database	308
Removing a Public Folder Database	309
Creating and Managing Public Folders in the Exchange Management Shell	311
Creating a Public Folder	311
Removing a Public Folder	311
Getting Information about a Public Folder	312
Managing Settings for a Public Folder	312
Summary	314
13 Creating and Managing Storage Groups	315
Review of Exchange Server 2007 Storage Architecture	315
Benefits of Using Storage Groups	317
Increased User Support	318
Individual Backup and Restore	319
Hosting of Multiple Businesses	319
Support for Special Mailboxes	320
Planning Storage Groups	320
Planning for Disk Space	321
Planning for Multiple Storage Groups	324
Planning for Backup and Restore Throughput	325
Managing Storage Groups	326
Creating Storage Groups	326
Modifying Storage Group Configuration	329
Removing Storage Groups	332
Managing Stores	333
Creating a Mailbox Store	333
Modifying Mailbox Database Configuration	335
Summary	343
14 Unified Messaging	345
Unified Messaging Overview	346

Unified Messaging Features	346
Exchange Server 2007 Unified Messaging Objects	348
Creating and Managing Unified Messaging Objects	350
Unified Messaging Dial Plans	350
Unified Messaging Mailbox Policy	357
Unified Messaging IP Gateways	363
Associating Servers with Dial Plans	366
Enabling Unified Messaging for Individual Mailboxes	367
Summary	370

Part V

Maintenance

15 Troubleshooting Exchange Server 2007	373
Using Troubleshooting Tools	373
Using Event Viewer	373
Using Diagnostics Logging	375
Inbox Repair Tool	379
RPinG Utility	380
Eseutil.exe Offline Tool	383
Best Practices Analyzer	385
Mail Flow Troubleshooter	387
Performance Troubleshooter	389
Other Useful Utilities	390
Finding Help	390
Product Documentation	391
Microsoft TechNet	391
Internet Newsgroups	391
Summary	392
16 Disaster Recovery	393
Backup and Restore Technologies	393
The Exchange Database	394
Volume Shadow Copy Service	399
Exchange Streaming Backup API	401
Other Exchange Server Components	405

Backup and Restore Strategies	406
Recovering an Exchange Mailbox Server	410
Recovering an Exchange Mailbox Database	414
Recovering a Single Exchange Mailbox	414
Backing up an Exchange Mailbox Server	416
Backing up an Exchange Mailbox Database	417
Backing up a Single Exchange Mailbox	418
Planning for Corruption	419
Implementing Backup Strategies	420
Operational Best Practices	425
Summary	426
17 Tuning Exchange Server 2007 Performance	427
Understanding How the Performance Snap-in Works	427
Performance Monitoring Concepts	428
Collecting Data with the Performance Snap-In	429
Viewing Collected Data	430
Evaluating the Four Main Subsystems in Windows	431
Evaluating Memory Usage	432
Evaluating Processor Usage	433
Evaluating Disk Usage	434
Evaluating Network Usage	436
Using the Performance Snap-in to Tune Exchange Server 2007	437
SMTP System Monitor Counters	437
Outlook Web Access	438
Unified Messaging Counters	439
Using Other Exchange Performance Tools	442
Microsoft Exchange Server Jetstress Tool	442
Exchange Load Generator	444
Summary	445

Part VI

Security

18 Security Policies and Exchange Server 2007	449
Why Are Information Security Policies Important?	450

Information Security Policies and Electronic Policies	452
Information Security Policies for Exchange Server 2007	453
Password Policies	453
Logon Policies	454
Acceptable Use Policies	455
Computer Viruses, Trojans, and Worms	456
Schema Extensions by Exchange Server 2007	457
Data Security	459
Legal Exposure to Unwanted E-Mail Content	460
Backing Up and Archiving Exchange Databases	461
E-Mail Integrity	462
Miscellaneous Elements to Consider	463
Related Resources	464
Summary	465
19 Exchange Server Security Basics	467
The Scope of Security	468
Motivations of a Criminal Hacker	469
How Hackers Work	470
Physical Security	474
Administrative Security	474
The Built-in Exchange Administrative Groups	475
The Add Exchange Administrator Wizard	477
SMTP Security	480
Computer Viruses	485
What Is a Virus?	485
Trojans	486
Worms	486
Junk E-Mail	487
Security Tools Provided by Microsoft	488
Summary	489
20 Antivirus and Anti-Spam	491
The Edge Transport Server at a Glance	491
Edge Transport Server Deployment	493
Verify the Edge Transport Server's DNS Suffix	493

Configure Firewalls to Pass Edge Traffic	494
Install Active Directory Application Mode	495
Install the Exchange Server 2007 Edge Transport Server Role	495
Subscribe the Edge Transport Server to the Exchange Server 2007 Organization	497
Managing Anti-Spam Features	502
Content Filtering	502
Connection Filtering: IP Allow List	506
Connection Filtering: IP Allow List Providers	508
Connection Filtering: IP Block List	509
Connection Filtering: IP Block List Providers	511
Recipient Filtering	514
Sender Filtering	515
Sender ID	517
Attachment Filtering	520
Managing Antivirus with Microsoft Forefront Security for Exchange Server ...	524
About Microsoft Forefront Security for Exchange Server	525
Installing Microsoft Forefront Security for Exchange Server	525
Managing Microsoft Forefront Security for Exchange Server	527
Other Microsoft Forefront Security for Exchange Server Benefits	529
Summary	530

21 Securing Exchange Server 2007 Messages 531

Windows Server 2003 Security Protocols	531
Understanding the Public Key Infrastructure in Windows Server 2003	532
Encryption and Keys	532
Encryption Schemes	533
Certificate Services in Windows Server 2003	534
Managing the Public Key Infrastructure	540
Installing and Configuring Certificate Services	540
Installing Web Enrollment Support	545
Using the Web Enrollment Pages	546
Viewing Information About Certificates	551
Securing Messaging in Outlook 2007	555
Initially Trusting a Certificate	556

Encryption and Outlook 2007	556
Digital Signatures and Outlook 2007	557
S/MIME and Outlook 2007	557
Configuring Outlook 2007 for Secure Messaging	558
Installing Exchange Certificate Templates	560
Understanding How Exchange Server 2007 Integrates with Windows Server 2003 Security	561
Summary	564

Part VII

Clients

22 Overview of Exchange Clients	567
Microsoft Office Outlook 2007	568
Windows Mail and Microsoft Outlook Express	570
Outlook Web Access	572
Standard Internet E-Mail Clients	573
Non-Windows Platforms	573
UNIX Clients	574
Macintosh Clients	574
Choosing a Client for Exchange Server	574
Summary	575
23 Deploying Microsoft Office Outlook 2007	577
Installing Outlook 2007	577
Standard Outlook Installation	578
Installing Outlook 2007 by Using the Office Customization Tool	579
Supporting Outlook 2007	579
Using Cached Exchange Mode	580
Enabling Multiple Users in Outlook 2007	586
Outlook Anywhere	590
Summary	593
24 Supporting Outlook Web Access	595
Features of OWA	595
Deploying OWA	596

Single-Server Scenario	596
Multiserver Scenario	597
ISA Server 2006 and OWA	600
Authentication Options	601
Configuring OWA Properties and Features	610
Managing Access to UNC Shares and SharePoint	
Document Repositories	611
OWA Segmentation	617
OWA User Features	622
Summary	624
25 Supporting Other Clients	625
Post Office Protocol Version 3	625
Enabling POP3	627
Administering POP3	627
Internet Messaging Access Protocol 4	632
Enabling IMAP4	633
Administering IMAP4	634
POP3/IMAP4 Considerations	639
Summary	640

Part VIII

Appendices

A Default Directory Structure for Exchange Server 2007	643
B Delivery Status Notification Codes	645
C Default Log File Locations	649
D Default Diagnostic Logging Levels for Exchange Processes	651
<i>Glossary</i>	<i>657</i>
<i>Index</i>	<i>669</i>

Chapter 2

Active Directory for Exchange Administrators

Brief Overview of Active Directory	19
Other Active Directory Components.....	27
Exchange Server 2007 and Active Directory	31
DNS Configuration.....	37
Summary	38

In the Chapter 1, “Overview of Microsoft Exchange Server 2007,” you learned about some of the basic components of an Exchange organization. This chapter builds on that knowledge by describing how Exchange Server 2007 integrates with Microsoft Windows Server 2003 and how it uses the services in Windows Server 2003 to its advantage. It begins with a brief overview of the Windows Server 2003 Active Directory service and finishes by describing how Exchange Server 2007 uses Active Directory and discussing some of the more important Internet information protocols.

Brief Overview of Active Directory

A full explanation of Active Directory is outside the scope of this book, but a brief overview is warranted. Because Exchange Server 2007 is heavily dependent on the underlying network operating system, it is important to have a basic understanding of Windows Server 2003 Active Directory.

More Info For a more thorough discussion of Active Directory and the other concepts discussed in this chapter, see *Microsoft Windows Server 2003 Administrator's Companion*, Second Edition, by Charlie Russel, Sharon Crawford, and Jason Gerend (Microsoft Press, 2006).

Directory Structure in Active Directory

Before beginning the discussion on what Active Directory is, you should first understand what a directory is. As an analogy, think of a generic file system. Perhaps in this file system, you have a C drive, and on that drive, you have a root folder named Memos. Under C:\Memos, you have a folder for each of the 12 months of the year, so you would find a folder in the structure named July. Under C:\Memos\July, you have a folder named Departments; the full pathname to Departments is C:\Memos\July\Departments. This is a hierarchy of folders in a file system.

A directory is no different from a folder list, except that the hierarchy consists not of folders but of objects. An *object* is an entity that is described by a distinct, named set of attributes. Instead of using Windows Explorer to search through this hierarchy of objects, you'll be using a protocol designed to search a directory, called the *Lightweight Directory Access Protocol* (LDAP).

Note The original protocol for accessing a directory was called Directory Access Protocol (DAP), but it had a high overhead and tended to be slow. *Lightweight Directory Access Protocol* (LDAP) is an improved version that is faster and requires less overhead.

With Active Directory, Microsoft has made significant improvements to the directory concept, such as dynamic DNS. The “Active” in Active Directory describes the flexibility and extensibility that have been built into Microsoft’s directory service.

Logical Structure of Active Directory

The components that form the logical structure of Active Directory include domains, organizational units, trees, and forests.

Domains

A *domain* is the core unit in Active Directory and is made up of a collection of computers that share a common directory database. The computers that share this common directory database are called domain controllers. A *domain controller* is a Windows Server 2003 server that has Active Directory installed. It can authenticate users for its own domain. Each domain controller holds a complete replica of the domain naming partition for the domain to which it belongs and a complete replica of the configuration and schema naming partitions for the forest. Dcpromo.exe is the utility used to promote a Windows Server 2003 server to a domain controller. Partitions are discussed later in this chapter.

All Active Directory domain names are identified by a DNS name as well as by a NetBIOS name. The following is an example of the two types of names:

- DNS-style domain name: contoso.com
- NetBIOS name: CONTOSO

Generally, the NetBIOS name is the same as the first naming component in the DNS name. However, a NetBIOS name can be only 15 characters in length, whereas each name in the DNS naming convention can have up to 64 characters. During installation, both names can be configured to meet your needs. In the initial release of Windows Server 2003, Active Directory names could be changed. Although there are tools available that allow you to change a domain name, it is a complex undertaking. It is better to be careful when initially creating your naming scheme.

More Info To read more about and download the Windows Server 2003 Active Directory Domain Rename Tools, visit <http://www.microsoft.com/technet/downloads/winsrvr/domainrename.mspx>.

The domain is also a security boundary in Active Directory. Administrators in a domain have the permissions and rights to perform administrative functions in that domain. However, because each domain has its own security, administrators must be given explicit permissions to perform administrative tasks in other domains. Members of the Enterprise Admins group have rights to perform administrative tasks in all domains across the forest. Hence, you can have domain administrators and a higher level of administration from the Enterprise administrators.

A Windows Server 2003 Active Directory domain can be in either mixed mode or native mode. The default installation is mixed mode. In mixed mode, a Windows Server 2003 domain controller acts like a Microsoft Windows NT 4 domain controller. Active Directory domains in mixed mode have the same limitations on the security accounts database as Windows NT 4 domain controllers. For example, in mixed mode, the size of the directory is limited to 40,000 objects, the same restriction imposed by Windows NT 4. These limitations allow Windows NT 4 backup domain controllers to exist on the network and connect to and synchronize with the Windows Server 2003 domain controllers.

Note Exchange Server 2007 requires that an Active Directory be in native mode prior to installation. You'll learn more about that in Chapter 6, "Installing Exchange Server 2007."

The PDC Emulator is one of the five Flexible Single Master Operation (FSMO) roles that make Windows Server 2003 look like a Windows NT 4 PDC. Only one Windows Server 2003 domain controller can act as the PDC Emulator. By default, the PDC Emulator role, like all other FSMO roles, is installed on one domain controller in each domain – by default, on the first domain controller of each domain. (FSMO roles are discussed in a moment.) To run Windows Server 2003 in native mode, you must not have any reason or desire to connect to a Windows NT 4 backup domain controller. In other words, when you decide to run Windows Server 2003 in native mode, you won't be able to use a Windows NT backup domain controller again on your network, and no applications running on your network will be able to use Windows NT to operate. The switch to native mode is a one-time, one-way switch and is irreversible. Native mode allows your Windows Server 2003 domain controllers to have millions of objects per domain. In addition, native mode allows the nesting of groups, something that is advantageous if you anticipate large distribution groups in Exchange Server 2007.

A Windows Server 2003 network running in native mode can accommodate Windows NT 4 stand-alone and member servers. Windows NT 4 workstations must be upgraded to Windows 2000 Professional, Windows XP Professional, or Windows Vista to participate in Active Directory, or you must install the Directory Service Client. Windows Server 2003 implements Active Directory in a multimaster model because objects in Active Directory can be modified on any domain controller, which accounts for the emphasis on directory replication between domain controllers. However, some roles are either too sensitive to security issues or too impractical to perform in a multimaster model because of potential conflicts that could arise from the replication traffic. An understanding of these roles is important: if a domain controller that is performing a particular role becomes unavailable, the function it performs is not available in Active Directory. These roles are schema master, domain naming master, relative identifier master, PDC emulator, and infrastructure master.

Schema Master

The *schema* is the set of object classes (such as users and groups) and their attributes (such as full name and phone number) that form Active Directory. The schema master controls all aspects of updates and modifications to the schema. To update the schema, you must have access to the schema master. There can be only one schema master in the forest at any given time.

Domain Naming Master

The domain naming master controls the addition and removal of domains in the forest. This is the only domain controller from which you can create or delete a domain. There can be only one domain naming master in the forest at any given time.

Relative Identifier Master

The relative identifier (RID) master allocates sequences of RIDs to each of the domain controllers in its domain. Whereas the schema master and domain naming master perform forestwide functions, one RID master is assigned per domain. Because each domain controller can create objects in Active Directory, the RID master allocates to each domain controller a pool of 500 RIDs from which to draw when creating the object. When a domain controller has used more than 400 RIDs, the RID master gives it another batch of 500 RIDs.

Whenever a new user, group, or computer object is created, the object inherits the security identifier (SID) of the domain. The RID is appended to the end of the domain SID to make up a unique SID for the object. In addition, when an object is moved from one domain to another, its SID changes, because it receives a new SID (made up of both the domain SID and the RID) in the destination domain. By allowing only the RID master to move objects between domains, Windows Server 2003 ensures SID uniqueness, even across domains. Objects maintain a SID history for security access to resources.

PDC Emulator

Each domain in the forest must have one domain controller that acts as the PDC emulator. If Active Directory is running in mixed mode with Windows NT 4 domain controllers on the same network, the PDC emulator is responsible for synchronizing password changes and security account updates between the Windows NT 4 servers and the Windows Server 2003 servers. Moreover, the PDC emulator appears to downlevel clients, such as Windows 95, Windows 98, and Windows NT 4, as the PDC of the domain. It functions as the domain master browser, is responsible for replication services to the BDCs, and performs directory writes to the Windows NT 4 domain security database.

In native mode, the PDC emulator receives the urgent updates to the Active Directory security accounts database, such as password changes and account lockout modifications. These urgent changes to user accounts are immediately replicated to the PDC emulator, no matter where they are changed in the domain. If a logon authentication fails at a domain controller, the credentials are first passed to the PDC emulator for authentication before the logon request is rejected.

Infrastructure Master

The infrastructure master is responsible for tracking group-to-user references whenever the user and the group are not members of the same domain. The object that resides in the remote domain is referenced by its GUID and SID. If an object is moved from one domain to another, it receives a new SID, and the infrastructure master replicates these changes to other infrastructure masters in other domains.

Organizational Units

An *organizational unit* (OU) is a container object that is used to organize other objects within a domain. An OU can contain user accounts, printers, groups, computers, and other OUs.

More Info The design of Active Directory is based on the X.500 standard, which can be procured from www.itu.org. The standard is rather short—around 29 pages—but reading it will give you a great background for understanding Active Directory and, for that matter, Novell Directory Services.

OUs are strictly for administrative purposes and convenience. They are transparent to the user and have no bearing on the user's ability to access network resources. OUs can be used to create departmental or geographical boundaries. They can also be used to delegate administrative authority to users for particular tasks. For example, you can create an OU for all of your printers and then assign full control over the printers to your printer administrator.

OUs can also be used to limit administrative control. For example, you can give your help desk support personnel the permission to change the password on all user objects in an OU without giving them permissions to modify any other attributes of the user object, such as group membership or names.

Because an Active Directory domain can hold millions of objects, upgrading to Windows Server 2003 allows companies to convert from a multiple-domain model to a single-domain model and then use organizational units to delegate administrative control over resources.

Trees and Forests

The first Windows Server 2003 domain that you create is the root domain, which contains the configuration and schema for the forest. You add domains to the root domain to form the tree. As Figure 2-1 illustrates, a *tree* is a hierarchical grouping of Windows Server 2003 domains that share a contiguous namespace. A *contiguous namespace* is one that uses the same root name when naming additional domains in the tree.

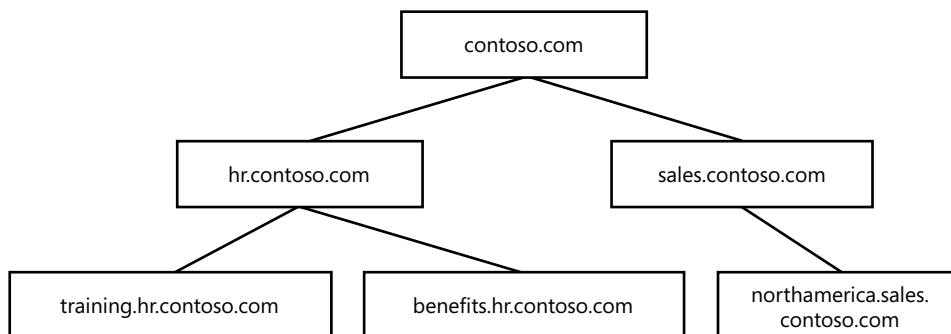


Figure 2-1 Fictitious tree of contoso.com

A collection of trees that does not share a contiguous namespace can be placed in the same forest. They then share a common configuration, schema, and Global Catalog (GC). By default, the name of the root domain becomes the name of the forest, even though other trees will not share the same name as the root domain.

Even though they don't share the same name, transitive trust relationships are automatically established between the root domain servers in each tree, as long as they are members of the same forest. Figure 2-2 shows two trees—contoso.com and litwareinc.com—in the same forest.




Figure 2-2 Forest consisting of contoso.com and litwareinc.com

The schema and configuration partitions for Active Directory are replicated to all domain controllers in each domain. Whereas a domain represents a boundary for security and the logical grouping of objects, a forest represents the boundary for Active Directory and the Exchange Server 2007 organization.

In addition, other domain names cannot be represented above the first domain name. For example, if your root domain name is sales.contoso.com, you can never install a domain named contoso.com in the same forest. You can join other domain names to the forest, such as litwareinc.com, as long as they are in a different namespace.

Groups

In Windows Server 2003, groups are used to reduce administrative effort and to enable the management of many user accounts simultaneously. Windows Server 2003 uses groups to reduce the number of objects that require direct administration.

There are basically two kinds of groups in Windows Server 2003. Each has its own advantages and restrictions that you must take into account when using them. Exchange Server 2007 uses both kinds of groups from Windows Server 2003:

- **Security groups** Security groups host security principles within Active Directory. They are used to group users or computers for the purpose of reducing the points of administration and providing permissions to network resources.
- **Distribution groups** Distribution groups are meant to perform distribution functions. You cannot use them to assign permissions to network resources.

Global Groups

Global groups, in mixed mode, can contain users only from the domain in which they are hosted. In native mode, they can contain users and global groups from the local domain in which they were created. However, they can be used to assign permissions to resources in any domain. Global groups can contain users, computers, and global groups from the local domain. They can be members of any other type of group.

Typically, you'll use global groups for administering user membership that has permissions to a network resource. The group itself is replicated as part of the Global Catalog, but its membership is not. This restriction means that adding user accounts to or removing user accounts from a global group does not trigger a new replication of the Global Catalog. Global groups can be converted to universal groups (discussed shortly) as long as the global groups do not contain other global groups and the domain is in native mode.

Domain Local Groups

Domain local groups in native mode can contain other domain local groups, users, global groups, and universal groups from any domain in the forest, but they can be granted permissions only in the domain in which they reside. In mixed mode, they can contain only user and global group accounts.

You grant permissions to domain local groups only for objects in the local domain. The existence of the domain local group is replicated to the Global Catalog server, but its membership is not replicated. Domain local groups are flexible in that you can use any other security principle inside the domain local group (when running in native mode) to reduce administrative effort. You can convert a domain local group to a universal group in native mode as long as it does not contain other domain local groups.

Universal Groups

Universal groups can contain users, global groups, and other universal groups from any Windows Server 2003 domain in the forest. The domain must be operating in native mode to create security groups with universal scope. You can grant permissions to resources anywhere in the forest to a universal group.

Universal group membership must be determined at the time of logon. Because the scope of the universal group is universal, this group is propagated through the Global Catalog. Hence, not only is the group itself propagated in the Global Catalog, but its membership is propagated as well. A universal group with a large membership generates additional replication overhead if the membership changes. Universal groups as security groups are available only in native mode. Table 2-1 summarizes group membership rules.

Table 2-1 Comparison of Various Types of Groups

Group scope	In mixed mode can contain	In native mode can contain	Can be a member of	Can be granted permissions for
Domain local	User accounts and global groups from any domain	User accounts, global groups, and universal groups from any domain in the forest, and domain local groups from the same domain	Domain local groups in the same domain	The domain in which the domain local group exists
Global	User accounts	User accounts and global groups from the same domain	Universal and domain local groups in any domain and global groups in the same domain	All domains in the forest
Universal	Not applicable	User accounts, global groups, and other universal groups from any domain in the forest	Domain local and universal groups in any domain	All domains in the forest

Other Active Directory Components

Active Directory is a complex system that includes far more than the basic logical structure just described. This section highlights several other components that play a critical role within Active Directory.

Naming Partitions

You can think of Active Directory as being divided into three distinct directories, or partitions: the domain partition, the configuration partition, and the schema partition. Each

partition is a self-contained section of Active Directory that can have its own properties, such as replication configuration and permissions structure. A Windows Server 2003 domain controller always holds three naming partitions in its database file (Ntds.dit). These are the default LDAP paths for these partitions:

- Configuration: cn=configuration,dc=sales,dc=contoso,dc=com
- Schema: cn=schema,cn=configuration,dc=sales,dc=contoso,dc=com
- Domain: dc=sales,dc=contoso,dc=com

In a multidomain structure, domain controllers belong to different domains. These servers share a common configuration and schema naming partition but have a unique domain naming partition. Exchange Server 2007 stores most of its information in the configuration naming partition, which is replicated throughout the forest.

Sites

A *site* within Active Directory is a collection of Internet protocol (IP) subnets that enjoy permanent, high-bandwidth connectivity. Active Directory assumes that all computers in the same site have permanent, high-speed connectivity with one another. Sites tend to map to the physical structure of your network: slow WAN links will be considered outside your sites, and high-speed links will form your sites.

Site and domain topologies are not dependent upon each other; a single domain can span multiple sites, or multiple domains can be located in a single site. Because the bandwidth between sites is assumed to be slow or unreliable, it stands to reason that some type of connector is needed to connect the sites. That connector is called a *site link*.

Site links are built manually by the administrator and form the physical topology of the network. To create replication paths between domain controllers across the site links (as well as between domain controllers within the same site), Windows Server 2003 employs the knowledge consistency checker (KCC), which runs automatically but can be configured manually, if necessary. The KCC creates *connection objects* on each domain controller in the configuration naming partition; these connection objects form the overall replication topology over which Active Directory information can be replicated. The KCC is a service that runs on each domain controller to create the connection objects for that domain controller.

Location Service Providers

In Windows Server 2003, DNS provides the role of server service locator, helping the client find the services it needs on the network. Dynamic DNS is supplied with Windows Server 2003 and is a standard part of the Active Directory installation. With dynamic

DNS, clients query DNS service (SRV) records to locate services on the network and can also update the DNS records if their own location changes.

Global Catalog Servers

In a multidomain environment, it is reasonable to assume that some users will need access to objects outside of their own domains. For example, a user in domain A might need access to a color printer located in domain B. Because domain controllers maintain only a replica of objects in their own domain, a special service is needed in the forest to gain access to objects located in remote domains. The Global Catalog server performs this function. This server holds a replica of all objects in the forest, with a limited set of attributes for those objects. The schema defines which attributes are listed for each object in the Global Catalog. The Global Catalog is not a separate file; it is instead held inside the NTDS.DIT file. The GC will be roughly 40 percent of the size of your active directory, or the size of the NTDS.DIT file on a non-GC domain controller.

Note By default, there is only one Global Catalog server in the entire forest, and that is the first domain controller installed in the first domain of the first tree. All other Global Catalog servers need to be configured manually. You can do this by opening the Active Directory Sites And Services snap-in, navigating to the NTDS settings on the server on which you want to install this service, right-clicking NTDS Settings, choosing Properties, and selecting the Global Catalog Server check box.

In addition to users needing access to services outside their domain, some applications need access to a forestwide listing of objects. Exchange Server 2007 is one of those applications. For example, a user might want to browse the Global Address List, which is generated by the Global Catalog server. The Global Catalog server gathers each mail-enabled object into a list and returns this list to the user inside the address book interface.

Even in a single-domain environment, Exchange clients are referred to the Global Catalog server for address book lookups. In this scenario, the default is to refer all those lookups to the root domain controller. You should plan for an increase in network traffic between Global Catalogs and Exchange Server 2007 servers. This increase may be sizeable if you are taking full advantage of all the new features and roles in Exchange Server 2007.

It is helpful to note that a Global Catalog server passes back different attributes depending on the TCP port used for the query. For example, a query to port 389 (the default LDAP port) allows a client to search for objects only within the home domain, with the full set of attributes for the object being returned. In contrast, a query over port 3268 allows a client to search for domain objects from all domains in the forest, including the home domain of the Global Catalog server. However, a query over this port returns only

a subset of the attributes available, even if the object is in the home domain of the Global Catalog server.

Client Authentication

When a client attempts to log on to the domain, it queries DNS SRV records to locate a domain controller. DNS attempts to match the client's IP address to an Active Directory site and then returns a list of domain controllers that can authenticate the client. The client chooses a domain controller at random from the list and then pings it before sending the logon request. In native mode, the authenticating domain controller passes the client's credentials to the local Global Catalog server so that the Global Catalog can enumerate universal security group access.

Active Directory Names

Both users and applications are affected by the naming conventions that a directory uses. To locate a network resource, you need to know its name or one of its properties. Active Directory supports many naming schemes for the different formats that can access Active Directory.

Distinguished Name

Each object in the directory has a *distinguished name* (DN) that identifies where the object resides in the overall object hierarchy. For example:

```
cn=dhall,cn=users,dc=contoso,dc=com
```

This example indicates that the user object dhall is in the Users container that is located in the contoso.com domain. If the dhall object is moved to another container, its DN changes to reflect its new position in the hierarchy. Distinguished names are guaranteed to be unique in the forest. You cannot have two objects with the same distinguished name.

Relative Distinguished Name

The *relative distinguished name* of an object is the part of the distinguished name that is an attribute of the object. In the former example, the relative distinguished name of the dhall object is dhall. The relative distinguished name of the parent organizational unit is Users. Active Directory does not allow two objects with the same relative distinguished name under the same parent container.

User Principal Name

The *user principal name* that is generated for each object is in the form username@DNS-domainname. Users can log on with their user principal name, and an administrator can define suffixes for user principal names if desired. User principal names are required to

be unique, but Active Directory does not enforce this requirement. It's best, however, to formulate a naming convention that avoids duplicate user principal names.

Globally Unique Identifier

Some applications require that an object be referred to by an identifier that remains constant. This is achieved by adding an attribute called the *globally unique identifier* (GUID), a 128-bit number that is guaranteed to be unique. A GUID is assigned to an object when it is created, and it never changes, even if the object is moved between containers in the same domain.

Exchange Server 2007 and Active Directory

Exchange Server 2007 is tightly integrated with the Windows Server 2003 Active Directory service. Integration with Windows Server 2003 provides several benefits, including the following:

- **Centralized object management** Administration is unified for Exchange Server 2007 and Windows Server 2003. Directory objects can be managed from one location, with one management tool, and by one team.
- **Simplified security management** Exchange Server 2007 uses the security features of Windows Server 2003, such as the discretionary access control list (DACL). Changes to security principles (such as user or group accounts) apply to data stored in both Exchange Server 2007 and Windows Server 2003 file shares.
- **Simplified creation of distribution lists** Exchange Server 2007 automatically uses Windows Server 2003 security groups as distribution lists, eliminating the need to create a security group for each department and a corresponding distribution group for the same department. Distribution groups can be created in those instances when e-mail distribution is the only desired function of the group.
- **Easier access to directory information** LDAP is the native access protocol for directory information.

Exchange Server 2007 and Active Directory Site Topology

In Exchange 2000 Server and Exchange Server 2003, Exchange requires the configuration of routing groups to govern how messages and other Exchange-related traffic are routed throughout an organization. Exchange Server 2007 no longer uses routing groups, instead relying on the routing technology built into Active Directory site topology.

All computers within a single Active Directory site are considered well connected, with a high-speed, reliable network connection. By default, when you first deploy Active Directory on a network, it creates a single site named (by default) *Default-First-Site-Name*. All server and client computers in the forest are made members of this first site. When you define more than one site, you must define the network subnets that are present on the network and associate each of those subnets with Active Directory sites.

In Active Directory, *IP site links* define the relationship between sites. An IP site link connects two or more Active Directory sites. Each IP site link is associated with a cost that helps dictate how Active Directory should consider using that link in relation to the costs of other available site links. You (or the Active Directory administrator) assign the cost to a link based on relative network speed and available bandwidth compared to other available connections.

Exchange Server 2007 uses the cost assignment for an IP site link to determine the lowest cost route for traffic to follow when multiple paths exist to a destination. The cost of the route is determined by aggregating the cost of all site links in the path. For example, assume that a computer in Site A must communicate with a computer in Site C. Site A is connected by an IP site link with a cost of 10 to Site B. Site B is connected by an IP site link to Site C with a cost of 5. The cost of the full route from Site A to Site C would be 15.

Active Directory clients assume site membership by matching their assigned IP address to a subnet associated with a particular site.

Because Exchange Server 2007 is now a site-aware application, it can determine its own Active Directory site membership and the Active Directory site membership of other servers. All Exchange Server 2007 server roles use site membership to determine which domain controllers and Global Catalog servers to use for processing Active Directory queries. In addition, Exchange Server 2007 also tries first to retrieve information about recipients from directory servers that are in the same site as the Exchange Server 2007 server.

The Exchange Server 2007 server roles use Active Directory site membership information as follows:

- The Mailbox server role uses Active Directory site membership information to determine which Hub Transport servers are located in the same Active Directory site as the Mailbox servers. The Mailbox server submits messages for routing and transport to a Hub Transport server that has the same Active Directory site membership as the Mailbox server. The Hub Transport server performs recipient resolution and queries Active Directory to match an e-mail address to a recipient account. The Hub Transport server delivers the message to the Mailbox server within its same Active Directory site, or it relays the message to another Hub Transport server for delivery to a Mailbox server that is outside the Active Directory site. If there are

no Hub Transport servers in the same Active Directory site as a Mailbox server, mail can't flow to that Mailbox server.

- Active Directory site membership and IP site link information is used to prioritize the list of servers that are used for public folder referrals. Users are directed first to the default public folder database for their mailbox database. If a replica of the public folder being accessed does not exist in the default public folder database, the Mailbox store where the default public folder database resides provides a prioritized referral list of Mailbox servers that hold a replica to the client. Public folder databases that are in the same Active Directory site as the default public folder database are listed first, and additional referral locations are prioritized based on Active Directory site proximity.
- The Unified Messaging (UM) server role uses Active Directory site membership information to determine which Hub Transport servers are located in the same Active Directory site as the Unified Messaging server. The Unified Messaging server submits messages for routing and transport to a Hub Transport server that has the same Active Directory site membership as the Unified Messaging server. The Hub Transport server delivers the message to a Mailbox server within its same Active Directory site, or it relays the message to another Hub Transport server for delivery to a Mailbox server that is outside the Active Directory Site.
- When the Client Access server role receives a user connection request, it queries Active Directory to determine which Mailbox server is hosting the user's mailbox. The Client Access server then retrieves the Active Directory site membership of that Mailbox server. If the Client Access server that received the initial user connection is not located in the same site as the user's Mailbox server, the connection is redirected to a Client Access server in the same site as the Mailbox server.
- Exchange Server 2007 Hub Transport servers retrieve information from Active Directory to determine how mail should be routed inside the organization. If the recipient's mailbox is located on a Mailbox server in the same Active Directory site as the Hub Transport server, the message is delivered directly to that mailbox. If the recipient's mailbox is located on a Mailbox server in a different Active Directory site, the message is relayed to a Hub Transport server in that site and then delivered to the Mailbox server.

Management Shell

You can use the **Set-AdSiteLink** cmdlet in the Exchange Management Shell to configure an Exchange-specific cost to an Active Directory IP site link. The Exchange-specific cost is a separate attribute that is used instead of the Active

Directory-assigned cost to determine the Exchange routing path. This configuration is useful when the Active Directory IP site link costs do not result in an optimal Exchange message routing topology.

Storing Exchange Server 2007 Data in Active Directory

It was mentioned earlier that Active Directory is divided into three naming partitions: configuration, schema, and domain. This section discusses how Exchange Server 2007 uses each of these partitions and which kind of data is stored in them.

Domain Naming Partition

In the domain naming partition, all domain objects for Exchange Server 2007 are stored and replicated to every domain controller in the domain. Recipient objects, including users, contacts, and groups, are stored in this partition. Exchange Server 2007 exploits Active Directory by adding attributes to user, group, and contact objects for messaging purposes.

Designing a Group Implementation Strategy

Exchange Server 2007 uses distribution groups to send the same message to a large number of recipients. Any user accounts that are placed inside the distribution group will receive the message. In Windows Server 2003 native mode, groups can be nested inside of other groups, effectively creating a multitiered distribution list. The two types of groups you will use most often for large distribution of a message are global and universal.

The largest downside to universal groups is that membership is fully replicated to each Global Catalog server, which means that replication traffic occurs whenever a Universal Group's membership changes. Therefore, it is best to populate a universal group with other global groups so that when membership changes in the global group, the universal group is not changed and traffic is not replicated.

Global groups can also be mail-enabled for message distribution. If you choose not to use universal groups, you can mail-enable global groups. Membership for a global group is not promoted to the Global Catalog server, which presents some issues to consider when working in a multidomain environment.

When a message is sent to a global group in a remote domain, the expansion server must connect to a domain controller in the group's home domain and retrieve the membership list. In addition, the expansion server must have IP connectivity to a domain controller in the group's home domain. If bandwidth between the two domains is slow or unreliable, retrieving membership from a remote domain might take time and slow down message delivery, which affects overall performance. It is best if Exchange Server 2007 is in the remote domain. Then you can set the expansion server to be the remote Exchange Server

2007 server instead of retrieving the membership remotely and expanding the group membership locally.

When deciding which group type to select, consider the following implications:

- **Are you using a single-domain or multiple-domain environment?** If you have a single domain, you don't need to use universal groups, because all of the domain objects are local. When you have multiple domains, use universal groups if the membership is fairly static (that is, global groups as opposed to individual users), and remember that users might not have access to all object attributes from other domains in universal groups.
- **Is direct IP connectivity possible between all domains?** If you have IP connectivity, use global groups when membership changes frequently or you have Exchange servers in each domain that can act as expansion servers. Otherwise, use universal groups, because membership is static, and the local expansion server can expand the list.
- **Will membership change frequently?** If membership changes often, use global groups. If membership changes infrequently, use universal groups.

Microsoft Outlook users will not be able to view the user memberships of a group that has been created in a remote domain. They can view membership only in global groups and domain local groups that have been created in their home domain.

An Expansion server, which has been mentioned several times, requires some explanation. When a message is sent to a mail-enabled group, the message needs to be expanded and individually addressed to each member of the group. By default, the local server performs the expansion and uses LDAP to contact the Global Catalog server to deliver the message to each member of the group. If the message is intended for a local group in the domain, the local Global Catalog server is contacted. If the local server is not available for expansion, another server in the site is used.

You can select a specific server in an organization to be an expansion server. The advantage of doing this is that you can offload the sometimes resource-intensive process of expanding large distribution groups to a dedicated server, removing the load from the mailbox server. The disadvantage of specifying an expansion server is that if that expansion server is not available, messages to the distribution group are not delivered, as Exchange does not try another server. For this reason, if you choose to designate an expansion server, take steps to ensure high availability for that server.

Configuration Naming Partition

The configuration partition of Active Directory stores information regarding how your Exchange Server 2007 system is organized. Because this information is replicated to all

domain controllers in the forest, the Exchange Server 2007 configuration is also replicated throughout the forest. The configuration information includes the Exchange Server 2007 topology, connectors, protocols, and service settings.

Schema Naming Partition

The schema partition contains all object types and their attributes that can be created in Active Directory. This information is replicated to all domain controllers in the forest. During the first installation of Exchange Server 2007 in the forest, the Active Directory schema is extended to include new object classes and attributes that are specific to Exchange Server 2007. These new classes start with “msExch” or “ms-Exch” and are derived from the LDAP Data Interchange Format (LDIF) information in the Exchange Server 2007 installation files.

Given that these extensions represent more than 1000 changes to the schema and that these changes are replicated to all the domain controllers in your forest, you should prepare the forest for Exchange Server 2007 at the beginning of a period of time when you anticipate that network activity will be relatively light—for example, on a Friday night. This schedule gives the domain controllers time to replicate all the schema changes into their own databases.

Note You can install Exchange Server 2007 using the */prepare AD* switch, which writes the new Exchange object classes and attributes to the schema but does not install Exchange itself. Plan on this activity taking anywhere from 30 to 90 minutes, depending on the speed and capacity of your hardware. Also, generally, the earlier in an Active Directory deployment that you extend schema, the better, because as domain controllers are added to the forest, they inherit the extended schema, thus reducing replication traffic when */prepare AD* is run. For more information on installing Exchange Server 2007, consult Chapter 6.

Exchange Server 2007 and Forest Boundaries

Because Exchange Server 2007 stores much of its information in the configuration naming partition, an Exchange Server 2007 organization cannot be extended past the boundaries of the forest. This is one area in which your Active Directory structure directly influences your Exchange topology. Having multiple forests in a company incurs the following limitations:

- You have separate Exchange organizations to administer.
- You have separate Global Address Lists, with no automatic directory replication between them.
- All e-mail system features are not available between forests.

Cross-Forest authentication is available, however. See Chapter 21, “Messaging Security,” for a discussion about this topic.

Although using a single forest is the recommended way to set up an Exchange Server 2007 organization, if you want to synchronize directory information among multiple forests, you can use one of two scenarios:

- **Resource forest** In this scenario, one forest is dedicated to running Exchange Server 2007 and hosting mailboxes. User accounts associated with those mailboxes are contained in separate forests. The disadvantage of using this scenario is higher cost associated with configuring the additional forest, domain controllers, and Exchange servers. You must also ensure that when objects are created in one forest, corresponding placeholder objects are created in the other forests.
- **Cross-forest** Exchange Server 2007 runs in multiple forests, and e-mail functionality is configured between forests. The primary disadvantage of this scenario is reduced e-mail functionality between forests.

Configuration Partition and Directory Data

The two Active Directory services that an Exchange server uses most often are the Global Catalog server for address book lookups and the configuration naming partition for routing information. It is possible that two different domain controllers will be referenced, depending on the type of request being made by the Exchange server.

When an Exchange server starts up, it establishes a number of LDAP connections to domain controllers and Global Catalog servers. If it needs routing information to route a message, it can contact any domain controller to obtain this information, because each domain controller in the forest has a full copy of the configuration naming partition. If the Exchange server needs to obtain the Global Address List, it contacts the closest Global Catalog server. Best practice is to place a Global Catalog server near the Exchange server and make sure that they are in the same site and domain.

DNS Configuration

On the Internet (or on any TCP/IP network, for that matter), every device is represented by an IP address—using a four-part dotted-decimal notation, such as 192.168.0.1. A device with a TCP/IP address is called a *host* and is assigned a host name, which is a character-based name that is easier for humans to recognize and remember than its numeric IP address. The format of the host name is *hostname.domain.com*. When a host name identifies a resource on a TCP/IP network, computers must translate that host name into an IP

address because computers communicate using only IP addresses. This translation is called *name resolution*.

Two basic methods exist for resolving host names to IP addresses on a TCP/IP network. The first involves using a Hosts file. The Hosts file is a single, flat file that simply lists hosts on a network and each host's IP address. To use the SMTP with a Hosts file, enter into that file the domain name and IP address of the hosts to which the IMS might need to transfer messages. As you might imagine, this process can be time consuming.

The second method of resolving names is more efficient. It involves the Domain Name System (DNS), a hierarchical, distributed database of host names and IP addresses. In order to run Exchange Server 2007, you must have already installed Windows Server 2003 Active Directory and DNS services on your network. Although host files are still available in Windows Server 2003, given the dynamic nature of the new implementation of DNS, there are few times when you'll want to use them.

You are likely to want outside SMTP hosts to be able to transfer messages to your SMTP service. To enable this capability, create two records in the DNS database so that those outside hosts can resolve your server's IP address. The first record you must create is an address record, or a record for your Exchange server. This can be registered dynamically with DNS in Windows Server 2003. The second record is a mail exchanger record, or MX record, which is a standard DNS record type used to designate one or more hosts that process mail for an organization or site. This record must be entered manually in your DNS tables.

More Info This chapter provides a simple discussion of configuring TCP/IP and DNS, but these topics actually encompass a monstrous amount of material. If you need more information about using TCP/IP and DNS in the Windows Server 2003 environment, see *Microsoft Windows Server 2003 Administrator's Companion*, Second Edition, by Charlie Russel, Sharon Crawford, and Jason Gerend (Microsoft Press, 2006).

Summary

This chapter described the ways in which Exchange Server 2007 is integrated with Windows Server 2003. It gave an overview of how Active Directory is structured and described how Exchange Server 2007 works with Active Directory. It also discussed the Internet information protocols installed with Windows Server 2003 as well as services available in Exchange Server 2007, such as Outlook Web Access. In the Chapter 3, "Exchange Server 2007 Architecture," you learn more about the Exchange Server 2007 architecture.