

**INCLUDES FREE
WEB-BASED TESTING!**

COVERS ALL
**100%
CERTIFIED**

EXAM OBJECTIVES

MCSSE

Exam 70-294: Planning, Implementing, and
Maintaining a Windows Server 2003
Active Directory Infrastructure


STUDY GUIDE & DVD TRAINING SYSTEM

DVD provides a "virtual classroom:" Get the benefits of instructor led training at a fraction of the cost and hassle

Guaranteed Coverage of all Exam Objectives:
If the topic is listed in the exam objectives,
it is covered here

Fully Integrated Learning: Includes a
Study Guide, DVD training and Web-based
practice exams

**Dr. Thomas W. Shinder, MD
Debra Littlejohn Shinder**





Syngress knows what passing the exam means to you and to your career. And we know that you are often financing your own training and certification; therefore, you need a system that is comprehensive, affordable, and effective.

Boasting one-of-a-kind integration of text, DVD-quality instructor-led training, and Web-based exam simulation, the Syngress Study Guide & DVD Training System guarantees 100% coverage of exam objectives.

The Syngress Study Guide & DVD Training System includes:

- **Study Guide with 100% coverage of exam objectives** By reading this study guide and following the corresponding objective list, you can be sure that you have studied 100% of the exam objectives.
- **Instructor-led DVD** This DVD provides almost two hours of virtual classroom instruction.
- **Web-based practice exams** Just visit us at www.syngress.com/certification to access a complete exam simulation.

Thank you for giving us the opportunity to serve your certification needs. And be sure to let us know if there's anything else we can do to help you get the maximum value from your investment. We're listening.

www.syngress.com/certification



MCSSE

Exam 70-294: Planning, Implementing, and
Maintaining a Windows Server 2003
Active Directory Infrastructure

STUDY GUIDE & DVD TRAINING SYSTEM

Michael Cross

Jeffery A. Martin

Todd A. Walls

Martin Grasdal Technical Reviewer

Debra Littlejohn Shinder Technical Editor

Dr. Thomas W. Shinder Technical Editor

Syngress Publishing, Inc., the author(s), and any person or firm involved in the writing, editing, or production (collectively “Makers”) of this book (“the Work”) do not guarantee or warrant the results to be obtained from the Work.

There is no guarantee of any kind, expressed or implied, regarding the Work or its contents. The Work is sold AS IS and WITHOUT WARRANTY. You may have other legal rights, which vary from state to state.

In no event will Makers be liable to you for damages, including any loss of profits, lost savings, or other incidental or consequential damages arising out from the Work or its contents. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

You should always use reasonable care, including backup and other appropriate precautions, when working with computers, networks, data, and files.

Syngress Media®, Syngress®, “Career Advancement Through Skill Enhancement®,” “Ask the Author UPDATE®,” and “Hack Proofing®,” are registered trademarks of Syngress Publishing, Inc. “Mission Critical™,” and “The Only Way to Stop a Hacker is to Think Like One™” are trademarks of Syngress Publishing, Inc. Brands and product names mentioned in this book are trademarks or service marks of their respective companies.

KEY	SERIAL NUMBER
001	TH33SLUGGY
002	Q2T4J9T7VA
003	82LPD8R7FF
004	Z6TDA3HVY
005	P33JEET8MS
006	3SHX65N\$RK
007	CH3W7E42AK
008	9EU6V4DER7
009	SUPACM4NFH
010	5BVF3MEV2Z

PUBLISHED BY
Syngress Publishing, Inc.
800 Hingham Street
Rockland, MA 02370

Planning, Implementing, and Maintaining a Windows Server 2003 Active Directory Infrastructure Study Guide & DVD Training System

Copyright © 2003 by Syngress Publishing, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

Printed in the United States of America

1 2 3 4 5 6 7 8 9 0

ISBN: 1-931836-94-9

Technical Editors: Debra Littlejohn Shinder
Thomas W. Shinder
Technical Reviewer: Martin Grasdal
Acquisitions Editor: Jonathan Babcock
DVD Production: Michael Donovan

Cover Designer: Michael Kavish
Page Layout and Art by: Patricia Lupien
Copy Editor: Beth Roberts
Indexer: Rich Carlson
DVD Presenter: Laura E. Hunter



Acknowledgments

We would like to acknowledge the following people for their kindness and support in making this book possible.

Will Schmied, the President of Area 51 Partners, Inc. and moderator of www.mcseworld.com for sharing his considerable knowledge of Microsoft networking and certification.

Karen Cross, Meaghan Cunningham, Kim Wylie, Harry Kirchner, Kevin Votel, Kent Anderson, Frida Yara, Jon Mayes, John Mesjak, Peg O'Donnell, Sandra Patterson, Betty Redmond, Roy Remer, Ron Shapiro, Patricia Kelly, Andrea Tetrick, Jennifer Pascal, Doug Reil, David Dahl, Janis Carpenter, and Susan Fryer of Publishers Group West for sharing their incredible marketing experience and expertise.

The incredibly hard working team at Elsevier Science, including Jonathan Bunkell, AnnHelen Lindeholm, Duncan Enright, David Burton, Rosanna Ramacciotti, Robert Fairbrother, Miguel Sanchez, Klaus Beran, and Rosie Moss for making certain that our vision remains worldwide in scope.

David Buckland, Wendi Wong, Daniel Loh, Marie Chieng, Lucy Chong, Leslie Lim, Audrey Gan, and Joseph Chan of Transquest Publishers for the enthusiasm with which they receive our books.

Kwon Sung June at Acorn Publishing for his support.

Jackie Gross, Gayle Voycey, Alexia Penny, Anik Robitaille, Craig Siddall, Darlene Morrow, Iolanda Miller, Jane Mackay, and Marie Skelly at Jackie Gross & Associates for all their help and enthusiasm representing our product in Canada.

Lois Fraser, Connie McMenemy, Shannon Russell, and the rest of the great folks at Jaguar Book Group for their help with distribution of Syngress books in Canada.

David Scott, Annette Scott, Delta Sams, Geoff Ebbs, Hedley Partis, and Tricia Herbert of Woodslane for distributing our books throughout Australia, New Zealand, Papua New Guinea, Fiji Tonga, Solomon Islands, and the Cook Islands.

Winston Lim of Global Publishing for his help and support with distribution of Syngress books in the Philippines.

A special thanks to Deb and Tom Shinder for going the extra mile on our core four MCSE 2003 guides. Thank you both for all your work.

Another special thanks to Daniel Bendell from Assurance Technology Management for his 24x7 care and feeding of the Syngress network. Dan manages our network in a highly professional manner and under severe time constraints, but still keeps a good sense of humor.



Contributors

Michael Cross (MCSE, MCP+I, CNA, Network+) is an Internet Specialist / Computer Forensic Analyst with the Niagara Regional Police Service. He performs computer forensic examinations on computers involved in criminal investigations, and has consulted and assisted in cases dealing with computer-related/Internet crimes. In addition to designing and maintaining their Web site at www.nrps.com and Intranet, he has also provided support in the areas of programming, hardware, network administration, and other services. As part of an information technology team that provides support to a user base of over 800 civilian and uniform users, his theory is that when the users carry guns, you tend to be more motivated in solving their problems.

Michael also owns KnightWare (www.knightware.ca), which provides computer-related services like Web page design, and Bookworms (www.bookworms.ca), where you can purchase collectibles and other interesting items online. He has been a freelance writer for several years, and is published over three dozen times in numerous books and anthologies. He currently resides in St. Catharines, Ontario Canada with his lovely wife Jennifer and his darling daughter Sara.

Eriq Oliver Neale is an Information Technology manager for a large manufacturing company headquartered in the southwest. His IT career spans 16 years and just about as many systems. He has contributed to a number of technical publications, including several MCSE exam preparation titles. His article on MIDI, still considered one of the seminal works on the topic, has been reprinted in hundreds of publications in multiple languages. Most recently, he has been focusing on electronic data privacy issues in mixed platform environments. When not working in and writing about information technology, Eriq spends time writing and recording music in his home studio for clients of his music publishing company. On clear nights, he can be found gazing at the moon or planets through his telescope, which he also uses for deep-space astrophotography.

Todd A. Walls (CISSP, MCSE) is a Senior Security Engineer for COACT, Inc., providing information security support to a government customer in Colorado Springs. Todd has over 19 years of IT experience spanning the range of micro, mini, and mainframe systems, running variants of UNIX, Windows, and proprietary operating systems. His security systems experience includes intrusion detection and prevention,

firewalls, biometrics, smart cards, password cracking, vulnerability testing, and secure-computing designs and evaluations. He is currently enrolled in graduate computer science studies at Colorado Technical University with a concentration in computer systems security.

Vinod Kumar is an author, developer and technical reviewer specializing in Web and mobile technologies using Microsoft solutions. He has been awarded the Microsoft's Most Valuable Professional (MVP) in .NET. He Currently works for Verizon. Vinod is a lead author for the forthcoming title *Mobile Application Development with .NET* and has co authored several other books. He had written many technical articles for sites like ASPToday, C# Today, and CSharp-Corner. Vinod runs a community site named www.dotnetforce.com which provides content related to .NET. In his free time he likes to spend time with his family and friends.

Brian Frederick is a Lead Network Analyst for Aegon USA, one of the top 5 insurance companies in the United States. Brian started working with computers on the Apple II+. Brian attended the University of Northern Iowa and is married with two adorable children. He is also a technical instructor at a local community college teaching MCSE, MCSA, A+, and Network+ certification courses. Brian owes his success to his parents and brother for their support and backing during his Apple days and in college, and to his wife and children for their support and understanding when dad spends many hours in front of the computer.

M. Troy Hudson (MCSE NT/2000, MCP, MCP+I, Master CNE, CNE-IW, CNE-4, CNE-5, CNE-GW4, CNE-GW5, A+) is the computer services manager for Sodexo at Granite School District Food Services in Salt Lake City, UT. He currently manages around 90 sites using a lot of remote management tools, inter-networking Microsoft Windows desktops with Novell networks and ZENworks for Desktops.

Troy has been a consultant, trainer, and writer since 1997 and has published items both on the Internet and with this publisher. He has authored student curricula and helped design training material and labs for students trying to pass the Microsoft MCSE exams. He holds a bachelor's degree from the University of Phoenix in e-Business. Troy currently resides in Salt Lake City, UT with his wife Kim and eight children: "My family is the reason for taking on extra projects and

I am grateful for their support! I love you Kim, Jett, Ryan, Rachael, James, McKay, Brayden, Becca and Hannah.”



Technical Editors

Debra Littlejohn Shinder (MCSE) is a technology consultant, trainer, and writer who has authored a number of books on networking, including *Scene of the Cybercrime: Computer Forensics Handbook*, published by Syngress Publishing (ISBN: 1-931836-65-5), and *Computer Networking Essentials*, published by Cisco Press. She is co-author, with her husband Dr. Thomas Shinder, of *Troubleshooting Windows 2000 TCP/IP* (ISBN: 1-928994-11-3), the best-selling *Configuring ISA Server 2000* (ISBN: 1-928994-29-6), and *ISA Server and Beyond* (ISBN: 1-931836-66-3). Deb is also a technical editor and contributor to books on subjects such as the Windows 2000 MCSE exams, the CompTIA Security+ exam, and TruSecure's ICSA certification. She edits the Brainbuzz A+ Hardware News and Sunbelt Software's WinXP News and is regularly published in TechRepublic's TechProGuild and Windowsecurity.com. Deb specializes in security issues and Microsoft products. She lives and works in the Dallas-Fort Worth area and can be contacted at deb@shinder.net or via the website at www.shinder.net.

Thomas W. Shinder M.D. (MVP, MCSE) is a computing industry veteran who has worked as a trainer, writer, and a consultant for Fortune 500 companies including FINA Oil, Lucent Technologies, and Sealand Container Corporation. Tom was a Series Editor of the Syngress/Osborne Series of Windows 2000 Certification Study Guides and is author of the best selling books *Configuring ISA Server 2000: Building Firewalls with Windows 2000* (Syngress Publishing, ISBN: 1-928994-29-6) and *Dr. Tom Shinder's ISA Server and Beyond* (ISBN: 1-931836-66-3). Tom is the editor of the Brainbuzz.com *Win2k News* newsletter and is a regular contributor to TechProGuild. He is also content editor, contributor and moderator for the World's leading site on ISA Server 2000, www.isaserver.org. Microsoft recognized Tom's leadership in the ISA Server community and awarded him their Most Valued Professional (MVP) award in December of 2001.



Technical Editor and Contributor

Jeffery A. Martin (MCSE, MCDBA, MCT, MCP+I, MCNE, CNI, CCNP, CCI, CCA, CTT, A+, Network+, I-Net+, Project+, Linux+, CIW, ADPM) has been working with computers and computer networks for over 15 years. Jeffery spends most of his time managing several companies that he owns and consulting for large multinational media companies. He also enjoys working as a technical instructor and training others in the use of technology.



Technical Reviewer

Martin Grasdal (MCSE+I, MCSE/W2K MCT, CISSP, CTT+, A+) is an independent consultant with over 10 years experience in the computer industry. Martin has a wide range of networking and IT managerial experience. He has been an MCT since 1995 and an MCSE since 1996. His training and networking experience covers a number of products, including NetWare, Lotus Notes, Windows NT, Windows 2000, Windows 2003, Exchange Server, IIS, and ISA Server. As a manager, he served as Director of Web Sites and CTO for BrainBuzz.com, where he was also responsible for all study guide and technical content on the CramSession.com Web site. Martin currently works actively as a consultant, author, and editor. His recent consulting experience includes contract work for Microsoft as a technical contributor to the MCP program on projects related to server technologies. Martin lives in Edmonton, Alberta, Canada with his wife Cathy and their two sons. Martin's past authoring and editing work with Syngress has included the following titles: *Configuring and Troubleshooting Windows XP Professional* (ISBN: 1-928994-80-6), *Configuring ISA Server 2000: Building Firewalls for Windows 2000* (ISBN: 1-928994-29-6), and *Dr. Tom Shinder's ISA Server & Beyond: Real World Security Solutions for Microsoft Enterprise Networks* (ISBN: 1-931836-66-3).



DVD Presenter

Laura E. Hunter (CISSP, MCSE, MCT, MCDBA, MCP, MCP+I, CCNA, A+, Network+, iNet+, CNE-4, CNE-5) is a Senior IT Specialist with the University of Pennsylvania, where she provides network planning, implementation and troubleshooting services for various business units and schools within the University. Her specialties include Microsoft Windows NT and 2000 design and implementation, troubleshooting and security topics. As an “MCSE Early Achiever” on Windows 2000, Laura was one of the first in the country to renew her Microsoft credentials under the Windows 2000 certification structure. Laura’s previous experience includes a position as the Director of Computer Services for the Salvation Army and as the LAN administrator for a medical supply firm. She also operates as an independent consultant for small businesses in the Philadelphia metropolitan area and is a regular contributor to the TechTarget family of Web sites.

Laura has previously contributed to the Syngress Publishing’s *Configuring Symantec Antivirus, Corporate Edition* (ISBN: 1-931836-81-7). She has also contributed to several other exam guides in the Syngress Windows Server 2003 MCSE/MCSA DVD Guide and Training System series as a DVD presenter, contributing author, and technical reviewer.

Laura holds a bachelor’s degree from the University of Pennsylvania and is a member of the Network of Women in Computer Technology, the Information Systems Security Association, and InfraGard, a cooperative undertaking between the U.S. Government and other participants dedicated to increasing the security of United States critical infrastructures.

MCSE 70-294 Exam Objectives Map and Table of Contents



All of Microsoft's published objectives for the MCSE 70-294 Exam are covered in this book. To help you easily find the sections that directly support particular objectives, we've listed all of the exam objectives below, and mapped them to the Chapter number in which they are covered. We've also assigned numbers to each objective, which we use in the subsequent Table of Contents and again throughout the book to identify objective coverage. In some chapters, we've made the judgment that it is probably easier for the student to cover objectives in a slightly different sequence than the order of the published Microsoft objectives. By reading this study guide and following the corresponding objective list, you can be sure that you have studied 100% of Microsoft's MCSE 70-294 Exam objectives.

Exam Objective Map

Objective Number	Objective	Chapter Number
1	Planning and Implementing an Active Directory Infrastructure	1
1.1	Plan a strategy for placing global catalog servers.	8
1.1.1	Evaluate network traffic considerations when placing global catalog servers.	8
1.1.2	Evaluate the need to enable universal group caching.	8
1.2	Plan flexible operations master role placement.	7
1.2.1	Plan for business continuity of operations master roles.	7
1.2.2	Identify operations master role dependencies.	7
1.3	Implement an Active Directory directory service forest and domain structure.	4
1.3.1	Create the forest root domain.	4
1.3.2	Create a child domain.	4

Objective Number	Objective	Chapter Number
1.3.3	Create and configure Application Data Partitions.	4
1.3.4	Install and configure an Active Directory domain controller.	7
1.3.5	Set an Active Directory forest and domain functional level based on requirements.	4
1.3.6	Establish trust relationships. Types of trust relationships might include external trusts, shortcut trusts, and cross-forest trusts.	5
1.4	Implement an Active Directory site topology.	6
1.4.1	Configure site links.	6
1.4.2	Configure preferred bridgehead servers.	6
1.5	Plan an administrative delegation strategy.	5
1.5.1	Plan an organizational unit (OU) structure based on delegation requirements.	5
1.5.2	Plan a security group hierarchy based on delegation requirements.	5
2	Managing and Maintaining an Active Directory Infrastructure	All chapters
2.1	Manage an Active Directory forest and domain structure.	4
2.1.1	Manage trust relationships.	5
2.1.2	Manage schema modifications.	8
2.1.3	Add or remove a UPN suffix.	8
2.2	Manage an Active Directory site.	6
2.2.1	Configure replication schedules.	6
2.2.2	Configure site link costs.	6
2.2.3	Configure site boundaries.	6
2.3	Monitor Active Directory replication failures. Tools might include Replication Monitor, Event Viewer, and support tools.	6
2.3.1	Monitor Active Directory replication.	6
2.3.2	Monitor File Replication service (FRS) replication.	6

Objective Number	Objective	Chapter Number
2.4	Restore Active Directory directory services.	11
2.4.1	Perform an authoritative restore operation.	11
2.4.2	Perform a nonauthoritative restore operation.	11
2.5	Troubleshoot Active Directory.	All chapters
2.5.1	Diagnose and resolve issues related to Active Directory replication.	6
2.5.2	Diagnose and resolve issues related to operations master role failure.	7
2.5.3	Diagnose and resolve issues related to the Active Directory database.	11
3	Planning and Implementing User, Computer, and Group Strategies	2
3.1	Plan a security group strategy.	3
3.2	Plan a user authentication strategy.	3
3.2.1	Plan a smart card authentication strategy.	3
3.2.2	Create a password policy for domain users.	3
3.3	Plan an OU structure.	5
3.3.1	Analyze the administrative requirements for an OU.	5
3.3.2	Analyze the Group Policy requirements for an OU structure.	5
3.4	Implement an OU structure.	5
3.4.1	Create an OU.	5
3.4.2	Delegate permissions for an OU to a user or to a security group.	5
3.4.3	Move objects within an OU hierarchy.	5
4	Planning and Implementing Group Policy	9
4.1	Plan Group Policy strategy.	9
4.1.1	Plan a Group Policy strategy by using Resultant Set of Policy (RSOP) Planning mode.	9
4.1.2	Plan a strategy for configuring the user environment by using Group Policy.	9
4.1.3	Plan a strategy for configuring the computer environment by using Group Policy.	9

Objective Number	Objective	Chapter Number
4.2	Configure the user environment by using Group Policy.	9
4.2.1	Distribute software by using Group Policy.	10
4.2.2	Automatically enroll user certificates by using Group Policy.	9
4.2.3	Redirect folders by using Group Policy.	9
4.2.4	Configure user security settings by using Group Policy.	9
4.3	Deploy a computer environment by using Group Policy.	
4.3.1	Distribute software by using Group Policy.	10
4.3.2	Automatically enroll computer certificates by using Group Policy.	9
4.3.3	Configure computer security settings by using Group Policy.	9
5	Managing and Maintaining Group Policy	9
5.1	Troubleshoot issues related to Group Policy application. deployment. Tools might include RSoP and the gpresult command.	9
5.2	Maintain installed software by using Group Policy.	10
5.2.1	Distribute updates to software distributed by Group Policy.	10
5.2.2	Configure automatic updates for network clients by using Group Policy.	10
5.3	Troubleshoot the application of Group Policy security settings. Tools might include RSoP and the gpresult command.	9

Contents

	Foreword	xxxii
	Chapter 1 Active Directory Infrastructure Overview	1
	Introduction	2
1	Introducing Directory Services	2
	Terminology and Concepts	5
	Directory Data Store	5
	Policy-Based Administration	9
	Directory Access Protocol	10
	Naming Scheme	11
	Installing Active Directory to Create a Domain Controller	15
1	Understanding How Active Directory Works	19
	Directory Structure Overview	19
	Sites	20
	Domains	21
	Domain Trees	22
	Forests	24
	Organizational Units	25
	Active Directory Components	26
	Logical vs. Physical Components	27
	Domain Controllers	28
	Schema	31
	Global Catalog	31
	Replication Service	32
1	Using Active Directory Administrative Tools	34
	Graphical Administrative Tools/MMCs	35
	Active Directory Users and Computers	38
	Active Directory Domains and Trusts	40
	Active Directory Sites and Services	44
	Command-Line Tools	45

	Cacls	46
	Cmdkey	47
	Csvde	47
	Dcgpofix	49
	Dsadd	49
	Dsget	49
	Dsmod	50
	Dsmove	50
	Ldifde	51
	Ntdsutil	53
	Whoami	54
1	Implementing Active Directory Security and Access Control	55
	Access Control in Active Directory	55
	Role-Based Access Control	60
	Authorization Manager	60
	Active Directory Authentication	61
	Standards and Protocols	62
	Kerberos	62
	X.509 Certificates	63
	LDAP/SSL	63
	PKI	64
1	What's New in Windows Server 2003 Active Directory?	65
	New Features Available on All Windows	
	Server 2003 Computers	68
	New Features Available Only with	
	Windows Server 2003 Domain/Forest Functionality	69
	Domain Controller Renaming Tool	70
	Domain Rename Utility	70
	Forest Trusts	70
	Dynamically Links Auxiliary Classes	70
	Disabling Classes	70
	Replication	70
	Summary of Exam Objectives	72
	Exam Objectives Fast Track	73
	Exam Objectives Frequently Asked Questions	75
	Self Test	76
	Self Test Quick Answer Key	81

	Chapter 2 Working with User, Group, and Computer Accounts	83
	Introduction	84
3	Understanding Active Directory Security Principal Accounts	84
	Security Principals and Security Identifiers	85
	Tools to View and Manage Security Identifiers	90
	Naming Conventions and Limitations	92
3	Working with Active Directory User Accounts	99
	Built-In Domain User Accounts	101
	Administrator	102
	Guest	103
	HelpAssistant	104
	SUPPORT_388945a0	104
	InterOrgPerson	104
	Creating User Accounts	105
	Creating Accounts Using	
	Active Directory Users and Computers	105
	Creating Accounts Using the DSADD Command	110
	Managing User Accounts	113
	Personal Information Tabs	115
	Account Settings	118
	Terminal Services Tabs	122
	Security-Related Tabs	126
3	Working with Active Directory Group Accounts	131
	Group Types	131
	Security Groups	132
	Distribution Groups	132
	Group Scopes in Active Directory	133
	Universal	134
	Global	134
	Domain Local	135
	Built-In Group Accounts	135
	Default Groups in Builtin Container	136
	Default Groups in Users Container	138
	Creating Group Accounts	140
	Creating Groups Using Active	
	Directory Users and Computers	141
	Creating Groups Using the DSADD Command	142

	Managing Group Accounts	143
3	Working with Active Directory Computer Accounts	150
	Creating Computer Accounts	150
	Creating Computer Accounts by Adding a Computer to a Domain	151
	Creating Computer Accounts Using Active Directory Users and Computers	152
	Creating Computer Accounts Using the DSADD Command	155
	Managing Computer Accounts	156
3	Managing Multiple Accounts	162
	Implementing User Principal Name Suffixes	162
	Moving Account Objects in Active Directory	164
	Moving Objects with Active Directory Users and Computers	164
	Moving Objects with the DSMOVE Command	165
	Moving Objects with the MOVETREE Command	166
	Troubleshooting Problems with Accounts	168
	Summary of Exam Objectives	170
	Exam Objectives Fast Track	171
	Exam Objectives Frequently Asked Questions	173
	Self Test	174
	Self Test Quick Answer Key	179
	Chapter 3 Creating User and Group Strategies	181
	Introduction	182
	Creating a Password Policy for Domain Users	182
	Creating an Extensive Defense Model	183
	Strong Passwords	184
	System Key Utility	185
	Defining a Password Policy	187
	Applying a Password Policy	187
	Modifying a Password Policy	190
	Applying an Account Lockout Policy	190
	Creating User Authentication Strategies	192
	Need for Authentication	193
	Single Sign-On	194
	Interactive Logon	194
	Network Authentication	195
	Authentication Types	195
	Kerberos	195

Understanding the Kerberos Authentication Process	196
Secure Sockets Layer/Transport Layer Security	197
NT LAN Manager	198
Digest Authentication	199
Passport Authentication	200
Educating Users	202
Planning a Smart Card Authentication Strategy	203
When to Use Smart Cards	205
Implementing Smart Cards	206
PKI and Certificate Authorities	206
Setting Security Permissions	208
Enrollment Stations	209
Enabling Certificate Templates	209
Requesting an Enrollment Agent Certificate	211
Enrolling Users	211
Installing a Smart Card Reader	212
Issuing Smart Card Certificates	213
Assigning Smart Cards	214
Logon Procedures	215
Revoking Smart Cards	215
Planning for Smart Card Support	216
Planning a Security Group Strategy	217
Understanding Group Types and Scopes	218
Security and Distribution Groups	218
Local, Domain Local, Global, and Universal Groups	219
Security Group Best Practices	224
Designing a Group Strategy for a Single Domain Forest	225
Designing a Group Strategy for a Multiple Domain Forest	226
Summary of Exam Objectives	230
Exam Objectives Fast Track	232
Exam Objectives Frequently Asked Questions	233
Self Test	235
Self Test Quick Answer Key	241
Chapter 4 Working with Forests and Domains	243
Introduction	244
Understanding Forest and Domain Functionality	244

	The Role of the Forest	246
	New Forestwide Features	247
	The Role of the Domain	254
	New Domainwide Features	256
	Domain Trees	259
	Forest and Domain Functional Levels	259
	Domain Functionality	260
	Forest Functionality	265
1.3.5	Raising the Functional Level of a Domain and Forest	270
	Domain Functional Level	270
	Forest Functional Level	272
	Optimizing Your Strategy for Raising Functional Levels	273
1.3/2.1	Creating the Forest and Domain Structure	275
	Deciding When to Create a New DC	275
	Installing Domain Controllers	276
1.3.1	Creating a Forest Root Domain	278
	Creating a New Domain Tree in an Existing Forest	285
1.3.2	Creating a New Child Domain in an Existing Domain	288
	Creating a New DC in an Existing Domain	293
	Assigning and Transferring Master Roles	300
1.3.3	Using Application Directory Partitions	313
	Establishing Trust Relationships	315
	Direction and Transitivity	315
	Types of Trusts	317
	Restructuring the Forest and Renaming Domains	318
	Domain Rename Limitations	318
	Domain Rename Limitations in a Windows 2000 Forest	319
	Domain Rename Limitations in a Windows Server 2003 Forest	319
	Domain Rename Dependencies	320
	Domain Rename Conditions and Effects	322
	Domain Rename Preliminary Steps	323
	Performing the Rename Procedure	334
	Steps to Take After the Domain Rename Procedure	354
	Implementing DNS in the Active Directory Network Environment	365
	DNS and Active Directory Namespaces	367
	DNS Zones and Active Directory Integration	367
	Configuring DNS Servers for Use with Active Directory	370

	Integrating an Existing Primary	
	DNS Server with Active Directory	370
	Creating the Default DNSApplication Directory Partitions	371
	Using dnscmd to Administer Application Directory Partitions	372
	Securing Your DNS Deployment	373
	Summary of Exam Objectives	374
	Exam Objectives Frequently Asked Questions	376
	Exam Objectives Fast Track	377
	Self Test	379
	Self Test Quick Answer Key	387
	Chapter 5 Working with Trusts and Organizational Units	389
	Introduction	390
1.3.6/	Working with Active Directory Trusts	390
2.1.1	Types of Trust Relationships	394
	Default Trusts	395
	Shortcut Trust	395
	Realm Trust	396
	External Trust	396
	Forest Trust	397
	Creating, Verifying, and Removing Trusts	398
	Securing Trusts Using SID Filtering	400
3.3.1/	Working with Organizational Units	401
3.4.3	Understanding the Role of Container Objects	402
3.4/ 3.4.1	Creating and Managing Organizational Units	402
	Applying Group Policy to OUs	406
3.4.2	Delegating Control of OUs	407
1.5/1.5.1/	Planning an OU Structure and Strategy for Your Organization	408
3.3/3.3.2	Delegation Requirements	409
	Security Group Hierarchy	410
	Summary of Exam Objectives	412
	Exam Objectives Fast Track	413
	Exam Objectives Frequently Asked Questions	414

	Self Test	416
	Self Test Quick Answer Key	423
	Chapter 6 Working with Active Directory Sites	425
	Introduction	426
	Understanding the Role of Sites	426
	Replication	427
	Authentication	427
	Interactive Logon Authentication	428
	Network Authentication	429
	Distribution of Services Information	429
	Relationship of Sites to Other Active Directory Components	431
	Relationship of Sites and Domains	431
	Physical vs. Logical Structure of the Network	433
	The Relationship of Sites and Subnets	433
1.4/2.2/	2.2.3 Creating Sites and Site Links	434
	Site Planning	434
	Criteria for Establishing Separate Sites	435
	Creating a Site	436
	Renaming a Site	438
	Creating Subnets	441
	Associating Subnets with Sites	444
1.4.1/2.2.2	Creating Site Links	446
1.4.1/2.2.2	Configuring Site Link Cost	449
2.2/2.2.1/	2.5.1 Understanding Site Replication	452
	Purpose of Replication	452
	Types of Replication	453
	Intrasite Replication	453
	Intersite Replication	454
1.4	Planning, Creating, and Managing the Replication Topology	455
	Planning Replication Topology	455
	Creating a Replication Topology	456
	Managing Replication Topology	456
	Configuring Replication between Sites	457
	Configuring Replication Frequency	457
	Configuring Site Link Availability	458

	Configuring Site Link Bridges	458
1.4.2	Configuring Bridgehead Servers	459
2.3	Troubleshooting Replication Failure	459
	Troubleshooting Replication	460
2.3.1	Using Replication Monitor	461
	Using Event Viewer	461
	Using Support Tools	462
2.3.2	Monitoring File Replication Service Replication	463
	Summary of Exam Objectives	465
	Exam Objectives Fast Track	465
	Exam Objectives Frequently Asked Questions	467
	Self Test	468
	Self Test Quick Answer Key	474
	Chapter 7 Working with Domain Controllers	475
	Introduction	476
1.3.4	Planning and Deploying Domain Controllers	476
	Understanding Server Roles	476
	Function of Domain Controllers	480
	Determining the Number of Domain Controllers	481
	Using the Active Directory Installation Wizard	484
	Creating Additional Domain Controllers	494
	Upgrading Domain Controllers	500
	Placing Domain Controllers within Sites	502
	Backing Up Domain Controllers	503
	Restoring Domain Controllers	504
1.2/2.5.2	Managing Operations Masters	505
	Understanding the Operation Masters Roles	505
	Forestwide Roles	506
	Domainwide Roles	507
1.2.1	Transferring and Seizing Operations Master Roles	509
	Transferring FSMOs	510
	Transferring the Schema FSMO	510
	Transferring Domain Naming FSMO	514
	Transferring RID, PDC, or Infrastructure FSMOs	516
1.2.1	Responding to OM Failures	516
	Seizing the PDC Emulator or Infrastructure FSMO	516
	Seizing the RID Master, Domain	

	Naming Master, and Schema Master FSMOs	517
	Summary of Exam Objectives	523
	Exam Objectives Fast Track	524
	Exam Objectives Frequently Asked Questions	526
	Self Test	528
	Self Test Quick Answer Key	537
	Chapter 8 Working with Global Catalog Servers and Schema	539
	Introduction	540
	Working with the Global Catalog and GC Servers	540
	Functions of the GC	541
2.1.3	UPN Authentication	541
	Directory Information Search	542
	Universal Group Membership Information	543
	Customizing the GC Using the Schema MMC Snap-In	543
	Creating and Managing GC Servers	545
	Understanding GC Replication	547
	Universal Group Membership	547
	Attributes in GC	547
1.1	Placing GC Servers within Sites	548
1.1.1	Bandwidth and Network Traffic Considerations	549
1.1.2	Universal Group Caching	550
	Troubleshooting GC Issues	552
2.1.2	Working with the Active Directory Schema	551
	Understanding Schema Components	553
	Classes	554
	Attributes	555
	Naming of Schema Objects	559
	Working with the Schema MMC Snap-In	560
	Modifying and Extending the Schema	561
	Deactivating Schema Classes and Attributes	562
	Troubleshooting Schema Issues	563
	Summary of Exam Objectives	564
	Exam Objectives Fast Track	565
	Exam Objectives Frequently Asked Questions	566
	Self Test	567
	Self Test Quick Answer Key	573

	Chapter 9 Working with Group Policy in an Active Directory Environment	575
	Introduction	576
4/4.2.1	4.3.1 Understanding Group Policy	576
	Terminology and Concepts	577
	Local and Non-Local Policies	577
	User and Computer Policies	577
	Group Policy Objects	580
	Scope and Application Order of Policies	580
	Group Policy Integration in Active Directory	583
	Group Policy Propagation and Replication	583
4/4.1	4.2.1/4.3.1 Planning a Group Policy Strategy	584
	Using RSoP Planning Mode	584
	Opening RSoP in Planning Mode	584
	Reviewing RSoP Results	587
	Strategy for Configuring the User Environment	588
	Strategy for Configuring the Computer Environment	590
4/4.2.1	4.3.1 Implementing Group Policy	596
	The Group Policy Object Editor MMC	595
	Creating, Configuring, and Managing GPOs	595
	Creating and Configuring GPOs	596
	Naming GPOs	597
	Managing GPOs	598
	Configuring Application of Group Policy	600
	General	600
	Links	601
	Security	601
	WMI Filter	602
	Delegating Administrative Control	604
	Verifying Group Policy	604
4/4.2.1	4.2.2/4.2.3 Performing Group Policy Administrative Tasks	608
4.3.1/4.3.2	Automatically Enrolling User and Computer Certificates	608

	Redirecting Folders	609
4.1.2/4.1.3	Configuring User and Computer Security Settings	612
4.2/4.2.4		
4.3/4.3.3		
	Computer Configuration	612
	User Configuration	613
	Using Software Restriction Policies	616
	Setting Up Software Restriction Policies	616
	Software Policy Rules	617
	Precedence of Policies	617
	Best Practices	618
4/4.2.1	Applying Group Policy Best Practices	619
4.3.1/5		
4/4.2.1	Troubleshooting Group Policy	621
4.3.1/5.1/		
5.3		
4.1.1	Using RSoP	622
	Using gpresult.exe	623
	Summary of Exam Objectives	628
	Fast Track	629
	Exam Objectives Frequently Asked Questions	631
	Self Test	633
	Self Test Quick Answer Key	639
4.2.1/4.3.1	Chapter 10 Deploying Software via Group Policy	641
	Introduction	642
	Understanding Group Policy Software Installation Terminology and Concepts	642
	Group Policy Software Installation Concepts	644
	Assigning Applications	644
	Publishing Applications	646
	Document Invocation	646
	Application Categories	647
	Group Policy Software Deployment vs. SMS Software Deployment	648
	Group Policy Software Installation Components	648
	Windows Installer Packages (.msi)	649
	Transforms (.mst)	650

	Patches and Updates (.msp)	651
	Application Assignment Scripts (.aas)	652
	Deploying Software to Users	652
	Deploying Software to Computers	653
5.2	Using Group Policy Software Installation to Deploy Applications ...	654
	Preparing for Group Policy Software Installation	654
	Creating Windows Installer Packages	654
	Using .zap Setup Files	656
	Creating Distribution Points	659
	Working with the GPO Editor	660
	Opening or Creating a GPO for Software Deployment	659
	Assigning and Publishing Applications	662
	Configuring Software Installation Properties	664
	The General Tab	665
	The Advanced Tab	665
	The File Extensions Tab	666
	The Categories Tab	666
5.2.1	Upgrading Applications	667
5.2.2	Automatically Configuring Required Updates	668
	Removing Managed Applications	669
	Managing Application Properties	670
	Categorizing Applications	673
	Adding and Removing Modifications for Application Packages	673
	Troubleshooting Software Deployment	675
	Verbose Logging	677
	Software Installation Diagnostics Tool	678
	Summary of Exam Objectives	679
	Exam Objectives Fast Track	679
	Exam Objectives Frequently Asked Questions	681
	Self Test	682
	Self Test Quick Answer Key	688
	Chapter 11 Ensuring Active Directory Availability	689
	Introduction	690
	Understanding Active Directory Availability Issues	690
	The Active Directory Database	690
	Data Modification to the Active Directory Database	692
	The Tombstone and Garbage Collection Processes	694

	System State Data	698
	Fault Tolerance and Performance	699
	RAID-1	700
	RAID-5	701
	Performing Active Directory Maintenance Tasks	701
	Defragmenting the Database	702
	Understanding Active Directory Database Fragmentation ...	702
	The Offline Defragmentation Process	703
	Moving the Database or Log Files	707
2.5.3	Monitoring the Database	711
	Using Event Viewer to Monitor Active Directory	711
	Using the Performance Console to Monitor Active Directory	713
	Backing Up and Restoring Active Directory	720
	Backing Up Active Directory	720
	Using the Windows Server 2003 Backup Utility	721
	Backing Up at the Command Line	733
2.4/2.4.1	Restoring Active Directory	733
2.4.2		
	Directory Services Restore Mode	733
	Normal Restore	734
	Authoritative Restore	741
	Primary Restore	743
2.5.3	Troubleshooting Active Directory Availability	745
	Setting Logging Levels for Additional Detail	745
	Using Ntdsutil Command Options	747
	Using the Integrity Command	747
	Using the recover Command	750
	Using the Semantic Database Analysis Command	752
	Using the esentutl Command	756
	Changing the Directory Services Restore Mode Password	758
	Summary of Exam Objectives	759
	Exam Objectives Fast Track	760
	Exam Objectives Frequently Asked Questions	762
	Self Test	764
	Self Test Quick Answer Key	769

Appendix Self Test Questions, Answers, and Explanations	771
Chapter 1: Active Directory Infrastructure Overview	772
Chapter 2: Working with User, Group, and Computer Accounts	781
Chapter 3: Creating User and Group Strategies	789
Chapter 4: Working with Forests and Domains	797
Chapter 5: Working with Trusts and Organizational Units	809
Chapter 6: Working with Active Directory Sites	819
Chapter 7: Working with Domain Controllers	826
Chapter 8: Working with Global Catalog Servers and Schema	840
Chapter 9: Working with Group Policy in an Active Directory Environment	847
Chapter 10: Deploying Software via Group Policy	855
Chapter 11: Ensuring Active Directory Availability	864
Index	873

Foreword

This book's primary goal is to help you prepare to take and pass Microsoft's exam number 70-294, *Planning, Implementing and Maintaining a Microsoft Windows Server 2003 Active Directory Infrastructure*. At the time of this writing, the exam is expected to be released in its beta version in June 2003. Our secondary purpose in writing this book is to provide exam candidates with knowledge and skills that go beyond the minimum requirements for passing the exam, and help to prepare them to work in the real world of Microsoft computer networking in an Active Directory domain environment.

What is Exam 70-294?

Exam 70-294 is one of the four core requirements for the Microsoft Certified Systems Engineer (MCSE) certification. Microsoft's stated target audience consists of IT professionals with at least one year of work experience on a medium or large company network. This means a multi-site network with at least three domain controllers, running typical network services such as file and print services, database, firewall services, proxy services, remote access services and Internet connectivity.

However, not everyone who takes Exam 70-294 will have this ideal background. Many people will take this exam after classroom instruction or self-study as an entry into the networking field. Many of those who do have job experience in IT will not have had the opportunity to work with all of the technologies covered by the exam. In this book, our goal is to provide background information that will help you to understand the concepts and procedures described even if you don't have the requisite experience, while keeping our focus on the exam objectives.

Exam 70-294 covers the basics of managing and maintaining the Active Directory infrastructure in a network environment that is built around Microsoft's Windows Server 2003. Objectives are task-oriented, and include the following:

- **Planning a strategy for placing global catalog servers**, including evaluating network traffic considerations and evaluating the need to enable universal group caching.

- **Planning the placement of flexible operations master roles**, including how to plan for business continuity of operations master roles and identifying operations master role dependencies.
- **Implementing an Active Directory directory service forest and domain structure**, including creating the forest root domain, creating a child domain, creating and configuring Application Data Partitions, and installing and configuring an Active Directory domain controller. This objective also includes setting an Active Directory forest and domain functional level based on requirements, and establishing trust relationships such as external trusts, shortcut trusts and cross-forest trusts.
- **Implementing an Active Directory site topology**, including configuring site links and configuring preferred bridgehead servers.
- **Planning an administrative delegation strategy**, including planning an organizational unit (OU) structure based on delegation requirements and planning a security group hierarchy based on delegation requirements.
- **Managing an Active Directory forest and domain structure**, including managing trust relationships, managing schema modifications, and adding or removing UPN suffixes.
- **Managing an Active Directory site**, including configuring replication schemes, configuring site link costs, and configuring site boundaries.
- **Monitoring Active Directory replication failures**, using tools such as Replication Monitor, Event Viewer and support tools to monitor Active Directory replication and File Replication Service (FRS) replication.
- **Restoring Active Directory directory services**, including performing both authoritative restore and nonauthoritative restore operations.
- **Troubleshooting Active Directory**, including diagnosing and resolving issues related to Active Directory replication, operations master role failure, and the Active Directory database.
- **Planning a security group strategy.**
- **Planning a user authentication strategy**, including planning a strategy for smart card authentication and creating a password policy for domain users.
- **Planning an OU structure**, including analyzing the administrative requirements for an OU and analyzing the Group Policy requirements for an OU structure.
- **Implementing an OU structure**, including creating an OU, delegating permissions for an OU to a user or a security group, and moving objects within the OU hierarchy.

- **Planning a Group Policy strategy**, including using Resultant Set of Policy (RSoP) planning mode, and strategies for configuring the user environment and computer environments using Group Policy.
- **Configuring the user environment with Group Policy**, including distributing software to users via Group Policy, automatically enrolling user certificates with Group Policy, redirecting folders via Group Policy and configuring user security settings using Group Policy.
- **Deploying a computer environment using Group Policy**, including distributing software to computers via Group Policy, automatically enrolling computer certificates with Group Policy, and configuring computer security settings using Group Policy.
- **Troubleshooting issues related to Group Policy application and deployment**, using tools such as RSoP and the gpreresult command.
- **Maintain installed software using Group Policy**, including distributing updates to software distributed by Group Policy and configuring automatic updates for network clients using Group Policy.
- **Troubleshoot the application of Group Policy security settings**, using tools such as RSoP and the gpreresult command.

Microsoft reserves the right to change the objectives and/or the exam at any time, so you should check the web site at <http://www.microsoft.com/traincert/exams/70-294.asp> for the most up-to-date version of the objectives.

Path to MCP/MCSA/MCSE

Microsoft certification is recognized throughout the IT industry as a way to demonstrate mastery of basic concepts and skills required to perform the tasks involved in implementing and maintaining Windows-based networks. The certification program is constantly evaluated and improved; the nature of information technology is changing rapidly and this means requirements and specifications for certification can also change rapidly. This book is based on the exam objectives as stated by Microsoft at the time of writing; however, Microsoft reserves the right to make changes to the objectives and to the exam itself at any time. Exam candidates should regularly visit the Certification and Training web site at <http://www.microsoft.com/traincert/> for the most updated information on each Microsoft exam.

Microsoft presently offers three basic levels of certification:

- **Microsoft Certified Professional (MCP)**: to obtain the MCP certification, you must pass one current Microsoft certification exam. For more information on exams that qualify, see <http://www.microsoft.com/traincert/mcp/mcp/requirements.asp>.

- **Microsoft Certified Systems Administrator (MCSA):** to obtain the MCSA certification, you must pass three core exams and one elective exam, for a total of four exams. For more information, see <http://www.microsoft.com/TrainCert/mcp/mcsa/requirements.asp>.
- **Microsoft Certified Systems Engineer (MCSE):** to obtain the MCSE certification on Windows Server 2003, you must pass six core exams (including four network operating system exams, one client operating system exam and one design exam) and one elective. For more information, see <http://www.microsoft.com/traincert/mcp/mcse/windows2003/>.

Passing Exam 70-294 will earn you the MCP certification (if it is the first Microsoft exam you've passed). Exam 70-294 also counts toward the MCSE. Exam 70-294 is *not* a requirement or elective for the MCSA.



NOTE

Those who already hold the MCSA in Windows 2000 can upgrade their certifications to MCSA 2003 by passing one upgrade exam (70-292). Those who already hold the MCSE in Windows 2000 can upgrade their certifications to MCSE 2003 by passing two upgrade exams (70-292 and 70-296).

Microsoft also offers a number of specialty certifications for networking professionals and certifications for software developers, including the following:

- **Microsoft Certified Database Administrator (MCDBA)**
- **Microsoft Certified Solution Developer (MCSD)**
- **Microsoft Certified Application Developer (MCAD)**

Exam 70-294 does not apply to any of these specialty and developer certifications.

Prerequisites and Preparation

There are no mandatory prerequisites for taking Exam 70-294, although Microsoft recommends that you meet the target audience profile described earlier, and many candidates will first take Exams 70-290, 70-291 and 70-293 in sequence before taking Exam 70-294 in their pursuit of the MCSE certification.

Preparation for this exam should include the following:

- Visit the web site at <http://www.microsoft.com/traincert/exams/70-294.asp> to review the updated exam objectives. Remember that Microsoft reserves the right to change or add to the objectives at any time, so new objectives might have been added since the printing of this book.

- Work your way through this book, studying the material thoroughly and marking any items you don't understand.
- Answer all practice exam questions at the end of each chapter.
- Complete all hands-on exercises in each chapter.
- Review any topics that you don't thoroughly understand
- Consult Microsoft online resources such as TechNet (<http://www.microsoft.com/technet/>), white papers on the Microsoft web site, and so forth, for better understanding of difficult topics.
- Participate in Microsoft's product-specific and training and certification newsgroups if you have specific questions that you still need answered.
- Take one or more practice exams, such as the one included on the CD with this book.

Exam Overview

In this book, we have tried to follow Microsoft's exam objectives as closely as possible. However, we have rearranged the order of some topics for a better flow, and included background material to help you understand the concepts and procedures that are included in the objectives. Following is a brief synopsis of the exam topics covered in each chapter:

- **Active Directory Infrastructure Overview:** In this chapter, we will start with the basics: defining directory services and providing a brief background of the directory services standards and protocols. You'll learn how the Active Directory works, and we will introduce you to the terminology and concepts required to understand the Active Directory infrastructure. We discuss the directory is structured into sites, forests, domains, domain trees, and organizational units, and you'll learn about the components that make up the Active Directory, including both logical and physical components. These include the schema, the global catalog, domain controllers and the replication service. You'll learn to use the Active Directory administrative tools, and we will discuss directory security and access control. Finally, we provide an overview of what's new for Active Directory in Windows Server 2003.
- **Working with User, Group and Computer Accounts:** This chapter introduces you to the concept of security principles – users, groups and computers – and the security identifiers that are used to represent them. You'll learn about the conventions and limitations for naming these objects. We show you how to work with Active Directory user accounts, including the built in accounts and those you create yourself. You'll also learn to work with group accounts, and you'll learn about group types and scopes. You'll learn to work with computer accounts, and

how to manage multiple accounts. We'll show you how to implement User Principal Name suffixes, and we'll discuss how to move objects within the Active Directory.

- **Creating User and Group Strategies:** This chapter deals with planning effective strategies for managing users and groups in Active Directory. We will discuss the creation of user authentication strategies, and we provide an overview of authentication concepts. You will learn to plan a smart card authentication strategy and find out what's new in smart card authentication for Windows Server 2003. We will also discuss how to create a password policy for domain users, and how to plan a security group strategy.
- **Working with Forests and Domains:** In this chapter, you will learn all about the functions of forests and domains in the Active Directory infrastructure, and we will walk you through the steps of creating a forest and domain structure for a network. You'll learn to install domain controllers, create the forest root domain and a child domain, and you'll find out how to name and rename domains and how to set the functional level of a forest and domain. We will then discuss the role of DNS in the Active Directory environment, and you'll learn about the relationship of the DNS and AD namespaces, how DNS zones are integrated into Active Directory, and how to configure DNS servers for use with Active Directory.
- **Working with Trusts and Organizational Units:** This chapter addresses two important components of Active Directory: trust relationships and organizational units (OUs). You'll learn about the different types of trusts that exist in the AD environment, both implicit and explicit, and you'll learn to create shortcut, external, realm and cross-forest trusts. You'll also learn to verify and remove trusts, and how to secure trusts using SID filtering. Then we discuss the creation and management of OUs and you learn to apply Group Policy to OUs and how to delegate control of an OU. We show you how to plan an OU structure and strategy for our organization, considering delegation requirements and the security group hierarchy.
- **Working with Active Directory Sites:** In this chapter, you learn about the role of sites in the Active Directory infrastructure, and how replication, authentication and distribution of services information work within and across sites. We discuss the relationship of sites and domains, the relationship of sites and subnets, and how to create sites and site links. You'll learn about site replication and how to plan, create and manage a replication topology. We walk you through the steps of configuring replication between sites, and discuss how to troubleshoot replication failures. We also address monitoring of the File Replication Service (FRS).
- **Working with Domain Controllers:** The focus of this chapter is the Active Directory domain controller (DC), and how to plan and deploy DCs on your net-

work. You'll learn about server roles, where domain controllers fit in, and how to create and upgrade DCs. We discuss placement of domain controllers within sites and how to back up your domain controllers. Then we get into the subject of operations master (OM) roles and you learn about the functions of all five OMs: the schema master, domain naming master, RID master, PDC emulator and infrastructure master. We talk about transferring and seizing master roles and role dependencies, and you'll learn to plan for the placement of OMs and how to respond to OM failures.

- **Working with Global Catalog Servers and the Schema:** In this chapter, we take a look at a special type of domain controller: the Global Catalog server. You'll learn about the role the Global Catalog (GC) plays in the network, and you'll find out how to customize the GC using the Schema MMC snap-in. We show you how to create and manage GC servers, and explain how GC replication works. You'll learn about the factors to consider when placing GC servers within sites. Next, we address the Active Directory schema itself. You'll learn about schema components: classes and attributes, and the naming of schema objects. We show you how to install and use the Schema management console, and you'll learn how to extend the schema and how to deactivate schema objects.
- **Working with Group Policy in an Active Directory Network:** This chapter starts with the basics of Group Policy terminology and concepts, introducing you to user and computer policies and Group Policy Objects (GPOs). We discuss the scope and application order of policies and you'll learn about Group Policy integration in Active Directory. We show you how to plan a Group Policy strategy, and then walk you through the steps of implementing Group Policy. We show you how to perform common Group Policy tasks, and discuss Group Policy propagation and replication. You'll also learn best practices for working with Group Policy, and we'll show you how to troubleshoot problems with Group Policy.
- **Deploying Software via Group Policy:** In this chapter, you will learn about Group Policy's software installation feature. We'll show you how to use the components of software installation: Windows installer packages, transforms, patches and application assignment scripts. You'll find out how to deploy software to users and to computers, by assigning or publishing applications. We walk you through the steps of preparing for GP software installation, working with the Group Policy Object Editor and setting installation options. You'll find out how to upgrade applications, configure automatic updates and remove managed applications. We'll also cover how to troubleshoot problems that can occur with Group Policy software deployment.
- **Ensuring Active Directory Availability:** The final chapter deals with how to maintain high availability of your Active Directory services. You'll learn about the

Active Directory database, and the importance of system state data to AD availability. We'll discuss fault tolerance plans as well as AD performance issues. You'll find out how to perform necessary maintenance tasks, such as defragging the database, and you'll learn how to monitor or move the database. We address backup and restoration of the Active Directory, and show you the different restoration methods that can be used and when each is appropriate. Finally, you'll learn to troubleshoot Active Directory availability.

Exam Day Experience

Taking the exam is a relatively straightforward process. Both Vue and Prometric testing centers administer the Microsoft 70-291 exam. You can register for, reschedule or cancel an exam through the Vue web site at <http://www.vue.com/> or the Prometric web site at <http://www.2test.com/index.jsp>. You'll find listings of testing center locations on these sites. Accommodations are made for those with disabilities; contact the individual testing center for more information.

Exam price varies depending on the country in which you take the exam.

Exam Format

Exams are timed. At the end of the exam, you will find out your score and whether you passed or failed. You will not be allowed to take any notes or other written materials with you into the exam room. You will be provided with a pencil and paper, however, for making notes during the exam or doing calculations.

In addition to the traditional multiple choice questions and the select and drag, simulation and case study questions introduced in the Windows 2000 exams, Microsoft has developed a number of innovative question types for the Windows Server 2003 exams. You might see some or all of the following types of questions:

- *Hot area* questions, in which you are asked to select an element or elements in a graphic to indicate the correct answer. You click an element to select or deselect it.
- *Active screen* questions, in which you change elements in a dialog box (for example, by dragging the appropriate text element into a text box or selecting an option button or checkbox in a dialog box).
- *Drag and drop* questions, in which you arrange various elements in a target area.

You can download a demo sampler of test question types from the Microsoft web site at <http://www.microsoft.com/traincert/mcpexams/faq/innovations.asp#H>.

Test Taking Tips

Different people work best using different methods. However, there are some common methods of preparation and approach to the exam that are helpful to many test-takers. In this section, we provide some tips that other exam candidates have found useful in preparing for and actually taking the exam.

- Exam preparation begins before exam day. Ensure that you know the concepts and terms well and feel confident about each of the exam objectives. Many test-takers find it helpful to make flash cards or review notes to study on the way to the testing center. A sheet listing acronyms and abbreviations can be helpful, as the number of acronyms (and the similarity of different acronyms) when studying IT topics can be overwhelming. The process of writing the material down, rather than just reading it, will help to reinforce your knowledge.
- Many test-takers find it especially helpful to take practice exams that are available on the Internet and with books such as this one. Taking the practice exams not only gets you used to the computerized exam-taking experience, but also can be used as a learning tool. The best practice tests include detailed explanations of why the correct answer is correct and why the incorrect answers are wrong.
- When preparing and studying, you should try to identify the main points of each objective section. Set aside enough time to focus on the material and lodge it into your memory. On the day of the exam, you be at the point where you don't have to learn any new facts or concepts, but need simply to review the information already learned.
- The value of hands-on experience cannot be stressed enough. Exam questions are based on test-writers' experiences in the field. Working with the products on a regular basis, whether in your job environment or in a test network that you've set up at home, will make you much more comfortable with these questions.
- Know your own learning style and use study methods that take advantage of it. If you're primarily a visual learner, reading, making diagrams, watching video files on CD, etc. may be your best study methods. If you're primarily auditory, classroom lectures, audiotapes you can play in the car as you drive, and repeating key concepts to yourself aloud may be more effective. If you're a kinesthetic learner, you'll need to actually *do* the exercises, implement the security measures on your own systems, and otherwise perform hands-on tasks to best absorb the information. Most of us can learn from all of these methods, but have a primary style that works best for us.
- Although it might seem obvious, many exam-takers ignore the physical aspects of exam preparation. You are likely to score better if you've had sufficient sleep the night before the exam, and if you are not hungry, thirsty, hot/cold or otherwise distracted

by physical discomfort. Eat prior to going to the testing center (but don't indulge in a huge meal that will leave you uncomfortable), stay away from alcohol for 24 hours prior to the test, and dress appropriately for the temperature in the testing center (if you don't know how hot/cold the testing environment tends to be, you may want to wear light clothes with a sweater or jacket that can be taken off).

- Before you go to the testing center to take the exam, be sure to allow time to arrive on time, take care of any physical needs, and step back to take a deep breath and relax. Try to arrive slightly early, but not so far in advance that you spend a lot of time worrying and getting nervous about the testing process. You may want to do a quick last minute review of notes, but don't try to "cram" everything the morning of the exam. Many test-takers find it helpful to take a short walk or do a few calisthenics shortly before the exam, as this gets oxygen flowing to the brain.
- Before beginning to answer questions, use the pencil and paper provided to you to write down terms, concepts and other items that you think you may have difficulty remembering as the exam goes on. Then you can refer back to these notes as you progress through the test. You won't have to worry about forgetting the concepts and terms you have trouble with later in the exam.
- Sometimes the information in a question will remind you of another concept or term that you might need in a later question. Use your pen and paper to make note of this in case it comes up later on the exam.
- It is often easier to discern the answer to scenario questions if you can visualize the situation. Use your pen and paper to draw a diagram of the network that is described to help you see the relationships between devices, IP addressing schemes, and so forth.
- When appropriate, review the answers you weren't sure of. However, you should only change your answer if you're sure that your original answer was incorrect. Experience has shown that more often than not, when test-takers start second-guessing their answers, they end up changing correct answers to the incorrect. Don't "read into" the question (that is, don't fill in or assume information that isn't there); this is a frequent cause of incorrect responses.
- As you go through this book, pay special attention to the Exam Warnings, as these highlight concepts that are likely to be tested. You may find it useful to go through and copy these into a notebook (remembering that writing something down reinforces your ability to remember it) and/or go through and review the Exam Warnings in each chapter just prior to taking the exam.
- Use as many little mnemonic tricks as possible to help you remember facts and concepts. For example, to remember which of the two IPsec protocols (AH and ESP) encrypts data for confidentiality, you can associate the "E" in encryption with the "E" in ESP.

Pedagogical Elements

In this book, you'll find a number of different types of sidebars and other elements designed to supplement the main text. These include the following:

- **Exam Warning** These focus on specific elements on which the reader needs to focus in order to pass the exam (for example, “Be sure you know the difference between symmetric and asymmetric encryption”).
- **Test Day Tip** These are short tips that will help you in organizing and remembering information for the exam (for example, “When preparing for the exam on test day, it may be helpful to have a sheet with definitions of these abbreviations and acronyms handy for a quick last-minute review”).
- **Configuring & Implementing** These are sidebars that contain background information that goes beyond what you need to know from the exam, but provide a “deep” foundation for understanding the concepts discussed in the text.
- **New & Noteworthy** These are sidebars that point out changes in W2003 Server from the old Windows 2000/NT family, as they will apply to readers taking the exam. These may be elements that users of W2K/NT would be very familiar with that have changed significantly in W2003 Server, or totally new features that they would not be familiar with at all.
- **Head of the Class** These are discussions of concepts and facts as they might be presented in the classroom, regarding issues and questions that most commonly are raised by students during study of a particular topic.

The book also includes, in each chapter, hands-on exercises in planning and configuring the features discussed. It is essential that you read through and, if possible, perform the steps of these exercises to familiarize yourself with the processes they cover.

You will find a number of helpful elements at the end of each chapter. For example, each chapter contains a *Summary of Exam Objectives* that ties the topics discussed in that chapter to the published objectives. Each chapter also contains an *Exam Objectives Fast Track*, which boils all exam objectives down to manageable summaries that are perfect for last minute review. *The Exam Objectives Frequently Asked Questions* answers those questions that most often arise from readers and students regarding the topics covered in the chapter. Finally, in the *Self Test* section, you will find a set of practice questions written in a multiple-choice form that will assist you in your exam preparation. These questions are designed to assess your mastery of the exam objectives and provide thorough remediation, as opposed to simulating the variety of question formats you may encounter in the actual exam. You can use the *Self Test Quick Answer Key* that follows the *Self Test* questions to quickly determine what information you need to review again. The *Self Test Appendix* at the end of the book provides detailed explanations of both the correct and incorrect answers.

Additional Resources

There are two other important exam preparation tools included with this Study Guide. One is the DVD included in the back of this book. The other is the practice exam available from our Web site.

- **Instructor-led training DVD provides you with almost two hours of virtual classroom instruction.** Sit back and watch as an author and trainer reviews all the key exam concepts from the perspective of someone taking the exam for the first time. Here, you'll cut through all of the noise to prepare you for exactly what to expect when you take the exam for the first time. You will want to watch this DVD just before you head out to the testing center!
- **Web based practice exams.** Just visit us at www.syngress.com/certification to access a complete Windows Server 2003 concept multiple choice review. These remediation tools are written to test you on all of the published certification objectives. The exam runs in both "live" and "practice" mode. Use "live" mode first to get an accurate gauge of your knowledge and skills, and then use practice mode to launch an extensive review of the questions that gave you trouble.

Chapter 1

MCSA/MCSE 70-294

Active Directory Infrastructure Overview

Exam Objectives in this Chapter:

- 1 Planning and Implementing an Active Directory Infrastructure
 - ☑ Summary of Exam Objectives
 - ☑ Exam Objectives Fast Track
 - ☑ Exam Objectives Frequently Asked Questions
 - ☑ Self Test
 - ☑ Self Test Quick Answer Key


Introduction

The Active Directory is the foundation of an enterprise-level Windows network, and Windows Server 2003 includes a number of improvements and enhancements to its directory services that will make a network administrator's job easier. Exam candidates must understand the basics of how directory services work and the role they play in the network, and specifically how the directory services concept is implemented in Microsoft's Active Directory.

In this chapter, we start with the basics by defining directory services and providing a brief background of the directory services standards and protocols. You'll learn how the Active Directory works, and be introduced to the terminology and concepts required to understand the Active Directory infrastructure.

We discuss how the directory is structured into sites, forests, domains, domain trees, and organizational units (OUs), and you'll learn about the components that make up the Active Directory, including both logical and physical components. These include the schema, the Global Catalog (GC), domain controllers (DCs), and the replication service. You'll learn to use the Active Directory administrative tools, and we discuss directory security and access control. Finally, we provide an overview of what's new for Active Directory in Windows Server 2003.

This chapter lays the groundwork for the specific Active Directory-related administrative tasks that you will learn to perform throughout the rest of the book.



Introducing Directory Services

As anyone familiar with networking knows, a network can be comprised of a vast number of elements, including user accounts, file servers, volumes, fax servers, printers, applications, databases, and other shared resources. Because the number of objects making up a network increases as an organization grows, finding and managing these accounts and resources becomes harder as the network gets bigger. To make a monolithic enterprise network more manageable, directory services are used to store a collection of information about users and resources, so they are organized and accessible across the network.

A directory allows accounts and resources to be organized in a logical, hierarchical fashion so that information can be found easily. By searching the directory, users can find the resources they need, and administrators are able to control and configure accounts and resources easily and effectively. Keeping this information in a centralized location ensures that users and administrators don't have to waste time looking at what's available on each server, they only have to refer to the directory.

At face value, the concept of directory services seems overwhelming. However, even if you're unfamiliar with directory services, you're probably familiar with directories in general. In a telephone directory, every account a telephone company manages is uniquely identified by a telephone number, and includes attributes such as the person's name and address. Each account needs to be uniquely identified, so one isn't mixed up with

another—you wouldn't want to dial one person's telephone number, only to be connected with someone else. To make it easier to find information, the telephone directory is structured to look up information in multiple ways. You can look up someone's name and then view his or her telephone number, or you can search for entries using other attributes, such as using the yellow pages and viewing categories of businesses. The same basic concept applies to directory services.

Any directory (regardless of what it's used for) is a structured source of information, consisting of objects and their attributes. As in the case of a telephone directory, a network directory contains uniquely identified objects with different attributes. Such data can be made available to applications, operating system services, network administrators, and other authorized users. Those who have access to the directory can look up an object, and then view its attributes. If they have sufficient rights (as in the case of an administrator), the object can be modified. These attributes can be used to provide information that's accessible to users, or control security at a granular level.

The objects and attributes in a directory can be used in various ways. For example, a user might need to use a color printer, but not know the printer's name. To find this printer, the user might know it is located on the second floor, and search the directory for an attribute with this information. In the same light, a user account can include attributes such as usernames, passwords, the user's name, address, telephone numbers, and other relevant information about the person. If a person has access to view another user object's attributes, he or she can access this data and find information on how to contact the other user.

Because a user can access account information from anywhere on the network, directory services allow a user to log on to multiple servers using a single logon. A single logon is an important feature to directory services, because without it, a user must log on to each server that provides needed resources. This is common on Windows NT networks, where the administrator must create a different account on each server the user needs to access. The user then needs to log on to each server individually. This is significantly different from the way Windows 2000/2003's directory services work, where a user logs on to the network once and can use any of the resources to which he or she has been given access.

Sophisticated directory services give administrators the ability to organize information, control security, and manage users and resources anywhere on the network. Information resides in a central repository that's replicated to different servers on the network. It allows the data to be accessed when needed and saves the administrator from having to visit each server to manage accounts. This lowers the amount of work needed to manage the network, while providing granular control over rights and permissions. The administrator only needs to modify a user account or other object once, and these security changes are replicated throughout the network.

Directory services have been used on different network operating systems for years, and have proven to be a useful and powerful technology. Following suit, Microsoft created its own implementation of directory services on Windows NT called NTDS, and then followed with Active Directory on newer versions of servers. NTDS used a flat namespace, which provided limited functionality in comparison with Active Directory's hierarchical

structure and feature set. Active Directory was first introduced in Windows 2000, and continues to provide directory services to the Windows Server 2003 family of servers. It can be installed on the Standard, Enterprise, and Datacenter Editions of Windows Server 2003, and provides a necessary foundation for any network using these servers.



NOTE

Installation of Active Directory on a Windows 2000 or Windows Server 2003 server makes that computer a DC. Windows Server 2003 Web Edition cannot function as a DC, and thus cannot have Active Directory installed.

Head of the Class...

A Brief History of Directory Services

Directory services have been around long since Microsoft's implementation of Active Directory. In 1984, Banyan offered customers the first directory service for enterprise networking. The product was called StreetTalk, and provided enterprise directory services for networks running Banyan VINES. Since then, many other network operating systems have evolved to use directory services as a method of storing information related to the components of a network.

Novell also provided directory services for its network operating system NetWare. In 1993, NetWare 4 introduced an object-oriented directory called NetWare Directory Services, which was later named Novell Directory Services (NDS). NDS used a hierarchical structure and provided the basis for new features in NetWare. NDS evolved into eDirectory, which provided greater features and interoperability between different operating systems that used its directory services.

Microsoft also incorporated a directory service in its network operating system. In 1993, the first release of Windows NT Server included the Windows NT Directory Services (NTDS). NTDS provided a single point of administration that allowed you to manage up to 25,000 users per domain. When Windows 2000 Server was released, the number of objects supported by directory services jumped dramatically. Active Directory theoretically supports up to 10 million objects per domain, with 1 million objects being a more practical estimate. In addition, Active Directory uses a hierarchical namespace and provides significantly more features than the directory services in Windows NT.

When directory services were still new to networking in the mid 1980s, there was a lack of common standards to control the development of directory services. Different standards were being used to determine how directories should function. The International Telecommunications Union (which was called the International Telegraph and Telephone Consultative Committee at the time) was developing a directory that allowed information (telephone numbers and other data) to be looked up from a directory service. The International Organization for

Continued

Standardization (ISO) was also developing standards for network functions and applications using the Open Systems Interconnect (OSI). These different specifications came together after a time to become X.500.

Part of the X.500 standard specified how information in directories was accessed. The Directory Access Protocol (DAP) was developed to allow communication between clients and directory servers, but the number of features in DAP made it inefficient with personal computers. By removing some of the code from DAP, and limiting the features, a more lightweight version of DAP was developed, called the Lightweight Directory Access Protocol (LDAP). This protocol and the X.500 standard are used by Active Directory and other directory services, and provide interoperability between different systems.

Terminology and Concepts

Before delving too far into the specifics of Active Directory, it is important to discuss a number of concepts and terms to appreciate the features and functionality of a directory service. As with anything dealing with technology, certain words and phrases associated with Active Directory and Windows Server 2003 are useful in identifying and defining specific components of the network. Whether you're new to Active Directory or experienced from using previous versions, the information provided here will help you to understand other topics that follow in this book.

In reading this section, it is important to realize that this is an overview of topics that we discuss later in greater detail. We define some of the terms used throughout this book, and look at concepts that we'll build on in later sections.

Some of the terms and concepts we discuss in the following subsections include:

- Directory data store
- Directory partitions
- Policy-based administration
- DAP and LDAP
- Naming schemes used in Active Directory

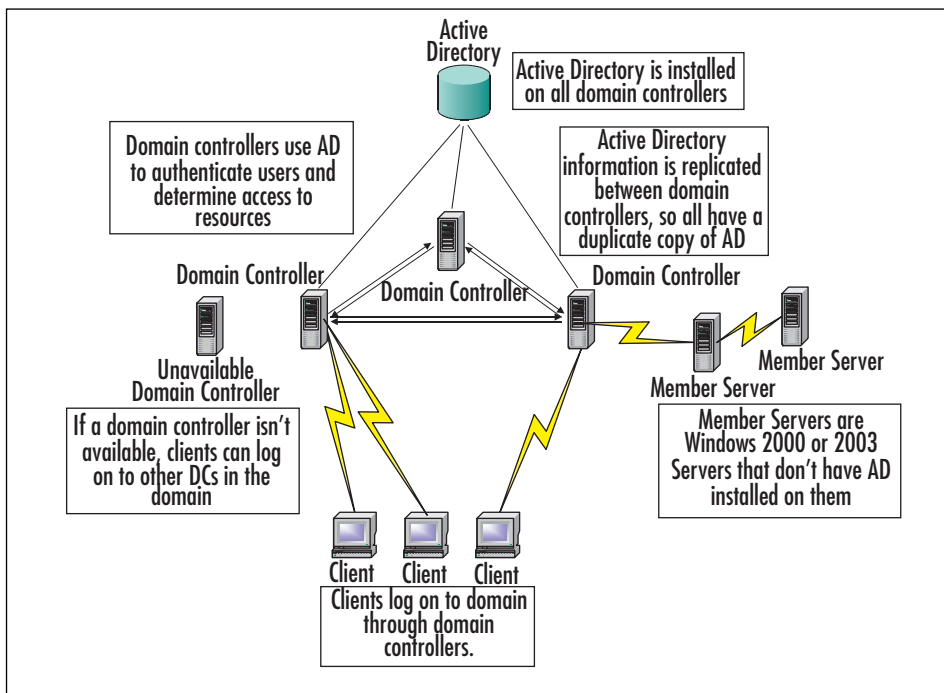
Directory Data Store

Active Directory isn't just a service that provides access to directory services; it's also a method of storing data about network elements. If you didn't have a place where configurations and directory data are saved, you'd lose this information every time you shut down your server. The data store contains a vast amount of information, including data dealing with users, groups, computers, the resources they can access, and other components of the network. Because the Active Directory data store is a database of all directory information, it is also referred to as the directory.

When you install the directory on a Windows Server 2003 server, the Active Directory data source is placed on the server's hard disk. The file used to store directory information is called NTDS.DIT, and is located in the NTDS folder in the systemroot (for example, C:\WINDOWS). Any changes made to the directory are saved to this file.

The presence of Active Directory's data store on a Windows Server 2003 server has a major impact on that server's role in the network. As shown in Figure 1.1, the directory is stored on DCs, which are servers with writable copies of the data store. A DC is used to manage domains, which are groups of computers, users, and other objects that share (or are included in) the same directory. Domains that use different Active Directory data sources can still communicate with one another, but (as we'll see later in this chapter) secure relationships between them must be configured.

Figure 1.1 Relationship Between Active Directory, Domain Controllers, Member Servers, and Clients

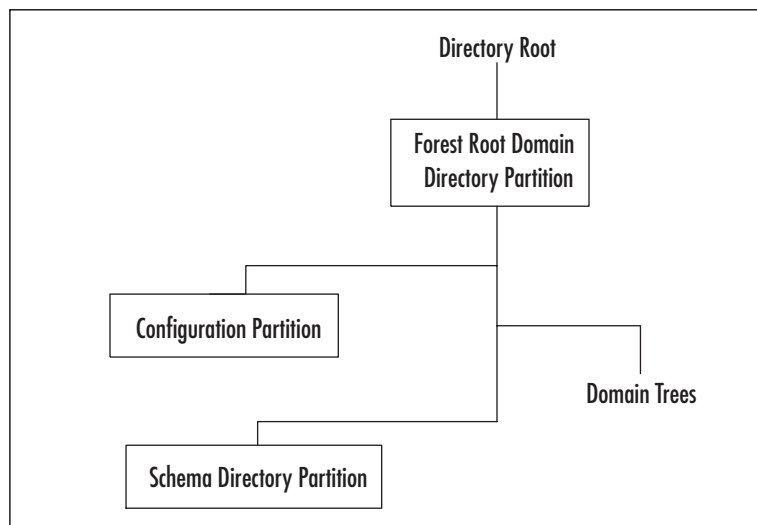


Each DC retains its own copy of the directory, containing information on the domain in which it is located. If one DC becomes unavailable, users and computers can still access the Active Directory data store on another DC in that domain. This allows users to continue logging on to the network even though the DC that's normally used is unavailable. It also allows computers and applications that require directory information to continue functioning while one of these servers is down.

Because a domain can have more than one DC, changes made to the directory on one DC must be updated on others. The process of copying these updates is called *replication*, and is used to synchronize information in the directory. Without replication, features in Active Directory would fail to function properly. For example, if you added a user on one DC, the new account would be added to the directory store on that server. This would allow the user to log on to that domain controller, but he or she still couldn't log on to other DCs until these changes to the directory were replicated. When a change is made on one DC, the changes need to be replicated quickly so that each DC continues to have an accurate duplicate copy of Active Directory.

Because replication is so important to making the directory consistent across the network, the data source is organized in a way to make replication more efficient. Not every piece of data is saved in the same location of the data source. As shown in Figure 1.2, information resides in different areas of the directory, called *directory partitions*. Because Active Directory is a logical, hierarchical structure, it has a treelike structure similar to that of the Windows Registry or folders on a hard disk. Data is stored within subtrees of the directory, much like data on your hard disk is stored within folders that are nested within one another. Each contiguous subtree in the directory is a partition. Any data that changes within a directory partition is replicated as a single unit to other DCs.

Figure 1.2 Active Directory Is a Hierarchical Structure



In Active Directory, three partitions exist on any DC and must be replicated, as these contain data that the Microsoft network needs to function properly:

- Domain partition
- Configuration partition
- Schema partition



NOTE

Directory partitions shouldn't be confused with disk partitions. Disk partitions are used to logically separate the hard disk so that multiple operating systems can run from a single hard drive, or to make it appear that a single hard disk is actually more than one. For example, if you have a single hard disk on your computer, the C: drive, you might partition it into two logical drives, so that the system now has C: and D: drives. While this is useful for organizing files stored on the hard disk, it is unrelated to how Active Directory is partitioned.

A directory is physically a database where information is stored within a table, and logically, a naming hierarchy made up of containers and objects. Information is partitioned in this manner to allow data to be stored in different areas. As we'll see in the paragraphs that follow, these partitions consist of containers storing domain, configuration, schema, and application information.

The domain partition contains information about the domain. This includes information about users, computers, resources, and attributes associated with each. Without this data being replicated, any changes would be limited to the server on which the changes were made, and other servers would use older settings. For example, if the domain data wasn't replicated and you disabled a user's account on one DC, the user would still be able to log on to other DCs. The domain partition is important because it contains information about objects and their attributes, which are fundamental elements of your network.

Configuration data deals with the topology of Active Directory, and includes information about how the domains, domain trees, and forests within a network are configured. A domain tree is a structure of domains. If more than one domain is in a domain tree, trusts are set up between those domains so that they can share data and resources between them. A forest also consists of multiple domains that share directory data. It consists of one or more trees that are connected through trusts. The configuration partition also includes information about the locations of DCs and the GC, which is a subset of the data contained in Active Directory that is used to provide search and logon functionality across multiple domains. We discuss each of these topics in greater detail later in this chapter.

Because Active Directory is made up of different objects, and each object has specific attributes, certain rules must be created to control what objects can exist in the directory, and the attributes of each. For example, a user account has attributes that include a password, an account name, and the first and last of the person to whom the account belongs. The types of objects that exist in Active Directory, and which attributes each type has, is determined by the schema. The *schema partition* contains information that defines object classes and attributes used within the domain. It determines what objects can exist within Active Directory, and what attributes each can have.

Windows Server 2003 servers can also create one or more application partitions, which are used to store data that is specific to different applications running on the network. Programs can use this partition to store settings that are needed while the programs are running on a server. We discuss this in greater detail later in the chapter.

Providing the Best Possible Protection for Your Directory

In addition to Windows 2000 servers, Active Directory can only be installed on Microsoft Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. When a server is configured to be a DC on any of these editions, a writable copy of the directory is stored on the server's hard disk. Because any file can be damaged, destroyed, or compromised (such as in the case of a hacking attempt or virus), you should take steps to ensure that the directory is safe on your server(s).

If only one DC is used, then only one NTDS.DIT file will exist, meaning there is only one copy of the directory for that domain. Failure of this server or damage to the NTDS.DIT file will disable the network. Users will be unable to log on, computers will be unable to access needed information from the directory, and any configurations on your network could be lost. Rather than hoping that nothing ever happens to your one DC, it is wise to use multiple DCs on your network.

If more than one DC exists in a domain, any updates to the NTDS.DIT will be replicated to other DCs. This will allow multiple copies of the directory to exist on the network, providing a level of fault tolerance if one server fails. If one fails, another can continue authenticating users, supplying services, and providing access to resources.

Because of the importance of the NTDS.DIT file, the drive on which it is stored should be formatted in NTFS format. NTFS is a file format that allows the best possible level of protection, allowing you to set permissions on who can access the directory and NTDS.DIT file locally and across the network. Such permissions cannot be set on hard disks that are formatted as FAT16 or FAT32. Limiting the access to this file lessens the chance that someone might accidentally or maliciously damage or delete the data source.

It is also important to remember that any measures you take to protect Active Directory from harm do not negate the need to perform regular backups. When backups are performed, the data on a computer is copied to other media (such as a tape, CD, or DVD), which can then be stored in another location. Should any problem occur, you can restore any files that were damaged or lost.

Policy-Based Administration

There can be hundreds or thousands of users and computers in a large network. Having to go through each account and configure settings can be an arduous task. For example, imagine having to go to each computer to change the desktop so that it displays a company logo as the background image. Rather than visiting each computer, it would be far easier to make such changes in one location, and have these settings apply to everyone. This is why policy-based administration is such a benefit to Active Directory: it makes managing accounts easier.

Group policies allow you to apply default settings to groups of users and groups. Policies can be used to:

- Control desktop settings that determine the display properties of a computer.
- Assign scripts that run at logon, logoff, startup, and shutdown.
- Enforce password security, such as by setting minimum password lengths, maximum length of time before a password must be changed, and so on.
- Redirect folders from the local computer to a folder on a networked computer, such as when the My Documents folder is redirected to use specific folders on a server.
- Deploy applications, so that certain members have programs available to them to install or have them automatically installed.

As we'll see in the chapters that follow, these are just a few of the options available to administrators in managing users and computers on a network.

When policies are created, they are stored as Group Policy Objects (GPOs) in Active Directory. The settings in a GPO can be applied to a site, domain, or OU. An OU is a container in Active Directory that can contain users, groups, computers, or other OUs. We'll discuss OUs in greater detail later in this chapter. Because GPOs can be applied at different levels, you can set different policies for different areas of your company. For example, you could create a group policy for users in Finance and another for the Sales department (by placing Finance users in one OU and Sales users in another). If you have different domains for different branch offices, you could have different settings for the Sales divisions in each domain. Using GPOs in this manner, you can configure which settings will be used for specific groups of users and computers.

Directory Access Protocol

For clients to search for objects, update information, and communicate with DCs when logging on to the network, a directory access protocol must be used. A protocol is a set of rules that dictate how data is sent over a network. A directory access protocol is used for the specific purpose of exchanging information with the directory service.

Active Directory uses LDAP for communications between clients and directory servers. LDAP is a version of the X.500 Directory Access Protocol (DAP), and is considered lightweight because it uses less code than DAP does.

X.500 is an established standard that defines directory services. It was developed by the International Telecommunication Union and published by the International Organization for Standardization. DAP is one of the protocols defined by X.500. Because it contains too many features to be efficient with personal computers, thin clients, and communication over the Internet, a more lightweight version was developed that contained fewer features. As mentioned, this light version of DAP is LDAP.

The Internet Engineering Task Force (IETF) established industry standards for LDAP, enabling LDAP to be used over local networks and the Internet by a variety of directory services. Many network operating systems that use directory services (including Novell NetWare, Windows 2000, and Windows Server 2003) implement LDAP for accessing the directory, while other products (such as Internet browsers) support it as a method for finding resources or managing the directory. Since its inception in 1994, there have been several versions of LDAP, with features being added to accommodate changing needs. Active Directory supports versions 2 and 3.

Naming Scheme

Active Directory supports several common formats for naming objects. By using different methods of naming objects, it allows objects to be accessed in a variety of ways. Providing different naming schemes also provides backward compatibility to older systems that might not support one or more of these formats. The naming schemes supported by Active Directory include:

- Domain Name System (DNS)
- User principal name (UPN)
- Universal Naming Convention (UNC)
- Uniform Resource Locator (URL)
- Lightweight Directory Access Protocol Uniform Resource Locator (LDAP URL)

DNS is used to resolve domain names to IP addresses. You're probably familiar with this method from using the Internet. It would be difficult or impossible to remember the IP address for every Web site or server that you want to access. DNS is a hierarchical, distributed database that allows clients to find a particular resource by entering a user-friendly domain name (such as `knightware.ca`), and then looks up the IP address for that domain.

In Active Directory, domains are usually given DNS names (such as `syngress.com`). Because Windows domains didn't use this naming scheme prior to Windows 2000, each domain is also given a name that's compatible with those used in Windows NT networks. These pre-Windows 2000 names are NetBIOS names, and are one-word names that users of older operating systems can use to log on to Active Directory. This allows clients to log on to domains by entering the domain name and username using the format: *domain name\username*.

UPNs are based on the IETF's RFC 822. Each user account in Active Directory has a logon name and UPN suffix. The logon name is the account name, and the UPN suffix is the domain that the user will log on to. The two are connected by the `@` symbol, making the logon appear like an Internet e-mail address (`username@domain`). After entering a username, the user will generally be required to enter a password to prove that he or she is authorized to use this account.

When the UPN is created for a user account, Active Directory also suggests a pre-Windows 2000 logon name that is used by the Security Account Manager (SAM) to log on to a server. The SAM is a service that stores information about user accounts and groups to which they belong. Local computer accounts use the SAM to store accounts that are used to access the local computer, and Windows NT servers use it for allowing network users access to resources on the server. Although you can create your own logon name, Active Directory will suggest a pre-Windows 2000 user logon name that's based on the first 20 bytes of the Active Directory logon name.

Every computer account that is created in Active Directory also has multiple names, so that the account can be identified and accessed in a variety of ways. When a computer account is created in Active Directory, you need to enter a name for the computer, which will uniquely identify it in the domain. This is the host name for the machine, which can be used by DNS to indicate its place in the domain, and can be used to help find the computer when clients search for it and its resources on the network.

In DNS, the host name is combined with the domain name to create the computer's fully qualified domain name (FQDN). This combines the host name with the domain name, and separates the two with a period. For example, if you have a computer named COMP100 in the domain called knightware.ca, the FQDN for this computer would be comp100.knightware.ca. No two computers in a domain can have the same name, as this would create conflicts.

When the computer account is created, it will also require the computer be given a pre-Windows 2000 name, so older clients and servers can identify and access it. As with user accounts, Windows Server 2003 will suggest a name, which is based on the first 15 bytes of the name used to create the account. If you don't want to use this default name, you can enter a new one at any time.

The UNC path is a tried-and-true method of accessing shared resources over a network. It uses the format of two backslashes, followed by the domain name or server name, the name of the share, and (where applicable) the name of the resource. The shared resource is often the name of a shared directory, and might be followed by the name of a file, application, or other resource on the server. In other words, the format would be `\\domain name\share\filename` or `\\servername\share\filename`. For example, if you were accessing a file named SPREAD-SHEET.XLS in a shared directory called XLS on a server named FS-GOTHAM, the UNC to access it would be `\\fs-gotham\xls\spreadsheet.xls`. You can use UNC names in the address bar of browsers, from the Run command of the Windows Start menu, or any other place where UNC names are allowed.

Another common method of accessing resources through a browser is by using URLs. If you've surfed the Web, you're likely to have used a URL, as they're entered into the browser's address bar to specify Web site addresses (for example, `http://knightware.ca/default.htm`). By using a URL, you can access files, applications, or other resources. A URL generally begins with `http` (for HyperText Transfer Protocol), a colon, and two forward slashes, followed by a server name such as `www`, a domain name such as `syngress.com`, and a filename path (which

can contain a directory name such as files, or just a filename such as file.htm or file.html for an HTTP document, file.asp for an Active Server Pages document, or file.jpg for a graphic in JPG format).



NOTE

Note the difference between a UNC path, which uses backslashes to separate the components of the name, and URLs, which use forward slashes.

The final naming scheme we'll discuss is LDAP URL. This method is similar to using URLs, but uses the X.500 naming structure to locate a resource. An LDAP URL uses the format LDAP://domain name/CN=common name/OU=organizational unit/DC=domain component. In this format, the common name is the name of an object in Active Directory, OU is the organizational unit, and DC is the DNS domain name in which the object exists. This allows you to specify an object that is uniquely identified in the directory. As we'll see in the sections that follow, this information is built on X.500/LDAP standards.

X.500/LDAP standards

Both the X.500 DAP and LDAP work by interacting with the directory. The directory is designed as a hierarchy, and has a tree-like structure called the *directory information tree*. Information in subtrees branch off the trunk, much as folders on the hard disk branch off a root directory. These subtrees contain objects that represent elements of the network, and are called *directory service entries*. Just as there can't be two files with the same name in a folder on your hard disk, each object must have a unique name in the directory structure.

Distinguished Name

To accommodate the need for each object being identified with a unique name in the directory, objects have a distinguished name (DN). A DN represents the exact location of an object within the directory. This is comparable to a file being represented by the full path, showing where it is located on the hard disk. With an object in the directory, several components are used to create this name:

- **CN** The common name of the object, and includes such things as user accounts, printers, and other network elements represented in the directory.
- **OU** The organizational unit. These are containers in the directory that are used to hold objects. To continue with our example of files on a hard disk, this would be comparable to a folder within the directory structure.
- **DC** A domain component. This is used to identify the name of the domain or server, and the DNS suffix (for example, .com, .net, .edu, and .gov).

When combined, these components of the DN are used to show the location of an object. Each DN can be used more than once to fully identify the object's place within the directory. For example, let's say a user account named BobSmith was stored in the Accounting OU in the syngress.com domain. In this case, the DN of this object would be:

```
CN=BobSmith, OU=Accounting, DC=syngress, DC=com
```

Relative Distinguished Name

An RDN is a portion of the DN, and is used to uniquely identify an object with a parent container. As each object must have a unique name with the directory structure, the RDN identifies an object within a particular OU. This is comparable to a file in a folder, where you specify the name of the file and not the full path to it. Just as a file in one folder might have the same name as a file in another folder, an object in one OU might have the same name as another object in another OU. While the RDN would be the same, the DN would indicate that each is in a different OU.

To illustrate this, let's look at the previous example, which used the DN /CN=BobSmith, /OU=Accounting, /DC=syngress, /DC=com. In this case, CN=BobSmith is the RDN of the object. It is a subset of the DN, and the only one by that name in the Accounting OU. However, you could have a user account named BobSmith in the Sales OU. Even though the RDNs are identical, the full DNs are unique.

DNS and RDNs apply to user accounts and any other objects within the directory. When a computer account is created, the name used for the computer is used by LDAP as the RDN. For example, if a computer were named COMP100, this would be its RDN.

Canonical Name

A canonical name is another way of showing the DN of an object. It contains the same information, but shows it in a way that is easier to read. Using the example of the BobSmith object, if we convert its DN to a canonical name, it would read:

```
syngress.com/Accounting/BobSmith
```

In the preceding example, the CN, OU, and DC components of the DN have been removed and replaced with slashes (similar to the way in which a pathname to a file on a DOS/Windows machine is notated with backslashes). The canonical format also reverses the information. Rather than beginning with the lowest level component of the DN (in other words, the object) and moving up through higher levels, it starts at the highest level of the directory structure and works its way down to the object's name. While it relates the DN of an object, it removes the extraneous notations in the name and makes it easier to read.



TEST DAY TIP

To avoid confusion between DNs, RDNs, and canonical names, try to remember that they're all related in some way. The DN shows the exact location in the direc-

tory, while the RDN identifies the object in relation to its location in the directory. Canonical names are similar to DNSs, but show the information in a different format.

Installing Active Directory to Create a Domain Controller

When Windows Server 2003 is installed on a computer, it doesn't mean that the directory is also installed. Active Directory is installed when you create a DC. It can be installed as part of the Windows Server 2003 installation, and can also be installed on member servers, which are computers running Windows Server 2003 that don't have Active Directory installed. A server without Active Directory installed on it can still deliver a variety of services, file storage, and access to other resources, but until Active Directory is installed, it can't authenticate users or provide the other functions of a DC. Once Active Directory is installed, the member server ceases to be a member server and becomes a DC.

To install Active Directory on a member server, the Active Directory Installation Wizard (DCPROMO) is used. DCPROMO is a tool that promotes a member server to DC status. Because a DC is a server with a writable copy of Active Directory installed on it, this tool will install a copy of the directory database on the server, and configure the structure of Active Directory based on your input. After Active Directory is installed, you can then perform other tasks that will allow users of your network to access resources on the domain.

EXERCISE 1.01

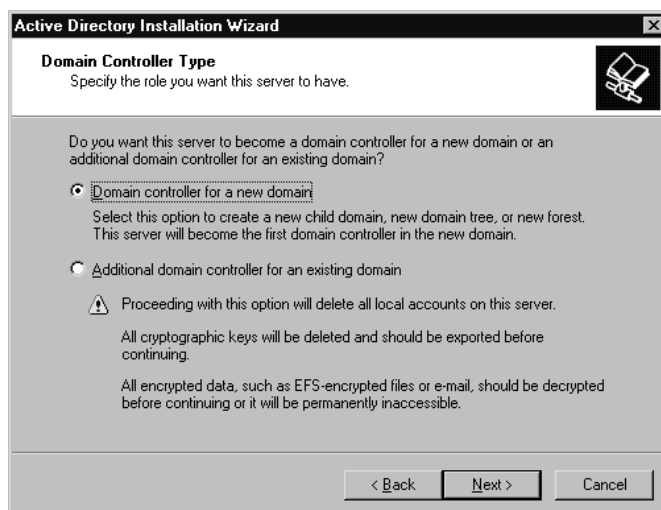
INSTALLING ACTIVE DIRECTORY

As with many of the exercises in this book, this exercise should not be performed on a production server. Moreover, while readers who have previous knowledge of Active Directory can perform this exercise, those who are new to Active Directory might want to read the next section to understand how Active Directory works before attempting to install it.

1. From the Run command on the Windows Start menu, type **DCPROMO** and then click **OK**.
2. A welcome screen will appear that identifies the program as the Active Directory Installation Wizard. Click **Next** to continue.
3. An information screen will appear, warning that clients running Windows 95 or Windows NT 4.0 SP3 and earlier won't be able to log on to Windows Server 2003 DCs or access domain resources. Click **Next** to continue.

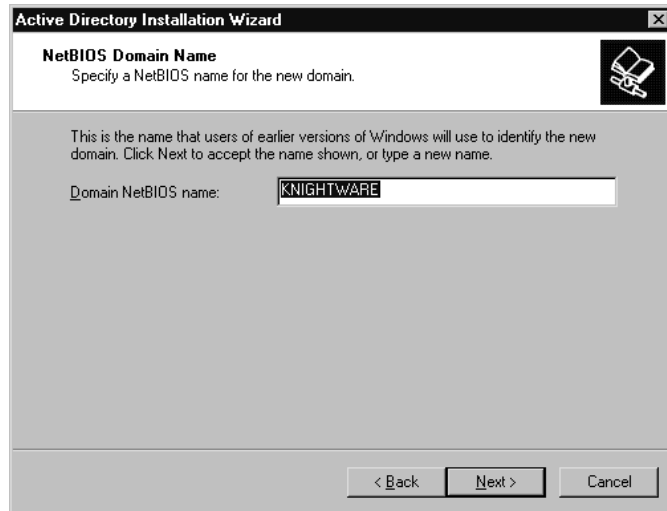
- The Domain Controller Type screen appears after this, allowing you to specify whether you want the server to be a DC for a new or existing domain (see Figure 1.3). Selecting the **Domain controller for a new domain** will allow you to create a new domain, while selecting **Additional domain controller for an existing domain** will add this server to a domain that already exists. Select the first of these options to create a new domain. Click **Next** to continue.

Figure 1.3 Domain Controller Type Screen of Active Directory Installation Wizard



- The next screen allows you to configure or install DNS on the server. If DNS is already running, then select **Yes, I will configure the DNS client**. If not, select **No, just install and configure DNS on this computer**. If you select **Yes** and DNS is not running, a warning screen will appear informing you of this. If DNS isn't running, select the second option (**No**), and click **Next** to continue.
- Enter the DNS name for the new domain (for example, syngress.com). Click **Next** to continue.
- As shown in Figure 1.4, the screen that appears next asks you to enter the NetBIOS name for this domain, which older versions of Windows will use to access the domain. Windows Server 2003 suggests a name based on your previously entered DNS name. Accept the default value, and click **Next** to continue.

Figure 1.4 NetBIOS Domain Name Screen of the Active Directory Installation Wizard



8. The next screen, shown in Figure 1.5, allows you to specify where the Active Directory database and log files will be stored. By default, this will be a directory called NTDS in the systemroot folder. Accept the default values and click **Next** to continue.

Figure 1.5 Database and Log Folders Screen of the Active Directory Installation Wizard

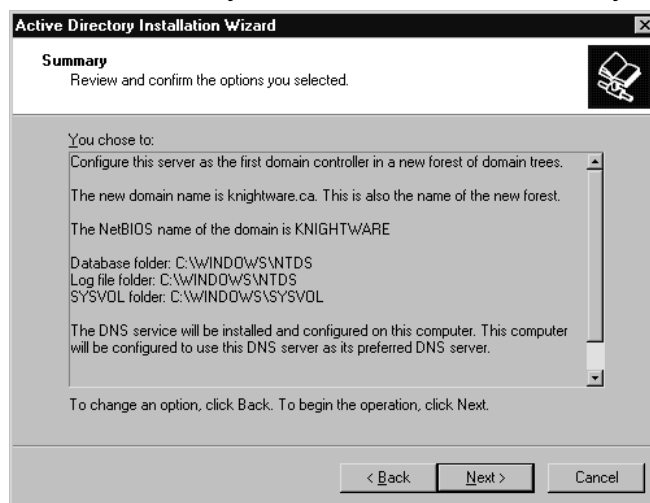


9. The next screen asks for the location of where public files that will be copied to other DCs will be stored. By default, this is stored in the

SYSVOL directory in the systemroot folder. Accept the default value and click **Next** to continue.


10. The next screen is used to set proper permissions based on whether you will be running server programs that were designed for pre-Windows 2000 domains. If this were the case, you would select the first option **Permissions compatible with pre-Windows 2000 Server operating systems**. Selecting this will allow anonymous users to read information on the domain, so it is best to select **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems** whenever possible. Assuming you will not be running such software, select the second option, and click **Next**.
11. The following screen asks that you enter a password used when the server is started in Directory Services Restore mode. This mode is used to restore Active Directory after it has become damaged. Enter a password in the first field, and then enter it in the field below to confirm your password. Click **Next** to continue.
12. The screen that appears next displays all the settings you chose for your installation of Active Directory (see Figure 1.6). Review the summary information that's shown on this screen, and then click **Next** to continue.

Figure 1.6 Summary Screen of the Active Directory Installation Wizard



13. The wizard will proceed to install and configure Active Directory based on your choices. Once this is done, click **Finish**.

14. To complete the installation, you will need to restart Windows Server 2003. A message box will appear informing you of this, and giving the options of restarting now or not. Click **Restart Now**.
-



Understanding How Active Directory Works

Active Directory provides the ability to manage your network through a single source of information. Using tools in Windows Server 2003, you can administer users, computers, printers, and a variety of other resources. Changes made to objects in the directory are replicated to other DCs. This ensures that each DC has an up-to-date copy of all directory objects and their attributes.



EXAM WARNING

Because of the logical, hierarchical structure of Active Directory, and the physical components required for it to work, most of the questions you see on the exam will indirectly or directly deal with the topics that follow. To understand Active Directory, it is important that you understand what these components are, so you can follow other issues discussed later in this book.

Directory Structure Overview

When you compare the directory structure of different organizations, you will find that they are different. Active Directory is organized in a hierarchical structure that is built from a variety of different components that represent elements of your network. For example, there are user objects, computer objects, and various containers to organize them. The way you organize these elements will make the hierarchical structure of Active Directory in your company different from other companies. The components that are part of this hierarchy (which we discuss in the sections that follow) include:

- Sites
- Domains
- Trees
- Forests
- Objects
- DCs

In addition to these, we will also look at the components of Active Directory that are used to organize and manage this hierarchy. These components are:

- GC
- Schema

Active Directory allows you to administrate your network by dealing with the physical and logical structure. The physical structure of your network consists of tangible elements that make up your network, while the logical structure is used to organize components into a hierarchy that matches the structure of your company. As we'll see in the sections that follow, sites represent the physical structure of a network, while domains, trees, and forests represent the logical structure.

Sites

A *site* is one or more IP subnets connected by a fast and reliable link. The term *subnet* is short for “subnetwork,” and refers to a group of neighboring computers that have been subdivided within the network. Computers in the subnet use a different network ID from those in other subnets, essentially becoming a smaller network within the network. Sites are used to store information about the topology of your network in Active Directory, so that the directory has information about the physical structure of the network.

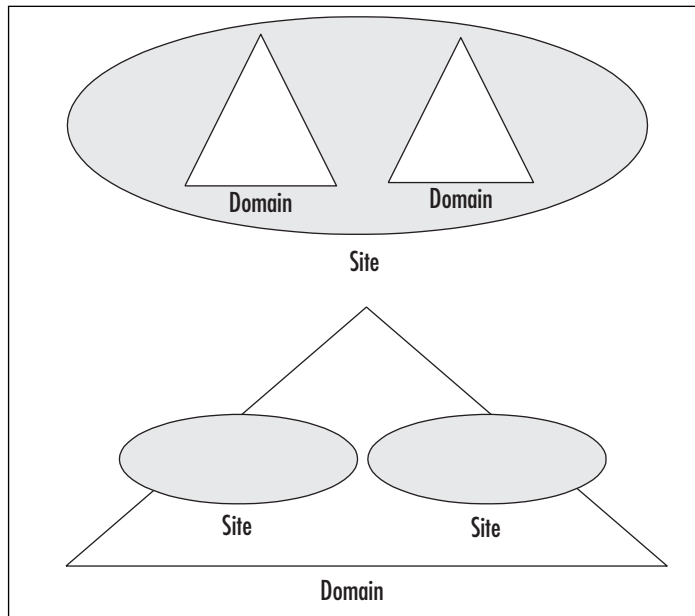
Active Directory uses information about the physical elements of a network in a number of ways. It allows Active Directory to determine the fastest connections between sites, so that updates in the directory can be replicated to other DCs. Sites contain computer and connection objects, which are used to configure replication between sites, allowing this information to be copied in the fastest, most effective way to DCs in other sites. It is also useful to users, as it will allow each user to be authenticated by the DC that's closest to that user.

Although not required, it is a good idea to have a DC in each site. When a client logs on to a domain, a DC must be contacted. The client will search the local site for a DC and then, if one is not found, attempt to connect to DCs in other sites. If the client has to connect to a DC in a different site, it might take a long time for the user to be authenticated. Creating different sites will group computers together, so they will authenticate to the DC that's closest to them.

An important feature of a site is that subnets are *well connected*. This means that the links between sites are reliable and fast. While determining what is fast can be subjective, Microsoft has traditionally defined a fast link as being at least 512 Kbps, while acknowledging that 128 Kbps or higher is sufficient. Because the bandwidth needed by an organization depends on the amount of data being transferred between sites, some companies will require a greater bandwidth to meet their needs.

As shown in Figure 1.7, there can be multiple domains in a site, or multiple sites in a domain. Because sites represent the physical structure, they are different from domains, trees, and forests (which we'll discuss next) that represent the logical structure. Sites are separate from these entities, and unfettered by issues that determine the logical structure of a Windows Server 2003 network.

Figure 1.7 Sites Can Contain Multiple Domains, and Domains Can Contain Multiple Sites



Domains

Domains have been a cornerstone of a Microsoft network since the days of Windows NT. A domain is a logical grouping of network elements, consisting of computers, users, printers, and other components that make up the network and allow people to perform their jobs. Because domains group these objects in a single unit, the domain acts as an administrative boundary, in which you can control security on users and computers. In Windows Server 2003, a domain also shares a common directory database, security policies, and (when other domains exist in the network) relationships with other domains. They are important logical components of a network, because everything is built upon or resides within the domain structure.

Sites and domains are different structures, and aren't bound by one another. Just as a site can include users and computers from multiple domains, domains can include multiple sites. This allows you to have objects from different areas of your network in the same domain, even if they're in different subnets or geographical locations.

In serving as an administrative boundary, each domain uses its own security policies. Group policies can be applied at a domain level, so that any users and computers within that domain are affected by it. This allows you to control access to resources, password policies, and other configurations to everyone within the domain. These security settings and policies only affect the domain, and won't be applied to other domains in the network. If large groups of users need different policies, you can either create multiple domains or apply settings in other ways (for example, using OUs, which we'll discuss later).

When a domain is created, a DNS domain name is assigned to identify it. DNS is used on the Internet and other TCP/IP networks for resolving IP addresses to user-friendly names. Because an Active Directory domain is integrated with DNS, this allows users, computers, applications, and other elements of the network to easily find DCs and other resources on the network.

As you can imagine, a significant number of objects can potentially exist within a domain. To allow for significant growth in a network, Microsoft designed Active Directory to support up to 10 million objects per domain. While Microsoft concedes this to be a theoretical estimate, the company provides a more practical estimate that each domain can support at least 1 million objects. In either case, chances are your domain will never reach either of these limits. If it does, you'll need to create additional domains, and split users, computers, groups, and other objects between them.

Earlier in this chapter, we mentioned that updates to the directory are replicated to other DCs, so that each has an identical copy of the directory database. We'll explain replication in greater detail later in this chapter, but for now it is important to realize that Active Directory information is replicated to every DC within a domain. Each domain uses its own directory database. Because the information isn't replicated to other domains, this makes the domain a boundary for replication as well as for administration and security.

Domain Trees

Although domains serve as boundaries for administration and replication, this does not mean that you should only use one domain until you reach the limit on the number of objects supported per domain. That depends on your organizational structure. You might want to use multiple domains for any of the following reasons:

- To decentralize administration
- To improve performance
- To control replication
- To use different security settings and policies for each domain
- If you have an large number of objects in the directory

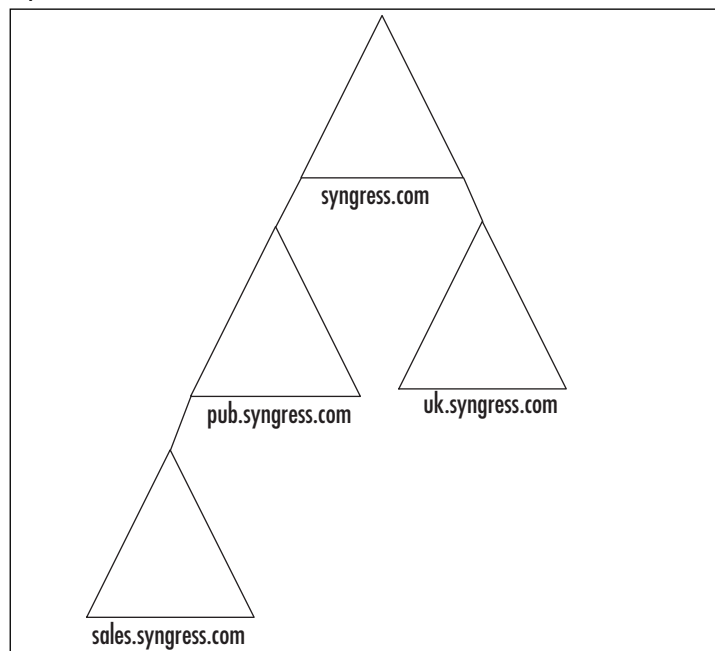
For example, your company might have branch offices in several countries. If there is only one domain, directory information will have to be replicated between DCs in each country, or (if no DCs resides in those locations) users will need to log on to a DC in

another country. Rather than replicating directory information across a WAN, and having to manage disparate parts of the network, you could break the network into several domains. For example, you might create one domain for each country.

Creating separate domains does not mean there will be no relationship between these different parts of your network. Active Directory allows multiple domains to be connected together in a hierarchy. As shown in Figure 1.8, a domain can be created beneath an existing domain in the hierarchy. The pre-existing domain is referred to as a “parent domain,” and the new domain created under it is referred to as a “child domain.” When this is done, the domains share a common namespace. They also share a schema, configuration, and GC, as do all domains in the same forest, whether or not they have a parent-child relationship (we’ll discuss these elements in greater detail later in this chapter).

As seen in Figure 1.8, domains created in this parent-child structure and sharing a namespace belong to a *domain tree*. Trees follow a DNS naming scheme, so that the relationship between the parent and child domains is obvious and easy to follow. To conform to this naming scheme, a child domain appends its name to the parent’s name. For example, if a parent domain used the domain name syngress.com, a child domain located in the United Kingdom might have the name uk.syngress.com. Names can also indicate the function of a domain, rather than its geographical location. For example, the child domain used by developers might use the name dev.syngress.com. Because domain trees use a contiguous namespace, it is easy to see which domains are child domains of a particular parent domain.

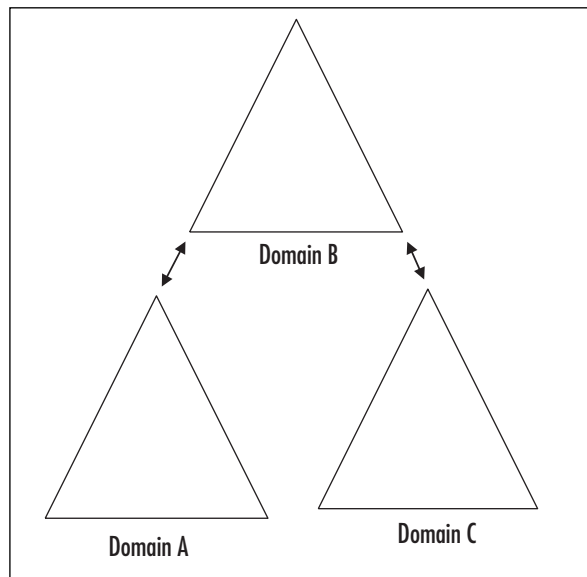
Figure 1.8 A Domain Tree Consists of Parent and Child Domains in a Contiguous Namespace



When a child domain is created, a two-way transitive trust relationship between the parent and child domains is automatically created. A trust relationship allows pass-through authentication, so users who are authenticated in a trusted domain can use resources in a trusting domain. Because the trust between a parent and child domain is bidirectional, both domains trust one another, so users in either domain can access resources in the other (assuming, of course, that the users have the proper permissions for those resources).

The other feature of the trust relationship between parent and child domains is that they are transitive. A transitive relationship means that pass-through authentication is transferred across all domains that trust one another. For example, in Figure 1.9, Domain A has a two-way transitive trust with Domain B, so both trust one another. Domain B has a two-way transitive trust with Domain C, so they also trust one another, but there is no trust relationship between Domain A and Domain C. With the two-way transitive trust, Domain C will trust Domain A (and vice versa) because both trust Domain B. This will allow users in each of the domains to access resources from the other domains. Trusts can also be manually set up between domains so that they are one-way and nontransitive, but by default, transitive bidirectional trusts are used in domain trees and forests. These trusts are also *implicit*, meaning that they exist automatically by default when you create the domains, unlike *explicit* trusts that must be created manually.

Figure 1.9 Adjoining Domains in a Domain Tree Use Two-Way Transitive Trusts

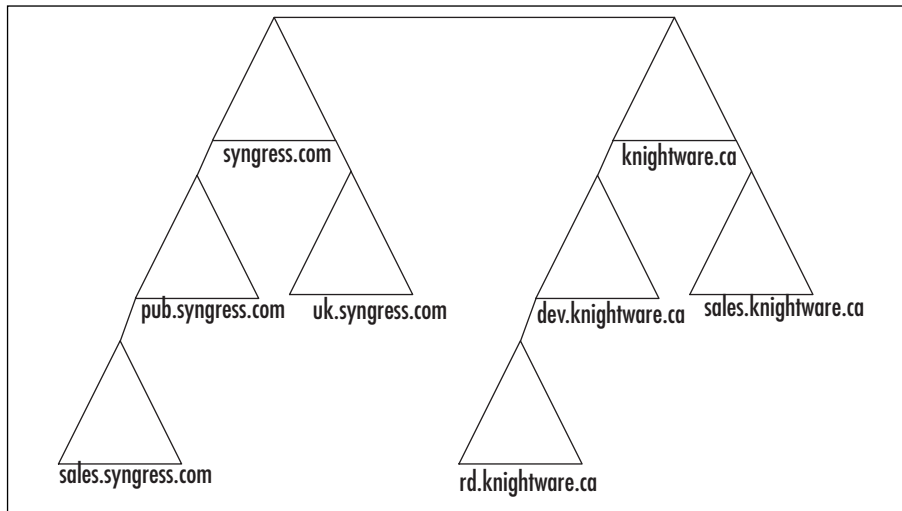


Forests

Just as domains can be interconnected into trees, trees can be interconnected into forests. A forest is one or more domain trees that share the same schema, GC, and configuration

information. As is the case with domain trees, domains in the same forests use two-way transitive trusts between the roots of all domain trees in the forest (that is, the top level domain in each tree) to allow pass-through authentication, so users can access resources in domains throughout the forest. As shown in Figure 1.10, although trees require a contiguous namespace, a forest can be made up of multiple trees that use different naming schemes. This allows your domains to share resources across the network, even though they don't share a contiguous namespace.

Figure 1.10 A Forest Allows Multiple Domain Trees to Be Connected and Share Information



Every Active Directory structure has a forest, even if it only consists of a single domain. When the first Windows Server 2003 DC is installed on a network, you create the first domain that's also called the *forest root domain*. Additional domains can then be created that are part of this forest, or multiple forests can be created. This allows you to control which trees are connected and can share resources with one another (within the same forest), and which are separated so that users can't search other domains sharing the GC (in separate forests).

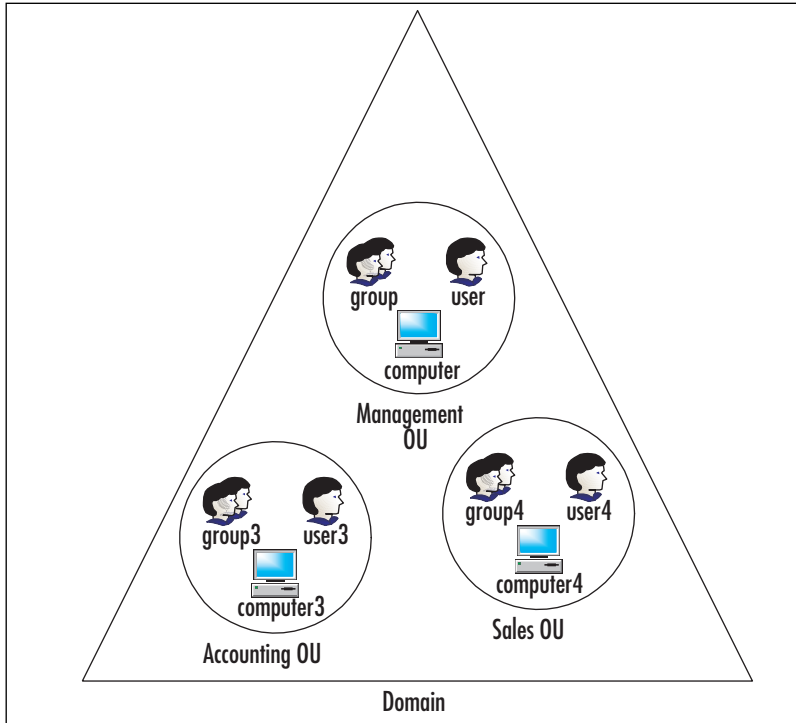
Organizational Units

When looking at domain trees, you might think that the only way to create a directory structure that mirrors the organization of your company is to create multiple domains. However, in many companies, a single domain is all that's needed. To organize Active Directory objects within this single domain, OUs can be used.

As we mentioned earlier, OUs are containers that allow you to store users, computers, groups, and other OUs. By placing objects in different OUs, you can design the layout of Active Directory to take the same shape as your company's logical structure, without creating separate domains. As shown in Figure 1.11, you can create OUs for different areas of

your business, such as departments, functions, or locations. The users, computers, and groups relating to each area can then be stored inside the OU, so that you can find and manage them as a single unit.

Figure 1.11 Organizational Units Can Contain Other Active Directory Objects



OUs are the smallest Active Directory unit to which you can delegate administrative authority. When you delegate authority, you give specific users or groups the ability to manage the users and resources in an OU. For example, you can give the manager of a department the ability to administer users within that department, thereby alleviating the need for you (the network administrator) to do it.

Active Directory Components

When looking at the functions of domains, trees, forests, and OUs, it becomes apparent that each serves as a container. These container objects provide a way to store other components of Active Directory, so that they can be managed as a unit and organized in a way that makes administration easier. OUs also provide the added feature of allowing nesting, so that you can have one OU inside another.

The bulk of components in Active Directory, however, are objects that represent individual elements of the network (in Novell's NDS structure, these are called *leaf objects*, in

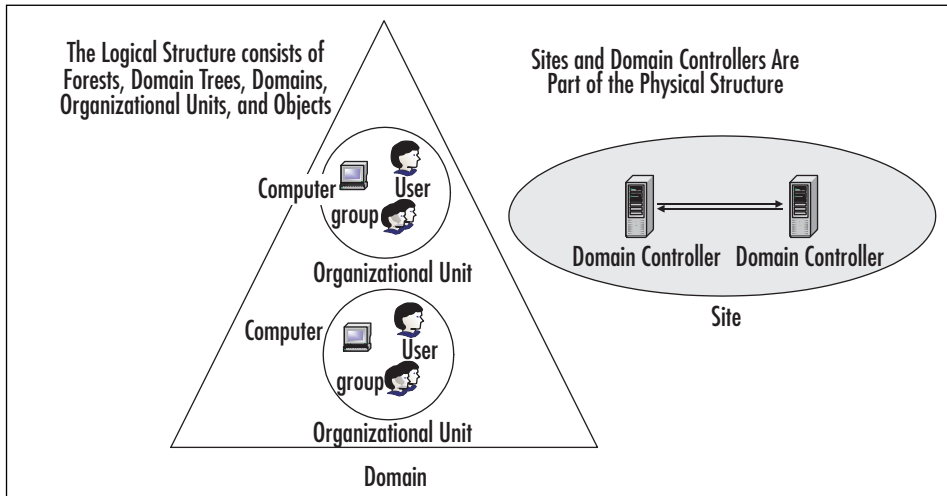
keeping with the tree analogy, because they are at the end of the hierarchical “branch” and don’t contain any other objects). Objects are divided into *classes*, and each object class includes a set of *attributes*, which are properties that hold data on characteristics and configurations. Just as people are defined by their characteristics (for example, eye and hair color, height, weight), attributes define an object. A printer object might have attributes that include the make, model, and configuration information related to that device, whereas a user object would include attributes such as username, password, and other data that defines the user. As we’ll see in Chapter 2, these various objects populate the directory, and are used to manage such things as user, computer, and group accounts.

Logical vs. Physical Components

The components making up Active Directory can be broken down into logical and physical structures. Logical components in Active Directory allow you to organize resources so that their layout in the directory reflects the logical structure of your company. Physical components in Active Directory are similarly used, but are used to reflect the physical structure of the network. By separating the logical and physical components of a network, users are better able to find resources, and administrators can more effectively manage them.

Many directories are designed to follow the logical structure of an organization. You’re probably familiar with organizational charts; maps that show the various departments in a company, and illustrate which departments are accountable to others. In such a map, a Payroll department might appear below the Finance department, even though they are physically in the same office. Just as the chart allows you to find where a department falls in the command structure of a company, the logical structure of a directory allows you to find resources based on a similar logical layout. As we saw earlier, you can organize your network into forests, trees, and domains, and then further organize users and computers into OUs named after areas of your business. A map of the directory structure can be organized to appear identical to the logical structure of the company.

Physical components are used to design a directory structure that reflects the physical layout, or *topology*, of the network. For example, as we saw earlier, a site is a combination of subnets, and a DC is a server that has a copy of the directory on it. DCs are physically located at specific locations in an organization, while subnets consist of computers using the same grouping of IP addresses. In both cases, you could visit a room or building and find these components. Thus, physical components can be used to mirror the physical structure of an organization in the directory. As illustrated in Figure 1.12, this makes the physical structure considerably different from the logical structure of a network.

Figure 1.12 Logical Structure vs. Physical Structure

Domain Controllers

DCs are used to manage domains. As mentioned, the directory on a DC can be modified, allowing network administrators to make changes to user and computer accounts, domain structure, site topology, and control access. When changes are made to these components of the directory, they are then copied to other DCs on the network.

Because a DC is a server that stores a writable copy of Active Directory, not every computer on your network can act as a DC. Windows Server 2003 Active Directory can only be installed on Microsoft Windows Server 2003, Standard Edition; Windows Server 2003, Enterprise Edition; and Windows Server 2003, Datacenter Edition. Servers running other than the Web Edition of Windows Server 2003 cannot be DCs, although they can be member servers that provide resources and services to the network.

When a DC is installed on the network, the first domain, forest, and site are created automatically. Additional domains, forests, and sites can be created as needed, just as additional DCs can be added. This allows you to design your network in a way that reflects the structure and needs of your organization.

While only one DC is required to create a domain, multiple DCs can (and usually should) be implemented for fault tolerance and high availability. If more than one DC is used and one fails, users will be able to log on to another DC that is available. This will allow users to continue working while the DC is down. In larger companies, a number of DCs can be added to accommodate significant numbers of users who might log on and log off at the same time of day or need to access resources from these servers.



EXAM WARNING

Windows Server 2003 computers can be promoted to DCs by installing Active Directory on them. To install Active Directory, the Active Directory Installation Wizard (invoked by running DCPROMO.EXE) is used. Information provided during the installation is used to add the server to an existing domain, or to create a new domain, forest, and site if the DC is the first one installed on a network.

Master Roles

Certain changes in Active Directory are only replicated to specific DCs on the network. Operations Masters are DCs that have special roles, keeping a master copy of certain data in Active Directory and copying data to other DCs for backup purposes. Because only one machine in a domain or forest can contain the master copy of this data, they are also referred to as Flexible Single Master Operations (FSMO) roles.

Five different types of master roles are used in an Active Directory forest, each providing a specific purpose. Two of these master roles are applied to a single DC in a forest (forestwide roles), while three others must be applied to a DC in each domain (domainwide roles). In the paragraphs that follow, we will look at each of these roles, and discuss how they are significant to Active Directory's functionality.

Forestwide master roles are unique to one DC in every forest. There are two master roles of this type:

- Schema Master
- Domain Naming Master

The *Schema Master* is a DC that is in charge of all changes to the Active Directory schema. As we'll see in the next section, the schema is used to define what object classes and attributes are used within the forest. The Schema Master is used to write to the directory's schema, which is then replicated to other DCs in the forest. Updates to the schema can be performed only on the DC acting in this role.

The *Domain Naming Master* is a DC that is in charge of adding new domains and removing unneeded ones from the forest. It is responsible for any changes to the domain namespace. Such changes can only be performed on the Domain Naming Master, thus preventing conflicts that could occur if changes were performed on multiple machines.

In addition to forestwide master roles, there are also domainwide master roles. There are three master roles of this type:

- Relative ID (RID) Master
- Primary domain controller (PDC) Emulator
- Infrastructure Master



TEST DAY TIP

Remember that there is only one forestwide master for each role in a forest, and one domainwide master for each role in a domain. There can only be one Schema Master and Domain Naming Master per forest. In other words, if there were two forests, then there would be one Schema Master and one Domain Naming Master in each forest. In the same way, there can only be one RID Master, PDC Emulator, and Infrastructure Master per domain. Although multiple domains can exist in a forest, there can only be one RID Master, PDC Emulator, and Infrastructure Master in each domain

The *RID Master* is responsible for creating a unique identifying number for every object in a domain. These numbers are issued to other DCs in the domain. When an object is created, a sequence of numbers that uniquely identifies the object is applied to it. This number consists of two parts: a domain security ID (SID) and a RID. The domain SID is the same for all objects in that domain, while the RID is unique to each object. Instead of using the name of a user, computer, or group, this SID is used by Windows to identify and reference the objects. To avoid potential conflicts of DCs issuing the same number to an object, only one RID Master exists in a domain, to control the allocation of ID numbers to each DC, which the DC can then hand out to objects when they are created.

The *PDC Emulator* is designed to act like a Windows NT primary DC. This is needed if there are computers running pre-Windows 2000 and XP operating systems, or if Windows NT backup domain controllers (BDCs) still exist on the network. The PDC Emulator is responsible for processing password changes, and replicating these changes to BDCs on the network. It also synchronizes the time on all DCs in a domain so servers don't have time discrepancies between them. Because there can only be one Windows NT PDC in a domain, there can be only one PDC Emulator.

Even if there aren't any servers running as BDCs on the network, the PDC Emulator still has a purpose in each domain. The PDC Emulator receives preferred replication of all password changes performed by other DCs within the domain. When a password is changed on a DC, it is sent to the PDC Emulator. The PDC Emulator is responsible for this because it can take time to replicate password changes to all DCs in a domain. If a user changes his or her password on one DC and then attempts to log on to another, the second DC he or she is logging on to might still have old password information. Because this DC considers it a bad password, it forwards the authentication request to the PDC Emulator to determine whether the password is actually valid. Whenever a logon authentication fails, a DC will always forward it to the PDC Emulator before rejecting it.

The *Infrastructure Master* is in charge of updating changes made to group memberships. When a user moves to a different domain and his or her group membership changes, it can take time for these changes to be reflected in the group. To remedy this, the infrastructure manager is used to update such changes in its domain. The DC in the Infrastructure Master role compares its data to the GC, which is a subset of directory information for all domains

in the forest. When changes occur to group membership, it then updates its group-to-user references and replicates these changes to other DCs in the domain.

Schema

The *schema* is a database that is used to define objects and their attributes. Information in the schema is used to control the types of objects (classes) that can be created in Active Directory, and the additional properties (attributes) associated with each. In other words, the schema determines what you can create in Active Directory, and the data that can be used to configure these objects.

The schema is made up of classes and attributes. Object classes define the type of object, and include a collection of attributes, which are used to describe the object. For example, the User class of object contains attributes made up of information about the user's home directory, first name, last name, address, and so on. While the object class determines the type of object that can be created in Active Directory, the attributes are used to provide information about it. An object's attributes are also known as its *properties*, and in most cases, you can configure its attributes by editing its properties sheet (usually accessed by right clicking the object and selecting **Properties**).

Active Directory comes with a wide variety of object classes, but additional ones can be created if needed. Because the schema is so important to Active Directory's structure, *extensions* (additions and modifications) to the schema can only be made on one DC in the forest. Modifications to the schema can only be made on the DC that's acting in the *Schema Master* role. Schema information is stored in a directory partition of Active Directory, and is replicated to all DCs in a forest.

Attributes are created using the Active Directory Schema snap-in for the Microsoft Management Console (MMC) (which we'll discuss later in this chapter). When a new class or attribute is added to the schema, it cannot be deleted. If a class or attribute is no longer needed, it can only be deactivated, so it cannot be used anymore. Should the class or attribute be needed later, you can then reactivate it.

Global Catalog

As anyone who's tried to search a large database can attest, the more data that's stored in a database, the longer it will take to search. To improve the performance of searching for objects in a domain or forest, the GC is used. The GC server is a DC that stores a copy of all objects in its host domain, and a partial copy of objects in other domains throughout the forest. The partial copy contains objects that are most commonly searched for. Because the GC contains a subset of information in Active Directory, less information needs to be replicated, and increases performance when users search for specific attributes of an object.

In addition to being used for searches, the GC is also used to resolve UPNs that are used in authentication. As discussed earlier, the UPN has a format like an e-mail address. If a user logs on to a DC in a domain that doesn't contain the account, the DC will use the GC to resolve the name and complete the logon process. For example, if a user logged on

with the UPN `myname@us.syngress.com` from a computer located in `ca.syngress.com`, the DC in `ca.syngress.com` would be unable to find the account in that domain. It would then use the GC to find and authenticate the user's account.

The GC is also used to store information on Universal Group memberships, in which users from any domain can be added and allowed access to any domain. We'll discuss groups in greater detail in Chapter 2. When a user who is a member of such a group logs on to a domain, the DC will retrieve his or her Universal Group membership from the GC. This is only done if there is more than one domain in a forest.

The GC is available on DCs that are configured to be GC servers. Creating a GC server is done by using the Active Directory Sites and Services snap-in for the MMC (which we'll discuss later in this chapter). After a GC server is configured, other DCs can query the GC on this server.

Replication Service

The Windows Server 2003 replication service is used to replicate Active Directory between DCs, so that each DC has an up-to-date copy of the directory database. Because each DC has an identical copy of the directory, they can operate independently, allowing users to be authenticated and use network resources if one of the DCs fails. This allows Windows Server 2003 DCs to be highly reliable and fault tolerant.

Multimaster replication is used to copy changes in the directory to other DCs. With multimaster replication, DCs work as peers to one another, so that any DC accepts and replicates these updates (with the exception of the special types of data for which an Operations Master is assigned). Rather than having to make changes on a primary DC, changes can be made to the directory from any DC.

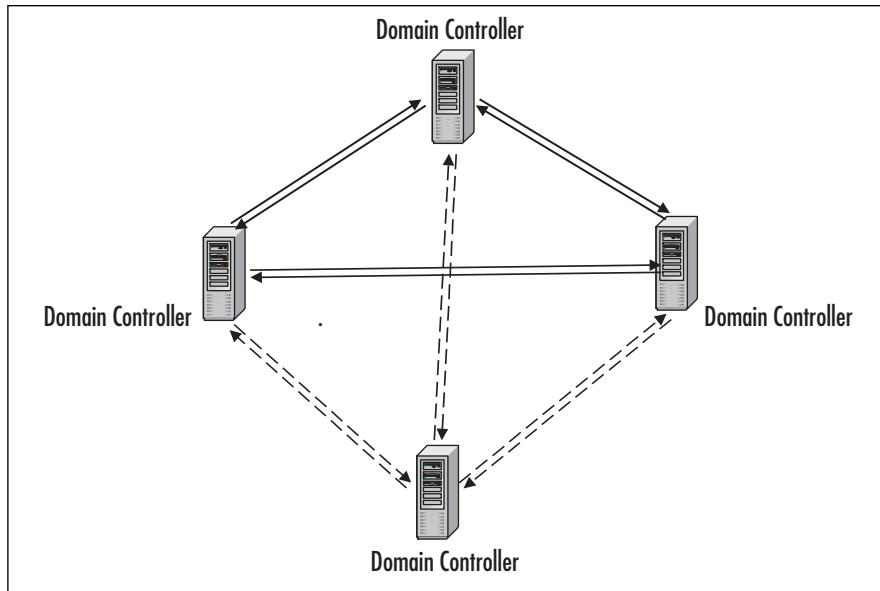
Replication occurs automatically between DCs, and generally, no additional configuration is required. However, because there are times when network traffic will be higher, such as when employees log on to DCs at the beginning of the workday, replication can be configured to occur at specific times. This will enable you to control replication traffic so it doesn't occur during peak hours.

To replicate the directory effectively, Windows Server 2003 uses the Knowledge Consistency Checker (KCC) to generate a replication topology of the forest. A *replication topology* refers to the physical connections used by DCs to replicate the directory to other DCs within the site and to DCs in other sites. After initially creating a replication topology, the KCC will review and modify the topology at regular intervals. This allows it to see if certain connections or DCs are unavailable, and if changes need to be made as to how replicated data will be transferred to other DCs.

Replication is handled differently within a site as opposed to when the directory is replicated to other sites. *Intra-site replication* (in which Active Directory is replicated within a site) is handled by using a ring structure. The KCC builds a bidirectional ring, in which replication data is passed between DCs in two directions. Because the data is only being transferred within the site, the replicated data isn't compressed.

The KCC creates at least two connections to each DC, so if one connection fails, the other can be used. For example, in Figure 1.13, connections that are functional are shown with a straight line, while broken connections are shown with dotted lines. Because one of the four servers in Figure 1.13 has failed, replication data cannot be passed through it, so another connection between the servers is used. Using multiple connections provides fault tolerance.

Figure 1.13 Replication Topology



Intra-site replication is automated to occur at regular intervals, and only occurs when DCs are notified of a change. By default, when a change is made on a DC, it will wait 15 seconds and then send notification to its closest replication partner. If it has more than one replication partner, it will send out notifications in three-second intervals to each additional partner. When a partner receives this notification, it will send out a request for updated directory information to the original DC, which then responds by sending the updated data. The exception to this process is when an account is locked out, the DC account is changed, or there are changes in account lockout policy or domain password policy. In these circumstances, there is no 15-second waiting period, and replication occurs immediately.


Replication between sites is called *inter-site replication*. Because the bandwidth between sites might be slower than that within a site, inter-site replication occurs less frequently and is handled differently. Rather than informing other DCs shortly after a change occurs, replication occurs at scheduled times. Information about *site link objects* is used to determine the best link to use for passing this data between sites.

Site links are used to define how sites replicate Active Directory information between one another. These objects store data controlling which sites are to replicate traffic between

one another, and which should be used over others. For example, you might have an ISDN connection between your offices and one located overseas. If the overseas link were slower and more costly to use than others, you could configure the link so it is only used as a last resort. Through the site link object, the fastest and least expensive connection between sites is used for replication.

A DC acts in the role of an inter-site topology generator in each site, and serves the purpose of building this topology. It considers the cost of different connections, whether DCs are available, and whether DCs have been added to sites. By gathering this information, the KCC can then update the topology as needed, and provide the method of passing data between the sites.

How often replication occurs is configurable, so that it occurs as frequently or infrequently as your needs dictate. By default, inter-site replication occurs every 180 minutes (three hours), and will use the site link to meet this schedule 24 hours a day, 7 days a week. The frequency of replications can be modified as needed to occur at certain times and days of the week.



Using Active Directory Administrative Tools

Just as organizations have the tendency to grow and change, so do the networks they use. In a Windows Server 2003 network, the number of domains, sites, OUs, users, computers, and other objects populating Active Directory can grow exponentially with a business. Every new employee needs a new account, and every new computer added to the network means another object added to the directory. Even when growth is limited, there can be a considerable amount of maintenance to these objects, such as when users change jobs, addresses, or other issues that involve changes to information and access. To aid administrators with these tasks, Active Directory provides a number of tools that make management easier.

Two types of administrative tools can be used to manage Active Directory. Windows Server 2003 provides a variety of new command-line tools that individually administer different aspects of the directory and its objects. By clicking on the Windows **Start** menu and clicking **Programs | Accessories | Command Prompt** (or simply clicking **Start | Run** and typing **cmd**), a prompt will appear allowing you to enter these commands and control objects and elements of Active Directory. The other method of managing Active Directory is with tools using a graphical user interface (GUI). These tools allow you to point and click through objects, and modify them using a graphical display. Most of the graphical tools are available through the **Start | Programs | Administrative Tools** menu.

Command-Line Tools

Windows Server 2003 includes a number of tools that weren't available for administration of Active Directory in Windows 2000. These tools run from the command prompt, and can perform tasks that could previously only be performed using the GUI consoles. Using these tools, you can connect to remote servers and make modifications to the directory without the added overhead of a GUI interface. Because many tasks can be performed through the MMC (which we'll discuss next) or text-based commands, this provides a greater freedom of choice for administrators on how these tasks are performed.

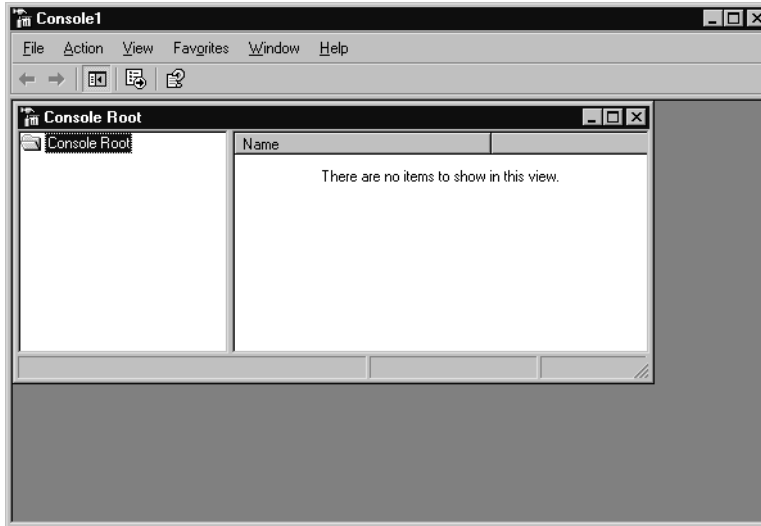
Command-line tools are particularly beneficial for administrators who were experienced with DOS or other network operating systems that used command-line tools (such as Novell NetWare and UNIX). While using these tools requires the user to switch from one tool to another and manually type each command, users might find that they can perform common tasks faster through this interface rather than with the GUI.

Command-line tools are also useful for administrators who want to schedule tasks at a certain time, or automate tasks in other ways. By running the commands through scripts or batch files, or invoking them through other applications developed in-house, you can automate certain tasks and make administration easier. Allowing these commands to be invoked in these ways can also be useful for allowing users to perform management tasks with which they aren't familiar.

Graphical Administrative Tools/MMCs

A primary administrative tool for managing Windows Server 2003 and Active Directory is the Microsoft Management Console (MMC). The MMC isn't a management tool in itself, but an interface that's used to load snap-ins that provide administrative functionality. Snap-ins provide a specific functionality, or a related set of functions. Because of the design of the MMC interface, you can load several snap-ins into one console, and create custom tools to deal with specific tasks. In addition, because these snap-ins run in the same environment, it becomes easier to learn how to use these tools because you don't have to learn a different interface for each.

MMCs can be started by opening pre-made consoles that are available under the **Administrative Tools** folder in the Windows **Start** menu. An empty MMC can be started by using the **Run** command in the Windows **Start** menu. By typing **MMC** in the **Run** command in the Windows **Start** menu, an empty MMC will start as shown in Figure 1.14.

Figure 1.14 Microsoft Management Console

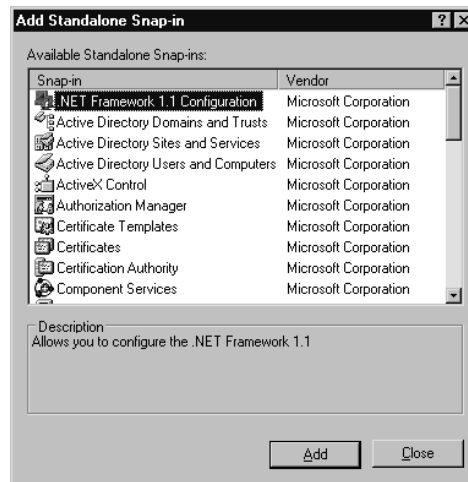
The windows appearing in the MMC are interfaces with individual snap-ins or custom console files. Each child window in the MMC has two panes. The left pane displays the console tree, which is a hierarchical display of tools available through the console. These can be multiple snap-ins that have been loaded into the MMC and saved as a custom console. The right pane is called the detail pane, and provides commands and information relating to what is selected in the console tree.

You can add snap-ins for specific tasks by clicking on the **File** menu and selecting **Add/Remove Snap-in**. When this is done, a new dialog box will appear with two tabs: Standalone and Extensions.

The Standalone tab is used for standalone snap-ins, which are designed to run without any additional requirements. The Extensions tab is used to load a special type of snap-in, called an extension snap-in. These are used to add additional functions to a standalone snap-in that's already been installed.

The Standalone tab is used to add or remove snap-ins from the console. As shown in Figure 1.15, clicking the **Add** button on this tab will display a list of available standalone snap-ins. After selecting the one you want to add, click the **Add** button on this dialog. Clicking **Close** will exit this screen, and return you to the previous one, which will now include your selected snap-ins in a list of ones to install in this console. Clicking **OK** confirms the selection, and installs them.

Figure 1.15 Add/Remove Snap-in Dialog Box



As you can see by in Figure 1.15, there are three snap-ins available for Active Directory:

- Active Directory Users and Computers
- Active Directory Domains and Trusts
- Active Directory Sites and Services

While we'll discuss each of these in the sections that follow, it is important to realize that these aren't the only snap-ins that you can use with Windows Server 2003. The MMC supplies these three snap-ins for use with Active Directory, but others are also available for specific purposes and management tasks. Each has an individual functionality or set of related functions for administering Windows Server 2003 and Active Directory.



NOTE

Note that although the three Active Directory-related snap-ins are available to be added to a custom MMC, each is already installed in a separate pre-configured MMC available through the Administrative Tools menu.

Because multiple snap-ins can be added and configured in the MMC, you can create custom consoles to perform specific tasks. After setting up a console, you can save it to a file that has the .msc extension. The console can be saved in one of two modes: Authoring and User.

Authoring mode is used to provide full access to the functions of an MMC console. When saved in this mode, users who open the console can add and remove snap-ins, create new Windows, create Favorites and taskpads, view everything in the console tree, and save consoles.

User mode is used to limit another user's ability to use certain functions of the console. If you were creating a console for users to perform a specific task, but didn't want them to access other functions, then User mode would be ideal. There are three access levels for User mode:

- **Full Access** The same as Author mode, except that snap-ins can't be added or removed, console settings can't be changed, and users can't create Favorites and taskpads.
- **Limited Access, Multiple Windows** Allows users to view parts of the console tree that were visible when the console was saved, and prohibits users from closing existing windows. Users can, however, create new windows.
- **Limited Access, Single Window** Also allows users to access parts of the console tree that were visible when the console was saved, but prohibits users from creating new windows.



TEST DAY TIP

The exam will test you on your knowledge of Active Directory, and how to use the tools that come with Windows Server 2003. Remember that the MMC allows you to load snap-ins to perform specific tasks, and can also be used to create custom consoles that provide limited functionality.

Active Directory Users and Computers

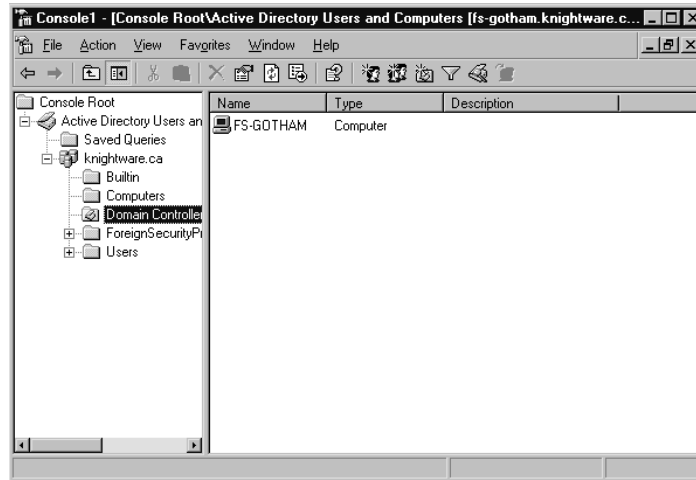
The Active Directory Users and Computers console is one of the MMC snap-ins for use with Active Directory. It allows you to administer user and computer accounts, groups, printers, OUs, contacts, and other objects stored in Active Directory. Using this tool, you can create, delete, modify, move, organize, and set permissions on these objects.

As shown in Figure 1.16, when this tool is loaded, a node will appear in the console tree (left pane) showing the domain. Expanding this node will show a number of containers that are created by default. While additional containers can be created, the ones that appear here after creating a DC are:

- Builtin
- Computers
- Domain Controllers
- Users

These containers store objects that can be managed with this tool, and allow you to view and modify information related to these different objects.

Figure 1.16 Active Directory Users and Computers



The *Builtin* container holds groups that were created by Windows Server 2003, and can be used to control access. You can add users to these Builtin groups to give them the ability to perform certain tasks. For example, rather than allowing everyone in the IT department to use the same Administrator account, users can be added to the built-in Administrators group. This gives them the ability to administer Windows Server 2003, but allows you to track which person with this level of security performed certain tasks.

The *Computers* container is used to store computer objects. These are (as the name implies) computers running on the network that have joined the domain and have accounts created in Active Directory. The Computers container can also include accounts used by applications to access Active Directory.

The *Domain Controllors* container contains objects representing DCs that reside in the domain. The ones shown in this container are ones running Windows 2000 Server and Windows Server 2003. Earlier versions are not displayed.

The *Users* container is used to store user accounts and groups. Users and Groups that appear in this container are ones that were created using application programming interfaces (APIs) that can use Active Directory, and ones that were created in Windows NT prior to upgrading.

Additional containers can be displayed when Active Directory Users and Computers is running with Advanced Features activated. You can enable Advanced Features by clicking on the menu item with this name, found in the View menu. When Advanced Features have been activated, *LostAndFound* and *System* containers are displayed in the left console tree.

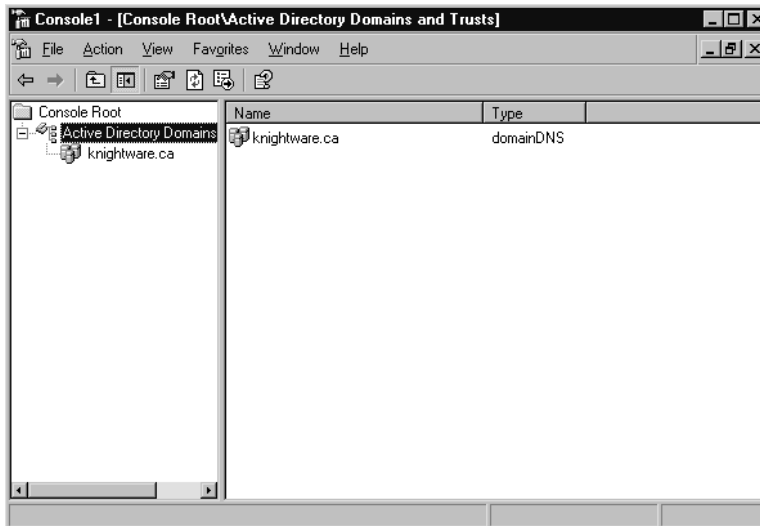
The *LostAndFound* container is used to store stray objects whose containers no longer exist. If an object is created at the same time its container is deleted, or if it is moved to a location that's missing after replication, the object is placed in this container. This allows you to manage the lost object, and move it to a container that does exist.

The *System* container is used for system settings. These are built-in settings for containers and objects used by Active Directory and Windows Server 2003.

Active Directory Domains and Trusts

The Active Directory Domains and Trusts console is used to manage domains and the trust relationships between them. As shown in Figure 1.17, the console tree of this tool includes a node for domains making up the network. By selecting the **Active Directory Domains and Trusts** node, a listing of domains will appear in the right pane. Using this tool, you can create, modify, and delete trust relationships between domains, set the suffix for UPNs, and raise domain and forest functional levels. This enables administrators to control how domains function, and how they interoperate.

Figure 1.17 Active Directory Domains and Trusts



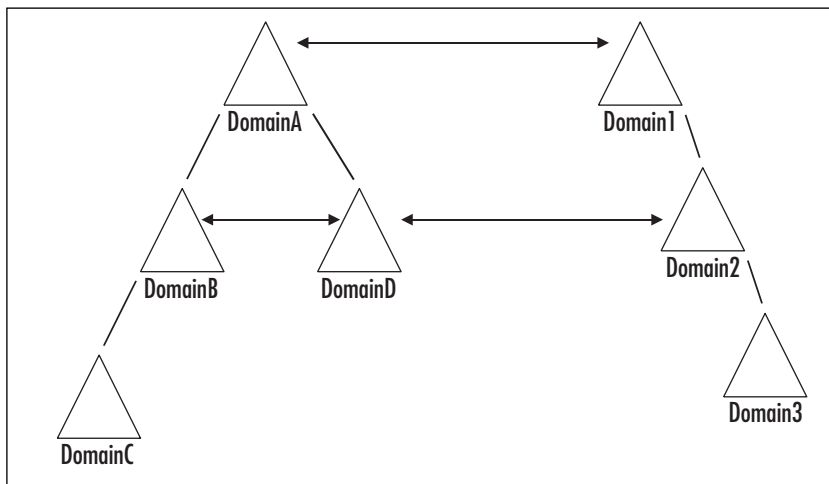
Using the Active Directory Domains and Trusts console, you can create a variety of different types of trusts between domains and forests. Earlier, we discussed how parent and child domains and domain trees use a two-way transitive trust to share resources between domains. The two-way transitive trust means that both domains trust one another, as well as any other domains with which they have similar trust relationships. In addition to this type of trust, additional trusts can be created:

- Shortcut trust
- Forest trust
- Realm trust
- External trust

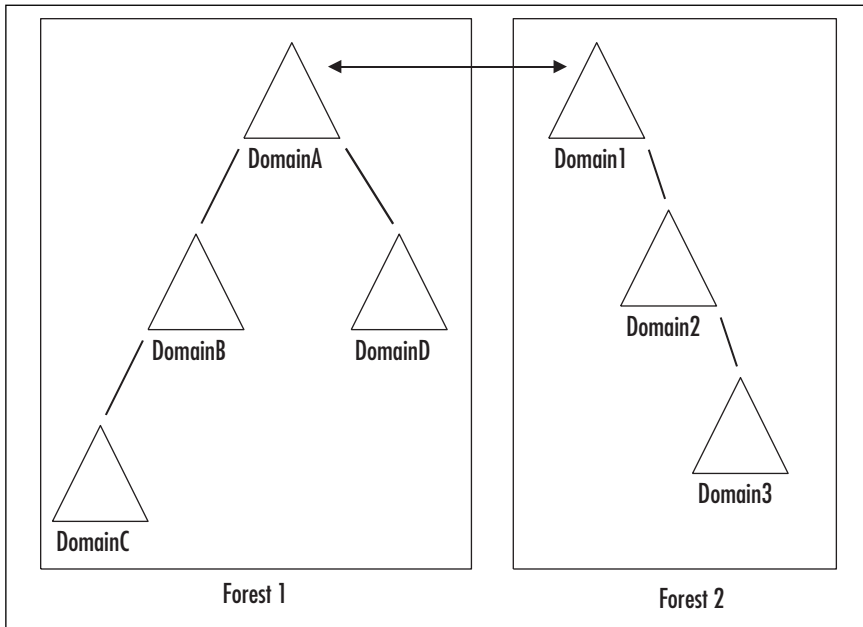
A *shortcut trust* is transitive, and can be either one-way or two-way. This means that either one domain can trust another but not vice versa, or both domains can trust each other. This type of trust is used to connect two domains in a forest, and is particularly useful when the domains are in different trees. By creating a shortcut, one domain can connect with another quickly, improving logon times between domains. Connection is quicker because, when two domains in different trees connect via the implicit trusts that exist by default, the trust path must go all the way up the tree to the root domain, across to the other tree's root domain, and back down the second tree. A shortcut trust, as its name indicates, creates a direct trust between the two domains in different trees.

To illustrate this, let's look at the situation in Figure 1.18. If a user in DomainD wanted to use resources in Domain2, he or she would be authenticating to a domain that is located in a different tree. Without a shortcut trust, the connection would go through DomainA, across the trust between the two trees to Domain1, and then to Domain2. With a shortcut trust, DomainD and Domain2 would have a direct trust between them that could be used for authentication. As we can also see in Figure 1.18, multiple shortcut trusts can exist, allowing users to be authenticated to other domains that they commonly need to access.

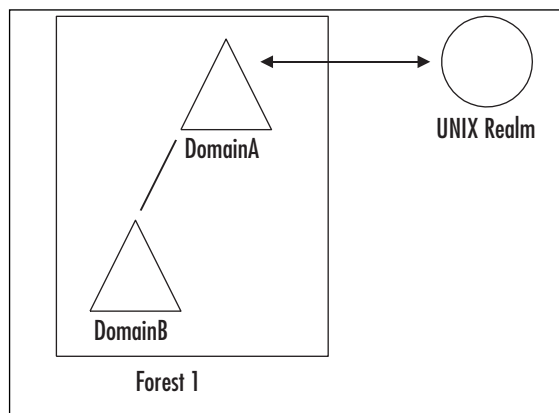
Figure 1.18 Shortcut Trusts



A *forest trust* is also transitive, and can be one-way or two-way. As shown in Figure 1.19, this type of trust is used to connect two different forests, so that users in each forest can use resources in the other. Using this type of trust, a user in a domain in one forest could be authenticated and access resources located in a domain that's in another forest. This allows different areas of the network to be interconnected, even though they are separated by administrative boundaries.

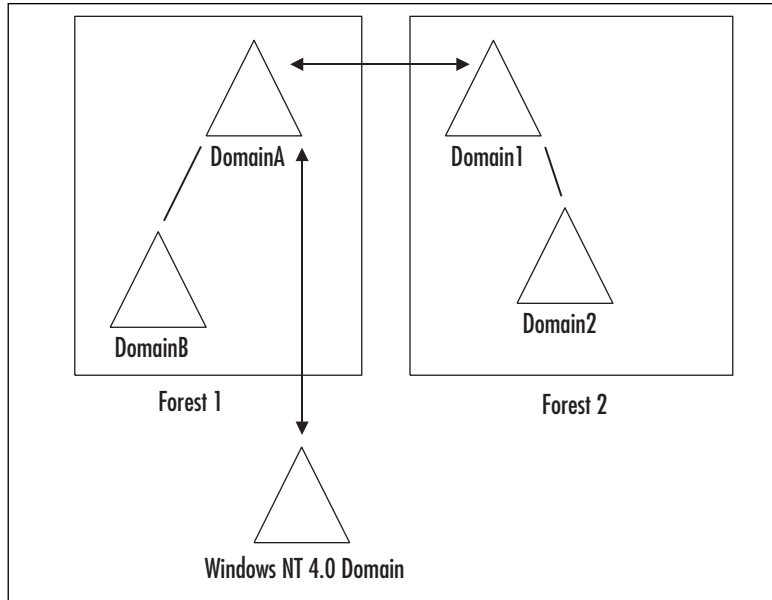
Figure 1.19 Forest Trust

A *realm trust* can be one-way or two-way, and can also be either transitive or nontransitive. *Nontransitive* means that the trust relationship doesn't extend beyond the two parties. For example, let's say DomainA trusts DomainB, and DomainB trusts DomainC. Because the trust is nontransitive, DomainA and DomainC don't trust one another because there isn't a trust relationship between them. As shown in Figure 1.20, the realm trust is used when a relationship needs to be created between a Windows Server 2003 domain and a non-Windows realm that uses Kerberos version 5 (such as one running UNIX).

Figure 1.20 Realm Trust

The final type of trust that can be created is an *external trust*. An external trust is always nontransitive, and can be either one-way or two-way. As shown in Figure 1.21, this type of trust is used to create a relationship between a Windows Server 2003 domain and one running Windows NT 4.0. It can also be used to connect two domains that are in different forests, and don't have a forest trust connecting them.

Figure 1.21 External Trust



The Active Directory Domains and Trusts console is also used for raising domain and forest levels, which enables additional features in Active Directory. Raising domain and forest functional levels depends on what operating systems are running on servers, and is something we discuss in greater detail later in this chapter.



EXAM WARNING

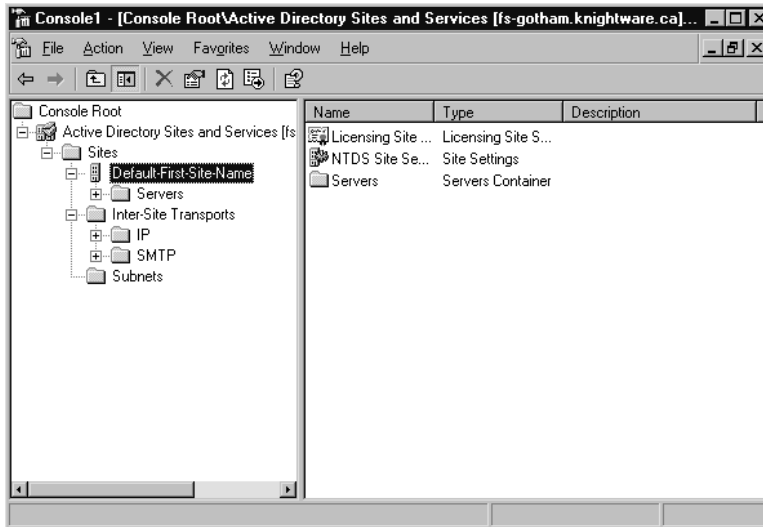
The Active Directory Domains and Trusts console allows you to create different types of trust relationships to share information and resources between forests, domains, and non-Windows Server 2003 networks. You can create one- and two-way transitive trusts, forest trusts, realm trusts, external trusts, and shortcut trusts. Each has a specific use, and cannot be used in all circumstances. You should familiarize yourself with the use of each type of trust.

Active Directory Sites and Services

Earlier in this chapter, we discussed how sites represent the physical structure of your network, and are important to replicating information in Active Directory. The Active Directory Sites and Services console is used to create and manage sites, and control how the directory is replicated within a site and between sites. Using this tool, you can specify connections between sites, and how they are to be used for replication.

As shown in Figure 1.22, the Active Directory Sites and Services console has a number of containers that provide information and functions on creating and maintaining sites. When a domain is first installed on a DC, a site object named Default-First-Site-Name is created. This container can (and should) be renamed to something that is meaningful to the business. As mentioned earlier, additional sites can be created to improve replication between sites, or domains can be added to this existing site.

Figure 1.22 Active Directory Sites and Services



The *Inter-Site Transports* container is used to create and store site links. A *site link* is a connection between sites. Links created under the IP container use the Internet Protocol (IP) as their transport protocol, while those created under SMTP use the Simple Mail Transfer Protocol (SMTP).

The *Subnets* container is used to create and store objects containing information about subnets on your network. Subnets are collections of neighboring computers that are subdivided within the network, using a common network ID. Using the Subnets container, you can group different subnets together to build a site.

Now that we've looked at the MMC, and the snap-ins used to manage Active Directory, let's get a little hands-on experience. In Exercise 1.02, we'll see how the MMC is used to load the snap-ins we'll use in future exercises.

EXERCISE 1.02

ADDING SNAP-INS TO THE MICROSOFT MANAGEMENT CONSOLE

1. From the **Run** command in the Windows **Start** menu, type **MMC**, and click **OK**.
2. When the MMC opens, click the **Add/Remove Snap-in** command on the **File** menu.
3. When the **Add/Remove Snap-in** dialog box appears, click the **Standalone** tab to select it. Click the **Add** button.
4. When the **Add Standalone Snap-in** appears, select **Active Directory Domains and Trusts** from the listing and then click the **Add** button. An entry for this snap-in should appear in the listing in the **Add/Remove Snap-in** dialog box.
5. Select **Active Directory Sites and Services** from the listing and then click the **Add** button. An entry for this snap-in should appear in the listing in the **Add/Remove Snap-in** dialog box.
6. Select **Active Directory Users and Computers** from the listing and then click the **Add** button. An entry for this snap-in should appear in the listing in the **Add/Remove Snap-in** dialog box.
7. Click **Close** to return to the previous screen. At this point, three entries should appear in the **Add/Remove snap-in** dialog box. Click **OK** to close the dialog.
8. The console tree in the MMC should now contain nodes for each snap-in. Expand each snap-in and notice that they contain objects and information relating to the aspects each snap-in deals with.

Command-Line Tools

Windows Server 2003 provides a number of command-line tools that you can use for managing Active Directory. These tools use commands typed in at the prompt, and can provide a number of services that are useful in administering the directory. The command-line tools for Active Directory include:

- **Cacls** Used to view and modify discretionary access control lists (DACLS) on files.
- **Cmdkey** Used to create, list, and delete usernames, passwords, and credentials.

- **Csvde** Used to import and export data from the directory.
- **Dcgpofix** Restores Group Policy Objects (GPOs) to the state they were in when initially installed.
- **Dsadd** Used to add users, groups, computers, contacts, and OUs.
- **Dsget** Displays the properties of an object in Active Directory.
- **Dsmmod** Used to modify users, groups, computers, servers, contacts, and OUs.
- **Dsmove** Renames an object without moving it, or moves an object to a new location.
- **Ldifde** Used to create, modify, and delete objects from Active Directory.
- **Ntdsutil** Used for general management of Active Directory.
- **Whoami** Provides information on the user who's currently logged on.

In the sections that follow, we will briefly discuss each of these tools, and show you how they can assist you in performing certain tasks when administering Active Directory.

Cacls

Cacls is used to view and modify the permissions a user or group has to a particular resource. *Cacls* provides this ability by allowing you to view and change DACLs on files. A DACL is a listing of access control entries (ACEs) for users and groups, and includes permissions the user has to a file. The syntax for using this tool is:

```
Cacls filename
```

Cacls also has a number of switches, which are parameters you can enter on the command line to use a specific functionality. Table 1.1 lists the switches for *Cacls*.

Table 1.1 Switches for the *Cacls* Tool

Parameter	Description
/t	Change the DACLs of files in the current directory and all subdirectories.
/e	Edit the DACL.
/r <i>username</i>	Revokes the users' rights.
/c	Ignore any errors that might occur when changing the DACL.
/g <i>username:permission</i>	Grants rights to a specified user. Rights that can be granted are: n (None), r (Read), w (Write), c (Change), and f (Full Control).

Continued

Table 1.1 Switches for the Cacs Tool

Parameter	Description
<code>/p username:permission</code>	Replaces the rights of a specified user. The rights that can be replaced are: n (None), r (Read), w (Write), c (Change), and f (Full Control).
<code>/d username</code>	Denies access to a specified user.

Cmdkey

Cmdkey is used to create, view, edit, and delete the stored usernames, passwords, and credentials. This allows you to log on using one account, and view and modify the credentials of another user. As with other command-line tools we'll discuss, *cmdkey* has a number of switches that provided needed parameters for the tool to function. Table 1.2 lists these parameters.

Table 1.2 Switches for the Cmdkey Tool

Parameter	Description
<code>/add:targetname</code>	Adds a username and password to the list, and specifies the computer or domain (using the <i>targetname</i> parameter) with which the entry will be associated.
<code>/generic</code>	Adds generic credentials to the list.
<code>/smartcard</code>	Instructs <i>cmdkey</i> to retrieve credentials from a smart card.
<code>/user: username</code>	Provides the username with which this entry is to be associated. If the <i>username</i> parameter isn't provided, you will be prompted for it.
<code>/pass:password</code>	Provides the password to store with this entry. If the <i>password</i> parameter isn't provided, you will be prompted for it.
<code>/delete: {targetname /ras}</code>	Deletes the username and password from the list. If the <i>targetname</i> parameter is provided, the specified entry will be deleted. If <i>/ras</i> is included, the stored remote access entry is deleted.
<code>/list: targetname</code>	Lists the stored usernames and credentials. If the <i>targetname</i> parameter isn't provided, all of the stored usernames and credentials will be listed.

Csvde

Csvde is used to import and export data from Active Directory. This data is comma delimited, so that a comma separates each value. Exporting data in this way allows you to then

import it into other applications (for example, Microsoft Office tools such as Access and Excel). Table 1.3 lists the parameters for this command.

Table 1.3 Switches for the Csvde Tool

Parameter	Description
-i	Used to specify the import mode.
-f <i>filename</i>	Specifies the filename to import or export data to.
-s <i>servername</i>	Sets the DC that will be used to import or export data.
-c <i>string1 string2</i>	Replaces the value of <i>string1</i> with <i>string2</i> . This is often used when importing data between domains, and the DN of the domain data is being exported from (<i>string1</i>) needs to be replaced with the name of the import domain (<i>string2</i>).
-v	Verbose mode.
-j <i>path</i>	Specifies the location for log files.
-t <i>portnumber</i>	The <i>portnumber</i> parameter is used to specify the LDAP port number. By default, the LDAP port is 389 and the GC port is 3268.
-d <i>BaseDN</i>	The <i>BaseDN</i> parameter is used to specify the DN of a search base for data export.
-p <i>scope</i>	Used to set the search scope. The value of the <i>scope</i> parameter can be <i>Base</i> , <i>OneLevel</i> , or <i>SubTree</i> .
-l <i>LDAPAttributeList</i>	Specifies a list of attributes to return in an export query. If this parameter isn't used, then all attributes are returned in the query.
-o <i>LDAPAttributeList</i>	Specifies a list of attributes to omit in an export query.
-g	Used to omit paged searches.
-m	Used to omit attributes that apply to certain objects in Active Directory.
-n	Specifies that binary values are to be omitted from an export.
-k	If errors occur during an import, this parameter specifies that <i>csvde</i> should continue processing.
-a <i>username password</i>	Specifies the username and password to be used when running this command. By default, the credentials of the user currently logged on are used.
-b <i>username domain password</i>	Specifies the username, domain, and password to use when running this command. By default, the credentials of the user currently logged on are used.

Dcgpofix

Dcgpofix is used to restore the default domain policy and default DC's policy to the way they were when initially created. By restoring these GPOs to their original states, any changes that were made to them are lost. This tool has only two switches associated with it:

- **/ignoreschema** Ignores the version number of the schema.
- **/target: {domain | dc | both}** Specifies the target domain, DC, or both.

When the */ignoreschema* switch is used, *dcgpofix* will ignore the version number of Active Directory's schema when it runs. This will allow it to work on other versions of Active Directory, as opposed to the one on the computer on which *dcgpofix* was initially installed. You should use the version of *dcgpofix* that was installed with your installation of Windows Server 2003, as GPOs might not be restored if versions from other operating systems are used.

Dsadd

Dsadd is used to add objects to Active Directory. The objects you can add with this command-line tool are users, computers, groups, OUs, contacts, and quota specifications. To add any of these objects, you would enter the following commands at the command prompt:

- ***dsadd user*** Adds a user to the directory
- ***dsadd computer*** Adds a computer to the directory
- ***dsadd group*** Adds a group to the directory
- ***dsadd ou*** Adds an OU to the directory
- ***dsadd contact*** Adds a contact to the directory
- ***dsadd quota*** Adds a quota specification to the directory

While the commands for this tool are straightforward, there is a variety of arguments associated with each. For full details on these arguments, type the command at the command prompt followed by */?*. This will display a list of parameters for each command.

Dsget

Dsget is used to view the properties of objects in Active Directory. The objects you can view with *dsget* are users, groups, computers, servers, sites, subnets, OUs, contacts, partitions, and quota specifications. To view the properties of these objects, enter the following commands:

- ***dsget user*** Displays the properties of a user
- ***dsget group*** Displays the properties of a group and its membership
- ***dsget computer*** Displays the properties of a computer

- *dsget server* Displays the properties of a DC
- *dsget site* Displays the properties of a site
- *dsget subnet* Displays the properties of a subnet
- *dsget ou* Displays the properties of an OU
- *dsget contact* Displays the properties of a contact
- *dsget partition* Displays the properties of a directory partition
- *dsget quota* Displays the properties of a quota specification

While the commands for this tool are straightforward, there is a variety of arguments associated with each. For full details on these arguments, type the command at the command prompt followed by */?*. This will display a list of parameters for each command.

Dsmod

Dsmod is used to modify existing objects in Active Directory. The objects you can modify using *dsmod* are users, groups, computers, servers, OUs, contacts, partitions, and quota specifications. To edit these objects, enter the following commands:

- *dsmod user* Modifies the attributes of a user in the directory
- *dsmod group* Modifies the attributes of a group in the directory
- *dsmod computer* Modifies a computer in the directory
- *dsmod server* Modifies the properties of a DC
- *dsmod ou* Modifies the attributes of an OU in the directory
- *dsmod contact* Modifies the attributes of a contact in the directory
- *dsmod partition* Modifies a directory partition
- *dsmod quota* Displays the properties of a quota specification

While the commands for this tool are straightforward, there is a variety of arguments associated with each. For full details on these arguments, type the command at the command prompt followed by */?*. This will display a list of parameters for each command.

Dsmove

Dsmove is used to either rename or move an object within a domain. Using this tool, you can rename an object without moving it in the directory, or move it to a new location within the directory tree.



EXAM WARNING

The *dsmove* tool can't be used to move objects to other domains.

Renaming or moving an object requires that you use the DN, which identifies the object's location in the tree. For example, if you have an object called JaneD in an OU called Accounting, located in a domain called syngress.com, the DN is:

```
CN=JaneD, OU=Accounting, DC=syngress, DC=com
```

The `-newname` switch is used to rename objects using the DN. For example, let's say you wanted to change a user account's name from JaneD to JaneM. To do so, you would use the following command:

```
Dsmove CN=JaneD, OU=Accounting, DC=syngress, DC=com -newname JaneM
```

The `-newparent` switch is used to move objects within a domain. For example, let's say the user whose name you just changed was transferred from Accounting to Sales, which you've organized in a different OU container. To move the user object, you would use the following command:

```
Dsmove CN=JaneM, OU=Accounting, DC=syngress, DC=com -newparent OU=Sales,
DC=syngress, DC=com
```

In addition to the `-newname` and `-newparent` switches, you can also use the parameters listed in Table 1.4 to control how this tool is used.

Table 1.4 Switches for Dsmove

Parameter	Description
<code>{-s Server -d Domain}</code>	Specifies a remote server or domain to connect to. By default, <code>dsmove</code> will connect to the DC in the domain you logged on to.
<code>-u Username</code>	Specifies the username to use when logging on to a remote server.
<code>-p {Password *}</code>	Specifies the password to use when logging on to a remote server. If you type the <code>*</code> symbol instead of a password, you are then prompted to enter the password.
<code>-q</code>	Sets <code>dsmove</code> to suppress output.
<code>{-uc -uco -uci}</code>	Specifies <code>dsmove</code> to format input and output in Unicode.

Ldifde

`Ldifde` is used to create, modify, and delete objects from the directory, and can also be used to extend the schema. An additional use for this tool is to import and export user and group information. This allows you to view exported data in other applications, or populate Active Directory with imported data. To perform such tasks, `ldifde` relies on a number of switches that enable it to perform specific tasks, listed in Table 1.5.

Table 1.5 Switches for Ldifde

Parameter	Description
-l	Sets Ldifde to import data. If this isn't specified, then the tool will work in Export mode.
-f <i>Filename</i>	Specifies the name of the file to import or export.
-s <i>Servername</i>	Specifies the DC that will be used to perform the import or export.
-c <i>string1 string2</i>	Replaces the value of <i>string1</i> with <i>string2</i> . This is often used when importing data between domains, and the DN of the domain data is being exported from (<i>string1</i>) needs to be replaced with the name of the import domain (<i>string2</i>).
-v	Verbose mode.
-j <i>path</i>	Specifies the location for log files.
-t <i>portnumber</i>	The <i>portnumber</i> parameter is used to specify the LDAP port number. By default, the LDAP port is 389 and the GC port is 3268.
-d <i>BaseDN</i>	The <i>BaseDN</i> parameter is used to specify the DN of a search base for data export.
-p <i>scope</i>	Used to set the search scope. The value of the scope parameter can be Base, OneLevel, or SubTree.
-r <i>LDAPfilter</i>	Specifies a search filter for exporting data.
-l <i>LDAPAttributeList</i>	Specifies a list of attributes to return in an export query. If this parameter isn't used, then all attributes are returned in the query.
-o <i>LDAPAttributeList</i>	Specifies a list of attributes to omit in an export query.
-g	Used to omit paged searches.
-m	Used to omit attributes that apply to certain objects in Active Directory.
-n	Specifies that binary values are to be omitted from an export.
-k	If errors occur during an import, this parameter specifies that Ldifde should continue processing.
-a <i>username password</i>	Specifies the username and password to be used when running this command. By default, the credentials of the user who's currently logged on are used.
-b <i>username domain password</i>	Specifies the username, domain, and password to use when running this command. By default, the credentials of the user who's currently logged on are used.

Ntdsutil

Ntdsutil is a general-purpose command-line tool that can perform a variety of functions for managing Active Directory. Using *Ntdsutil*, you can:

- Perform maintenance of Active Directory
- Perform an authoritative restore of Active Directory
- Modify the Time To Live (TTL) of dynamic data
- Manage domains
- Manage data in the directory and log files
- Block certain IP addresses from querying the directory, and set LDAP policies
- Remove metadata from DCs that were retired or improperly uninstalled
- Manage Security Identifiers (SIDs)
- Manage master operation roles (Domain Naming Master, Schema Master, Infrastructure Master, PDC Emulator, and RID Master)

Typing **ntdsutil** at the command prompt will load the tool and the prompt will change to **ntdsutil:**. As shown in Figure 1.23, by typing **help** at the command line, you can view different commands for the tasks being performed. After entering a command, typing **help** again will provide other commands that can be used. For example, typing **metadata cleanup** after first starting **ntdsutil**, and then typing **help** will display a list of commands relating to metadata cleanup. This allows you to use the command as if you were navigating through menus containing other commands. You can return to a previous menu at any time, or exit the program by typing **Quit**.

Figure 1.23 NTDSUTIL

```

C:\WINDOWS\system32>ntdsutil
ntdsutil: help

?
Authoritative restore          - Show this help information
Configurable Settings         - Authoritatively restore the DIT database
Domain management             - Manage configurable settings
Files                         - Prepare for new domain creation
Help                           - Manage NTDS database files
LDAP policies                  - Show this help information
Metadata cleanup              - Manage LDAP protocol policies
Popups %s                      - Clean up objects of decommissioned servers
Quit                           - (en/dis)able popups with "on" or "off"
Roles                          - Quit the utility
Security account management   - Manage NTDS role owner tokens
D Cleanup                      - Manage Security Account Database - Duplicate SI
Semantic database analysis     - Semantic Checker
Set DSRM Password             - Reset directory service restore mode administra
tor account password

ntdsutil:

```


Whoami

Whoami is a tool for displaying information about the user who is currently logged on. Using this tool, you can view your domain name, computer name, username, group names, logon identifier, and privileges. The amount of information displayed depends on the parameters that are entered with this command. Table 1.6 lists the available parameters.

Table 1.6 Switches for Whoami

Parameter	Description
/upn	Displays the UPN of the user currently logged on.
/fqdn	Displays the FQDN of the user currently logged on.
/logonid	Displays the Logon ID.
/user	Displays the username of the user currently logged on.
/groups	Displays group names.
/priv	Displays privileges associated with the currently logged-on user.
/fo <i>format</i>	Controls the format of how information is displayed. The <i>format</i> parameter can have the value of: <i>table</i> (to show output in a table format), <i>list</i> (to list output), or <i>csv</i> to display in a comma-delimited format.
/all	Displays username, groups, SIDs, and privileges for the user currently logged on.

EXERCISE 1.03

USING WHOAMI

1. From the Windows **Start** menu, click **Command Prompt**.
2. When the **Command Prompt** opens, type **WHOAMI** at the prompt and then press the **Enter** key. The output will show the account you are currently logged on with.
3. Type **WHOAMI /UPN** and then press **Enter**. The UPN of the currently logged-on user will be displayed on the screen.
4. Type **WHOAMI /FQDN** and then press **Enter**. The FQDN of the user that's currently logged on will appear on the screen.
5. Type **WHOAMI /PRIV** and then press **Enter**. A listing of privileges associated with the account you are currently logged on with should appear on the screen.

5. Type **WHOAMI /ALL** and then press **Enter**. As shown in Figure 1.24, a listing of information relating to the account you're currently logged on with will be listed on the screen.

Figure 1.24 Results of Using the *WHOAMI /ALL* Command

```

C:\Documents and Settings\Administrator>whoami /all

USER INFORMATION
-----
User Name                SID
=====
knighware\administrator S-1-5-21-2331013139-2269155295-1045398719-500

GROUP INFORMATION
-----
Group Name                Attributes                Type                SID
=====
Everyone
Mandatory group, Enabled by default, Enabled group S-1-1-0
BUILTIN\Administrators  Mandatory group, Enabled by default, Enabled group, Group owner S-1-5-32-544
BUILTIN\Users           Mandatory group, Enabled by default, Enabled group Alias S-1-5-32-545
  
```

EXAM
70-294
OBJECTIVE
1

Implementing Active Directory Security and Access Control

Security is an important part of Windows Server 2003 and Active Directory. Two primary methods of implementing security are user authentication and access control. *Authentication* is used to verify the identity of a user or other objects, such as applications or computers. After it's been determined they are who or what they say they are, the process continues by giving them the level of access they deserve. *Access control* manages what users (or other objects) can use, and how they can use them. By combining authentication and access control, a user is permitted or denied access to objects in the directory.

Access Control in Active Directory

In Active Directory, permissions can be applied to objects to control how these objects are used. Permissions regulate access by enforcing whether a user can read or write to an object, has full control, or no access. Three elements determine a user's access, and define the permissions they have to an object:

- Security descriptors
- Object Inheritance
- Authentication



NOTE

Active Directory permissions are separate from share permissions (also called shared folder permissions) and NTFS permissions (also called file-level permissions), and work in conjunction with both.

Objects in Active Directory use security descriptors to store information about permissions, and control who has access to an object. The security descriptor contains information that's stored in access control lists (ACLs), which define who can access the object and what they can do with it. There are two different types of ACLs in the security descriptor:

- Security access control list (SACL)
- Discretionary access control list (DACL)

The SACL is used to track an object's security based on how a user or group accesses the object. For example, you can audit whether a user was able to access the object using a particular permission (such as Read, Write, or Full Control). Information about what to audit is kept in ACEs, which are stored within the SACL. These entries control what is audited, and contain information about the events to be logged. In doing this, records can be kept on the security of objects, and whether specific users or groups are able to successfully access them.

As we saw earlier, when we discussed command-line tools for Active Directory, a DACL is a listing of ACEs for users and groups, and includes information about the permissions that a user or group has to a file. The DACL controls whether a user is granted or denied access to an object. ACEs in the DACL explicitly identify individual users and groups, and the permissions granted to each. Because only users and groups identified in the DACL can access an object in Active Directory, any user or group that isn't specified is denied access.

Active Directory places the permissions you can apply to objects into two categories: standard permissions and special permissions. *Standard permissions* are those that are commonly applied to objects, whereas *special permissions* provide additional access control. For most objects in Active Directory, five permissions are available as standard permissions:

- **Full Control** Allows the user to change permissions, take ownership, and have the abilities associated with all other standard permissions.
- **Read** Allows the user to view objects, attributes, ownership, and permissions on an object.

- **Write** Allows the user to change attributes on an object.
- **Create All Child Objects** Allows the user to add objects to an OU.
- **Delete All Child Objects** Allows the user to delete objects from an OU.

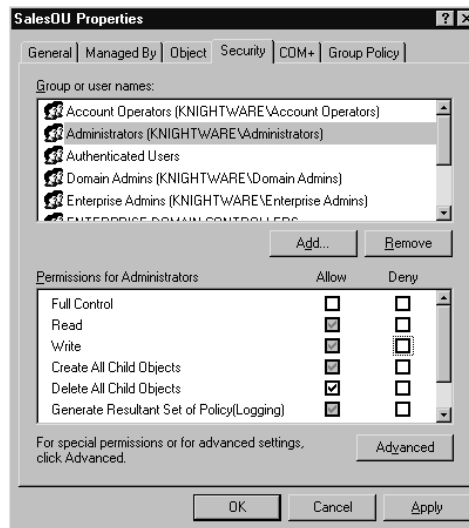
Permissions can be set on objects by using the **Active Directory Users and Computers** snap-in for the MMC. As shown in Figure 1.25, you can set permissions by using the **Security** tab of an object's **Properties** dialog box. The **Security** tab is hidden in the **Properties** dialog box, unless the **Advanced Features** menu item is toggled on the **View** menu first. After this is done, you can then bring up the **Properties** dialog box by selecting an object and clicking **Properties** on the **Action** menu, or right-clicking on the object and selecting **Properties**.



EXAM WARNING

Because changing permissions can cause major problems if done incorrectly, by default the **Security** tab is hidden and needs to be enabled by turning on the **Advanced Features** for Active Directory Users and Computers. Until this is done, you will not be able to modify permissions.

Figure 1.25 Permissions Are Set on the Security Tab of the Object's Properties



The top pane of the Security tab lists users and groups, and the lower pane lists the various permissions that can be applied to these users and groups. You can set permissions by selecting one of these users and groups, and checking the applicable permissions. Special permissions can be set for objects by clicking the **Advanced** button, which displays a dialog box where additional permissions can be applied.

Because it would take a while to assign permissions to every object in Active Directory, object inheritance can be used to minimize how often and where permissions are assigned. *Object inheritance* refers to how the permissions of a parent object are inherited by child objects. When permissions are applied to a container, they are propagated to objects within that container. For example, if a group had Full Control permissions on an OU, the group would also have Full Control of any of the printer objects within that OU. The permissions of one object flow down to any objects within the hierarchy, so child objects have the same permissions as their parents.

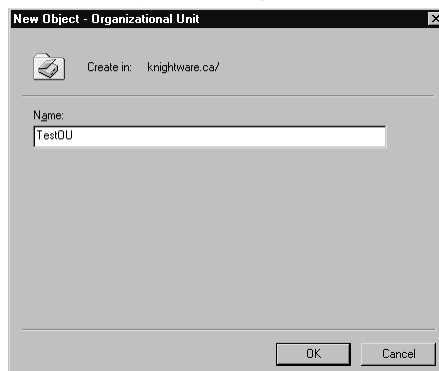
Since there might be times when you don't want the permissions from a parent to propagate to child objects, inheritance can be blocked. By clearing the **Allow Inheritable Permissions From Parent To Propagate To This Object** check box, the permissions from containers higher in the hierarchy are blocked. When this is done, any permissions that are modified on parent objects don't apply to the child. Permissions for the child object must be explicitly assigned.

EXERCISE 1.04

SETTING PERMISSIONS ON ACTIVE DIRECTORY OBJECTS

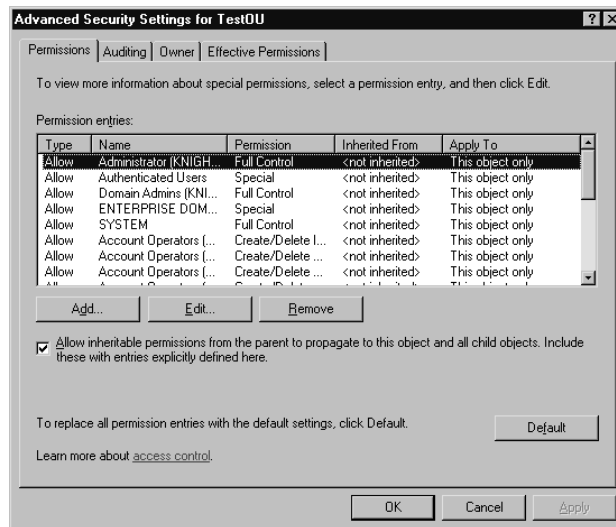
1. Open **Active Directory Users and Computers** by clicking selecting **Administrative Tools** in the Windows **Start** menu, and then clicking on the **Active Directory Users and Computers** menu item.
2. When the MMC opens with this snap-in installed, expand the console tree so that your domain and the containers within it are visible.
3. Select your domain from the console tree. From the **Action** menu, select **New** and then click the **Organizational Unit** menu item. As shown in Figure 1.26, when the dialog box appears, name the new OU **TestOU**, and then click **OK**. A new OU with this name should now appear in the console tree beneath your domain.

Figure 1.26 New Object Dialog Box



4. In the **View** menu, click **Advanced Features**.
5. Select the **TestOU** OU. From the **Action** menu, click **Properties**.
6. When the Properties dialog box appears, click the **Security** tab. In the list of usernames, select the name of the account you're currently logged on with.
7. In the pane below the list of usernames and groups, click the **Full Control** check box under **Allow**, so that a check mark appears in it. You now have full control of the OU.
8. Click the **Advanced** button to display the **Advanced Security Settings** dialog box. When the dialog box appears, click the **Permissions** tab. As shown in the Figure 1.27. Ensure that the **Allow inheritable permissions from the parent to propagate to this object and all child objects** check box is checked. This will allow inheritable permissions to be applied to this OU, and any within the container. Click **OK** to return to the previous screen.

Figure 1.27 Advanced Settings Dialog Box



9. Click **OK** to exit the Properties dialog box.

Role-Based Access Control

Access control can be managed based on the *role* an Active Directory object plays in an organization. Since objects represent users, computers, and other tangible elements of an organization, and these people and things serve different purposes in a company, it makes sense to configure these objects so that they reflect the tasks they perform. *Role-based administration* is used to configure object settings, so that computers and users have the necessary permissions needed to do their jobs based on the roles they fill.

The roles that users and computers are assigned correspond to the functions they serve in a company. Two categories of roles can be used for role based access control: authorization and computer configuration.

Authorization roles are based on the tasks a person performs as part of his or her job. For example, Help Desk personnel would need the ability to change passwords, while accountants would need to be able to access financial information and audit transactions. Using role-based access control, you can give each person the access he or she needs to perform these tasks.

Authorization roles are similar to security groups, to which users can become members and acquire a level of security that gives them the ability to perform certain tasks. However, authorization roles differ in that they are used for applications. Role-based access can be applied to a single application, set of applications, or a scope within the application. Another important difference is that role-based authorization can be dynamic, so that users become part of a group membership as an application runs. This is different from security groups that require membership to be set beforehand.

In the same way that users have different purposes in a company, so do computers. A business might have DCs, mail servers, file servers, Web servers, and any number of other machines providing services to users and applications in an organization. *Computer configuration roles* are used to control which features, services, and options should be installed and configured on a machine, based on the function it serves in the company.

Authorization Manager

Authorization Manager is a snap-in for the MMC that allows you to configure role-based access for applications. By using roles, you ensure that users only have access to the functions and resources they need to perform their jobs, and are prohibited from using other features and resources they're not authorized to use. For example, personnel in Payroll would need to view information on employees (so they can be paid), but wouldn't need to access administrative features that allow them to modify passwords.

In Authorization Manager, roles are designed based on the tasks that are supported by the application. After the role is developed, users and groups can then be assigned to the role so they have the access necessary to perform these tasks. The tasks that are available for users to use depend on the application, as the ability to support roles and the functions available are part of the software design.

Active Directory Authentication

When you log on to a Windows Server 2003 domain, a single logon gives access to any resources you're permitted to use, regardless of their location on the network. A user doesn't need to re-enter a password every time the user accesses a server or other resources, because any authentication after initially logging on is transparent. Because only one logon is needed, the system needs to verify a person is who he or she claims to be, before any access is given.

Authentication is used to verify a user's logon credentials. The primary method of determining the identity of a user is by logging on to the local computer and network, where a person enters a username and password. If these don't match the username and password for the local computer or Active Directory account, the person isn't able to gain access.

Operating systems such as Windows NT, 2000, and Server 2003 store account information in the SAM database. The SAM stores credentials that are used to access the local machine. When a user logs on to a computer with a local user account that's stored in the SAM, the user is authenticated to the local machine. The user's access is limited to just that computer when logging on to the machine.

When users log on to the Windows Server 2003 domain, an account in Active Directory is used to access network resources located within the domain, or in other trusted domains. When a user logs on, the Local Security Authority (LSA) is used to log users on to the local computer. It is also used to authenticate to Active Directory. After validating the user's identity in Active Directory, the LSA on the DC that authenticates the user creates an access token and associates a SID with the user.

The access token is made up of data that contains information about the user. It holds information about the user's name, group affiliation, SID, and SIDs for the groups of which he or she is a member. The access token is created each time the user logs on. Because the access token is created at logon, any changes to the user's group membership or other security settings won't appear until after the user logs off and back on again. For example, if the user became a backup operator, he or she would have to log off and log back on before these changes affected the user's access.



TEST DAY TIP

Access control and authentication are vital parts of Active Directory's security, so it is important that you understand the features and controls of Active Directory. The initial security feature that users will experience is the interactive logon. When users log on, an access token is created to indicate the user's security capabilities. When changes are made to a user's account, they will not apply to the user until that user logs on to the domain.

Standards and Protocols

Authentication relies on standards and protocols that are used to confirm the identity of a user or object. Windows Server 2003 supports several types of network authentication:

- Kerberos
- X.509 certificates
- Lightweight Directory Access Protocol/Secure Sockets Layer (LDAP/SSL)
- Public Key Infrastructure (PKI)

As we'll see in the paragraphs that follow, some of these standards and protocols not only provide a method of authenticating users, but also the ability to encrypt data. By encrypting data, you ensure that unauthorized users and applications won't be able to view or modify the data. The data is encoded at one end, and decoded at the other. By providing encryption/decryption features, the privacy of information is better maintained.

Kerberos

Kerberos version 5 is an industry standard security protocol that Windows Server 2003 uses as the default authentication service. It is used to handle authentication in Windows Server 2003 trust relationships, and is the primary security protocol for authentication within domains.

Kerberos uses mutual authentication to verify the identity of a user or computer, and the network service being accessed. Each side proves to the other that they are who they claim to be. Kerberos does this through the use of *tickets*.

A Kerberos ticket is encrypted data that's issued for authentication. Tickets are issued by a *Key Distribution Center* (KDC), which is a service that runs on every DC. When a user logs on, the user authenticates to Active Directory using a password or smart card. Because the KDC is part of Active Directory, the user also authenticates to the KDC and is issued a session key called a *Ticket Granting Ticket* (TGT). The TGT is generally good for as long as the user is logged on, and is used to access a ticket granting service that provides another type of ticket: *service tickets*. A service ticket is used to authenticate to individual services, by providing the ticket when a particular service is needed.

Kerberos Made Easy

What with all the different elements making up the Kerberos process of authentication, it can be a little difficult wrapping your head around everything that's going on. A good way of understanding and remembering something is to compare it to something familiar to you.

Being authenticated by Kerberos is a little like going to a theme park. The TGT allows you to get into the park, where you can now get tickets to go on the rides. These secondary tickets allow you to use services, and identify that you're allowed to use them.

With this analogy in mind, let's take a second look at how Kerberos works:

1. The user logs on, and authenticates to the KDC.
2. A TGT is acquired from the KDC, which is then handed to the ticket granting service.
3. The ticket granting service issues a service ticket to the client.
4. The service ticket is handed to the network service you want to access.

X.509 Certificates

X.509 is a popular standard for digital certificates, published by the International Organization for Standardization (ISO). X.509 certificates are used to verify that the user is who he or she claims to be. Digital certificates work as a method of identifying the user, much as your birth certificate is used to identify you as a person. They can also be used to establish the identity of applications, network services, computers, and other devices.

X.509 specifies the syntax and format of digital certificates; in other words, it explains what is to be included in a digital certificate. An X.509 certificate includes information about the user to whom the certificate was issued, information about the certificate itself, and can include information about the issuer of the certificate (referred to as the certification authority (CA)). To prevent the certificate from being used indefinitely, it also contains information about the time period during which the certificate is valid.

LDAP/SSL

LDAP is used by Active Directory for communication between clients and directory servers. LDAP allows you to read and write data in Active Directory, but isn't secure by default. To extend security to LDAP communications, LDAP can be used over Secure Sockets Layer/Transport Layer Security. Secure Sockets Layer (SSL) and Transport Layer Security (TLS) provide data encryption and authentication. TLS is the successor to SSL, and is more secure. It can be used by clients to authenticate servers, and by servers to authenticate clients. Communication using TLS allows messages between the client and server to be encrypted, so data being passed between the two isn't accessible by third parties.

PKI

Public Key Infrastructure (PKI) provides a means for organizations to secure their communications and transactions through the use of digital certificates and public key cryptography. Certificate Authorities (CAs) are an integral part of a PKI and are used to create and manage the digital certificates and public keys that are throughout the enterprise. Public key cryptography is used in combination with digital certificates for a variety of purposes, which include authentication, authorization, confidentiality of data, verification of data integrity, and non-repudiation. Public key cryptography uses two types of keys: a private key and a public key.

For data confidentiality, the public key is used to encrypt session keys and data, and the private key is used for decryption. The public key is openly available to the public, while the private key is secret and known only to the person for whom it is created. The members of a key pair are mathematically related, but you cannot extrapolate the private key by knowing the public key. Using the two keys together, messages can be encrypted and decrypted using public key cryptography. Furthermore, only the possessor of the private key can decrypt the message encrypted with the public key.


For authentication, the roles of the public and private keys are reversed. The private key is used for encryption, and the public key is used for decryption. The private key is unique to the person being identified, so each user has his or her own private key for authentication purposes. Because each private key has a corresponding public key, the public key is used to decrypt information used for authenticating the user.

The public and private keys are generated at the same time by a CA. The CA creates and manages keys, binding public and private keys to create certificates, and vouching for the validity of public keys belonging to users, computers, services, applications, and other CAs.

In addition to a CA, a registration authority (RA) can also be used to request and acquire certificates for others. The RA acts as a proxy between the user and the CA, and relieves the CA of some of the burden of verification. When a user makes a request to a CA, the RA can intercept the request, authenticate it, and then pass it on to the CA. When the CA responds to the request, it sends it to the RA, which then forwards it to the user.

Private and public keys are created when someone or something needs to establish the validity of his, her, or its identity. When the public and private keys are created, the private key is given to the person or entity that wants to establish the credentials, and a public key is stored so that anyone who wants to verify these credentials has access to it. When a person wants to send a message using public key cryptography with the data encrypted so that it cannot be read by anyone but the holder of the private key, the public key is acquired from the CA and used to encrypt the message. When a person who holds the private key receives the message, the public key is validated with the CA. Since the CA is trusted, this validates the authenticity of the message. After this is done, the private key is then used to decrypt the message.

Conversely, if a person wants to send a message and ensure that he or she is the actual sender, that person can encrypt the message with his or her private key. Then, the recipient decrypts it with the sender's public key, thereby proving that the message really did come from that sender.



What's New in Windows Server 2003 Active Directory?

A number of enhancements and new features in the Windows Server 2003 Active Directory weren't available in Windows 2000 Server. These improvements allow various tasks and network operations to be performed more efficiently. However, although there are many new features, the availability of a number of them depends on the environment in which DCs are running.

When a Windows Server 2003 DC is created on a network, Active Directory is installed with a basic set of features. Additional features can be enabled, but this is dependent on the operating systems running as DCs and the functional level (formerly called the mode) that's configured for the domain or forest. There are four different levels of functionality for Active Directory:

- Windows 2000 mixed
- Windows 2000 native
- Windows 2003 interim
- Windows 2003

If you're upgrading from Windows 2000 Server on your network, you're probably familiar with the first two levels. Each of these appeared in Windows 2000, and provided backward compatibility to older operating systems such as Windows NT 4.0, and allowed control of what features were available in Active Directory. Windows Server 2003 interim and Windows Server 2003 functionality are new to Active Directory, and weren't available in previous versions.

Windows 2000 mixed allows domains to contain Windows NT BDCs that can interact with Windows 2000 and Windows Server 2003 servers. In this level, the basic features of Active Directory are available to use. However, you aren't able to nest groups within one another, use Universal Groups that allow access to resources in any domain, or use Security ID Histories (SIDHistory). Because it accommodates the widest variety of servers running on your network, this is the default level of functionality when a Windows Server 2003 DC is installed.

Windows 2000 native is the highest mode available for Windows 2000 and the next highest level for Windows Server 2003 DCs. Windows 2000 native removes support for replication to Windows NT BDCs, so these older servers are unable to function as DCs. In this level, only Windows 2000 and Windows Server 2003 DCs can be used in the domain, and support for Universal Groups, SIDHistory, and group nesting becomes available.

Windows 2003 interim is a new level that's available in Windows Server 2003. This level is used when your domain consists of Windows NT and Windows Server 2003 DCs. It provides the same functionality as Windows 2000 mixed mode, but is used when you are upgrading Windows NT domains directly to Windows Server 2003. If a forest has never had Windows 2000 DCs, then this is the level used for performing an upgrade.

The highest functionality level for Active Directory is *Windows 2003*. The Windows 2003 level is used when there are only Windows Server 2003 DCs in the domain. When this level is set for the domain, a considerable number of features are enabled. We discuss these features later in this chapter, when we discuss new features that are available with domain and forest functionality.

The number of features available for Active Directory is also dependent on whether the functionality level has been raised for the domain or the entire forest. With domain-level functionality, all servers in the domain are running Windows Server 2003. With this level, different domains in a forest can be set to use different functionality levels. With forest-level functionality, all domains in the forest are running Windows Server 2003 and have their domain functionality raised to Windows Server 2003. As stated previously, there are four different levels for Windows Server 2003 domain functionality.

Forest functionality can also be raised to enable features that apply to all domains in the forest. With forest functionality, there are three different levels available:

- Windows 2000
- Windows 2003 interim
- Windows 2003

Windows 2000 level allows Windows NT, Windows 2000, and Windows Server 2003 DCs on the network, and is the default level for a forest. The other two levels are the same as the domain levels, in that Windows 2003 interim supports Windows Server 2003 DCs and NT BDCs, while Windows 2003 level supports only Windows Server 2003 DCs on the network. When the default level is raised to either of these other levels, additional features in Active Directory become available.

To raise the forest functionality, you must first raise the functionality of domains within the forest. Each domain in the forest must be raised to either Windows 2000 native or Windows 2003 before the forest functionality can be raised to Windows 2003. When the forest functional level is then raised to Windows 2003, any DCs in the forest's domains will have their domain functional level automatically raised to Windows 2003.



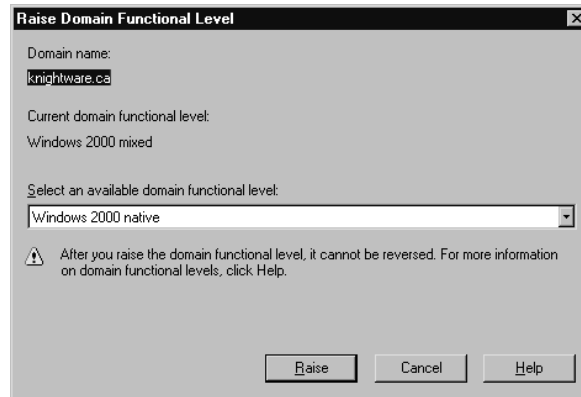
TEST DAY TIP

New features might be dependent on first raising the functional level of the domain or forest. Remember which operating systems are allowed to exist at specific levels, and which features are available when all DCs are running Windows Server 2003.

The tool used to raise domain and forest functional levels is *Active Directory Domains and Trusts*. Raising domain levels is done by right-clicking the domain in the left console pane and then clicking **Raise Domain Functional Level** from the menu that appears. As shown in Figure 1.28, you then select the level to which you want to raise the domain, and

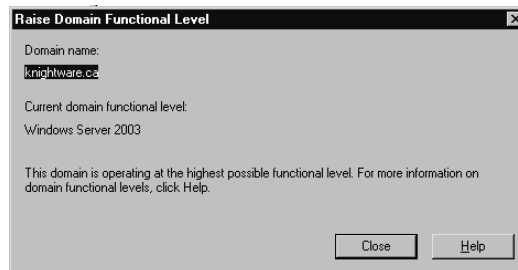
then click the **Raise** button. Raising forest functional levels is done similarly. To raise the forest level, right-click the **Active Directory Domains and Trusts** node, and then click **Raise Forest Functional Level** from the menu that appears (see Figure 1.28). Select the level to which you want to raise the forest, and click **Raise** to complete the task.

Figure 1.28 Raise Domain Functional Level Dialog Box



When raising the forest or domain functional levels, it is important to remember that it is a one-way change. After raising the level, you cannot lower it again later. For example, if you raise the domain from Windows 2000 mixed to Windows 2003, you cannot return the level to Windows 2000 mixed again. This means that you can't add Windows NT BDCs or Windows 2000 DCs to your domain after the upgrade, and any existing DCs need to be upgraded or permanently removed from service. If you attempt to change the domain or forest level after raising it to Windows 2003, a screen similar to Figure 1.29 will appear.

Figure 1.29 Raise Domain Functional Level Dialog Box After Raising the Domain Functional Level





EXAM WARNING

Raising the functional level is a one-way change. You cannot change your mind later and then lower the level to accommodate older operating systems being used as DCs. Before raising the level, ensure that operating systems not running Windows Server 2003 can function in the new level or are upgraded or taken offline permanently before the change.

New Features Available on All Windows Server 2003 Computers

Before we look at the individual features that become available when you raise the domain or forest level, let's first discuss the new features available regardless of whether the domain or forest level has been raised. The features and tools we'll discuss next are available on all versions of Windows Server 2003 that can act as DCs.

As discussed earlier in this chapter, a number of command-line utilities for Active Directory enable administrators to perform certain tasks from the command prompt. This allows administrators to manually enter commands to run operations from a command prompt, or use these commands in batch files or scripts that can be scheduled to run at certain times.

We also saw earlier that the directory uses partitions to separate data into different collections, and that the application partition is used to store data that's needed by specific applications. Because this application-specific data is stored in its own partition, you can configure Active Directory to replicate only this information to other DCs. Not replicating the entire directory cuts down on the amount of time and network traffic needed to copy data to other DCs.

Another new capability provided in Windows Server 2003 is that DCs can be created from backups. Backups are used to copy data to other media such as tapes, and can be used to restore lost data if problems arise. For example, if the hard disk on a server fails, you can use the backup to restore the data to another disk and have the server up and running again. This same process can be used to restore Active Directory to a new DC, avoiding the need to replicate the entire directory to the DC across the network. Allowing additional DCs to be added to an existing domain through the use of backups reduces the time it takes to set up new DCs on the network.

You can use encryption to protect information that is being transmitted across the network. As previously discussed, LDAP can be used over SSL to encrypt data and ensure that data isn't tampered with. This protection prevents unauthorized users from accessing data over the network.

Active Directory allows you to select multiple user objects, so that you can change the attributes of more than one object at a time. After selecting two or more user objects in **Active Directory Users and Computers**, you can bring up the properties and modify

the attributes that are common to each object. This capability makes it faster to manage users, because you don't have to make changes to each account individually.

Active Directory also provides the capability to drag and drop objects into containers. By selecting an object with your mouse, you can then hold down your left mouse button to drag the object to another location (such as another OU). Releasing the left button drops the object into the container. This capability also makes it easy to add user and group objects to groups. Dragging and dropping a user or group into another group adds it to the group membership.

As we'll see in the next chapter, a new object class has been added to Active Directory called `InetOrgPerson`. `InetOrgPerson` is a type of object that's used to represent users in non-Microsoft directory services, and used just as a user object. The presence of this type of class is important when directory information is migrated to Active Directory from these directories.

To prevent users, computers, and groups from creating an unlimited number of objects in Active Directory, Windows Server 2003 has added quotas. Active Directory quotas are used to limit how many objects are owned in a particular directory partition. While quotas can be applied to almost every user, computer, and group, Domain Administrators and Enterprise Administrators are exempted from these limits.

The quotas that are used to limit the ability of a user, computer, or group from creating too many objects in Active Directory should not be confused with disk quotas, which are also available on Windows Server 2003 servers (regardless of the functionality level being used). Disk quotas can be used to limit the amount of hard disk space that can be used on a volume that's formatted in NTFS. The NTFS file system is more advanced than other file systems such as FAT or FAT32, which can also be used to format volumes. By using disk quotas on an NTFS volume, administrators can prevent users from filling up the hard disk with an unlimited number of files.

Finally, searching for objects in Active Directory is easier and more efficient in Windows Server 2003. Active Directory uses object-orientated searches to minimize network traffic, and provides the capability to save queries so that they can be reused repeatedly. The capability to save commonly used queries in **Active Directory Users and Computers** is a topic we'll look at in detail in Chapter 2.

New Features Available Only with Windows Server 2003 Domain/Forest Functionality

When the domain or forest functional levels have been raised so that all DCs are running Windows Server 2003, a number of new features become enabled. These features allow you to modify elements of both your domain and forest, and provide advanced functions that aren't available until functionality levels are raised. In the paragraphs that follow, we will look at the new features available in Active Directory when all DCs have been upgraded to Windows Server 2003, and the functionality has been raised to Windows 2003.

Domain Controller Renaming Tool

The DC renaming tool allows you to rename a DC without having to demote it first. This can be useful when you need to restructure the network, or simply want to use a more meaningful name for a particular DC. When this tool is used, the DC name changes, and any Active Directory and DNS entries are automatically updated.

Domain Rename Utility

Domains can also be renamed. Using the domain rename utility (`rendom.exe`), you can change the NetBIOS and DNS names of a domain, including any child, parent, domain-tree, or forest root domains (from which all others branch off in the hierarchy). By renaming domains in this manner, you can thereby move them in the hierarchy. For example, you can change the name of `dev.web.syngress.com` to `dev.syngress.com`, making the `web.syngress.com` and `dev.syngress.com` domains on the same level of the hierarchy. You could even rename the domain so that it becomes part of a completely different domain tree. The only domain that you can't reposition in this manner is the forest root domain.

Forest Trusts

As we saw earlier, forest trusts can also be created, so that a two-way transitive trust relationship exists between two different forests. In creating such a trust, the users and computers in each forest are able to access what's in both forests. This expands the network, so users are able to use services and resources in both forests.

Dynamically Links Auxiliary Classes

Additional features have also been added to the schema. Windows Server 2003 supports *dynamically linked auxiliary classes*, which allow additional attributes to be added to individual objects. For example, you can have an auxiliary class that has attributes that are used for the Accounting department, and others that are useful for the Sales department. By applying the auxiliary classes to the objects, only those objects are affected. Rather than adding attributes to an entire class of objects, dynamically linking auxiliary classes allows you to apply additional attributes to a selection of objects.

Disabling Classes

Because certain objects in Active Directory might no longer be needed after a specific point, you can disable classes and attributes that are no longer needed in the schema. Classes and attributes can be disabled, but cannot be deleted. If schema objects are no longer required, you can deactivate them, and reactivate them later if the situation changes.

Replication

Improvements have also been made in how Active Directory replicates directory data. Rather than having the entire group membership replicated as a single unit, individual

members of groups can now be replicated to other DCs. In addition, changes have been made to GC replication. When there is an extension of a partial attribute set, only the attributes that have been added are replicated. These improvements decrease the amount of network traffic caused by replication because less data is transmitted across the network.

EXERCISE 1.05

RAISING DOMAIN AND FOREST FUNCTIONALITY

This exercise should not be performed on a production network. It assumes that all DCs in the domain are running Windows Server 2003. After raising the functional levels, you will not be able to roll back to a previous level.

1. From the Windows **Start** menu, select **Administrative Tools**, and then click the **Active Directory Domains and Trusts** menu item.
2. When Active Directory Domains and Trusts opens, expand the **Active Directory Domains and Trusts** node, and select your domain.
3. From the **Action** menu, click **Raise Domain Functional Level**.
4. When the **Raise Domain Functional Level** dialog box appears, select **Windows Server 2003** from the drop-down list. Click the **Raise** button.
5. A warning message will appear, informing you that this action will affect the entire domain, and after you raise the domain functional level, it cannot be reversed. Click **OK**.
6. After you raise the level, a message box will inform you that the action was successful. Click **OK** to continue.
7. In the context pane of Active Directory Domains and Trusts, select the **Active Directory Domains and Trusts** node.
8. From the **Action** menu, click **Raise Forest Functional Level**.
9. When the **Raise Forest Functional Level** dialog box appears, select **Windows Server 2003** from the drop-down list. Click the **Raise** button.
10. A warning message will appear, informing you that this action will affect the entire forest, and after you raise the forest functional level, it cannot be reversed. Click **OK**.
11. After you raise the level, a message box will inform you that the action was successful. Click **OK** to continue.

Summary of Exam Objectives

Active Directory is a database with a hierarchical structure, storing information on accounts, resources, and other elements making up the network. This information is stored in a data source located on the server and replicated to other DCs on the network. The information pertaining to Active Directory is organized into the schema, domain, and configuration partitions, and can also have additional information for programs stored in the application partition. This data can be accessed over the network using LDAP.

To identify objects within the directory structure, Active Directory supports a variety of different naming schemes. These include the Domain Name System (DNS), user principal name (UPN), Universal Naming Convention (UNC), Uniform Resource Locator (URL) and Lightweight Directory Access Protocol Uniform Resource Locator (LDAP URL). Distinguished names (DNs), relative distinguished names (RDNs) and canonical names, based on X.500 specifications, are also used to identify objects.

A variety of objects build the directory's hierarchical structure, including users, computers, printers, other objects, and container objects that store them. In addition, other components are used to make up the physical and logical structure of Active Directory. Sites represent the physical structure of a network, while domains, trees, and forests represent the logical structure. Together, they are the building blocks that make up Active Directory.

A primary administrative tool for managing Windows Server 2003 and Active Directory is the Microsoft Management Console (MMC). Using this tool, you can load snap-ins that are used to administer different aspects of Windows Server 2003 and Active Directory. Three snap-ins are predominantly used to manage Active Directory: Active Directory Users and Computers, Active Directory Domains and Trusts, and Active Directory Sites and Services. In addition to these graphical tools, new command-line tools can be used to perform administrative tasks.

Active Directory also provides mechanisms for access control and authentication. Permissions can be applied to objects to control how they are used, while security descriptors, object inheritance, and authentication are used to determine a user's access and the permissions set on objects. Authentication methods that are supported include Kerberos, X.509 certificates, LDAP over SSL, and PKI. Through these methods, Windows Server 2003 and Active Directory are secured from unauthorized access.

Windows Server 2003 provides a number of new features and tools. For some of these to be available, the functional level of the domain and/or forest must be raised first. The functional level is similar to the domain modes used in Windows Server 2000, where backward-compatible features become deactivated and new features that older operating systems can't use become available as you raise the level.

A good understanding of the purpose and function of directory services and the infrastructure and topology of Active Directory are key elements in getting the most out of this powerful database. In this chapter, we provided the overview that is necessary to fully understanding the more specific topics covered in the rest of the book.

Exam Objectives Fast Track

Introducing Directory Services

- ☑ The Active Directory data store is a database of all directory information, and is also referred to as the *directory*. It is a file called NTDS.DIT, and is located in the NTDS folder in the systemroot.
- ☑ When Active Directory is installed, three partitions exist on each DC: the domain partition, the configuration partition, and the schema partition. There can also be one or more application partitions.
- ☑ Active Directory uses LDAP for communications between clients and directory servers. LDAP is a light version of the X.500 Directory Access Protocol (DAP).

Understanding How Active Directory Works

- ☑ Domains are logical groupings of network elements, consisting of computers, users, printers, and other objects making up the network.
- ☑ Active Directory allows multiple domains to be connected together in a hierarchy called a *domain tree*, consisting of parent and child domains.
- ☑ Active Directory has two forestwide master roles, and two domainwide master roles that store master copies of information. The Schema Master and Domain Naming Master roles are unique to one DC per forest, while the RID Master, PDC Emulator, and Infrastructure Master roles are all unique to one DC per domain.

Using Active Directory Administrative Tools

- ☑ Active Directory Users and Computers allows you to administer user and computer accounts, groups, printers, organizational units (OUs), contacts, and other objects stored in Active Directory. Using this tool, you can create, delete, modify, move, organize, and set permissions on these objects.
- ☑ Active Directory Domains and Trusts is used to manage domains and the trust relationships between them. Using this tool, you can create, modify, and delete trust relationships between domains, set the suffix UPNs, and raise domain and forest functional levels.
- ☑ The Active Directory Sites and Services tool is used to create and manage sites, and control how the directory is replicated within a site and between sites. Using this tool, you can specify connections between sites, and how they are to be used for replication.

Implementing Active Directory Security and Access Control

- ☑ Active Directory divides the permissions you can apply to objects into two categories: standard and special. Standard permissions are commonly applied to objects, while special permissions provide additional access control.
- ☑ Kerberos version 5 is an industry-standard security protocol that's used by Windows Server 2003 as the default authentication service. It is used to handle authentication in Windows Server 2003 trust relationships, and is the primary security protocol for authentication within domains.
- ☑ Public Key Infrastructure (PKI) is a method of authentication that uses private and public keys to provide authentication and encryption. For data confidentiality, the public key is available to the public and is used to encrypt session keys and data, while the private key is only known to the person for whom it is created, and is used for decryption. For authentication, the private key is used for encryption, and the public key is used for decryption.

What's New in Windows Server 2003 Active Directory?

- ☑ Domain functional levels can be raised to enable additional features in Active Directory. There are four different levels of domain functionality: Windows 2000 mixed, Windows 2000 native, Windows 2003 interim, and Windows 2003.
- ☑ Forest functional levels can also be raised to enable additional features in Active Directory. There are three different levels of forest functionality: Windows 2000, Windows 2003 interim, and Windows 2003.
- ☑ Windows Server 2003 provides a number of command-line utilities that allow administrators and users to manage and interact with Active Directory.

Exam Objectives

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

Q: Which editions of Windows Server 2003 can be used as DCs?

A: Active Directory can be installed on Windows Server 2003 Standard Edition, Windows Server 2003 Enterprise Edition, and Windows Server 2003 Datacenter Edition. When Active Directory is installed on any of these editions, it will serve as a DC. Active Directory cannot be installed on Windows Server 2003 Web Edition.

Q: How do I install Active Directory on a Windows Server 2003 member server, to make it become a DC?

A: Use **DCPROMO**. DCPROMO invokes the Active Directory Installation Wizard, and can be used to promote a member server to a DC. You can run it by clicking **Start | Run** and typing **dcpromo**, or you can use the **Configure Your Server** wizard to start the Active Directory installation.

Q: Why do Windows Server 2003 DCs use NetBIOS names when other naming schemes are used?

A: NetBIOS names are used to provide backward support. NetBIOS names are used by pre-Windows 2000 servers and clients, and allow users of those operating systems to log on to Windows Server 2003 domains.

Q: I am creating a new Windows Server 2003 network, and have just installed the first DC on the network. What must I do to create my first site, forest, and domain?

A: Nothing. When a DC is installed on the network, the first domain, forest, and site are automatically created. Additional domains, forests, and sites can be created as needed, just as additional DCs can be added. However, the first domain, forest, and site are created based on information you provided when you installed Active Directory.

Q: I want to set permissions on objects in Active Directory, so that unauthorized access to these objects is prevented. What snap-in do I use?

A: The **Active Directory Users and Computers** snap-in for the Microsoft Management Console. Using this tool, you can modify permissions and control access. This snap-in is already preconfigured in a console that you can access via **Start | Programs | Administrative Tools**.

Q: I want to make security changes to a user account, but when I bring up the permissions using the **Active Directory Users and Computers** snap-in for the MMC, the **Security** tab doesn't appear.

A: The **Security** tab is hidden in the **Properties** dialog box, unless the **Advanced Features** menu item is selected on the **View** menu first. After this is done, the **Security** tab will appear when you bring up the properties for an object.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Introducing Directory Services

1. An employee has retired from the company, and you have just disabled his account so no one can log on to the domain as this user. When this change is made, where will it be stored in the directory?
 - A. Domain partition
 - B. Configuration partition
 - C. Schema partition
 - D. Application partition
2. Your company's employees are represented by two unions. Management has a union that represents the managers' interests, while others in the company belong to another union. Each union requires that dues be deducted from paychecks to pay for their representation. The Finance department has requested that a field be added to each user account, so that a code can be entered on the account to show which union each employee belongs to. They have asked you to create this field. When this new attribute has been added to user objects, where will it be stored in the directory?

- A. Domain partition
 - B. Configuration partition
 - C. Schema partition
 - D. Application partition
3. You perform a search of Active Directory over the network, in search of an object stored in the directory. In performing this search, what protocol will be used?
- A. IPX/SPX
 - B. Directory Access Protocol (DAP)
 - C. Lightweight Directory Access Protocol (LDAP)
 - D. X.500
4. A user has the username JohnB. He wants to access a Microsoft Access database called db.mdb that's located on a DC called syngress.com, in a directory called DB. Using the URL, what will this user enter into a browser to access the database?
- A. JohnB@syngress.com
 - B. syngress.com
 - C. http://syngress.com/DB/db.mdb
 - D. \\syngress.com\DB\db.mdb
5. A user with the username of JaneD works in the Sales department. Her account is located in the syngress.com domain. Based on this information, what canonical name would be used to identify this object in Active Directory?
- A. CN=JaneD
 - B. /CN=JaneD /OU=Sales /DC=syngress /DC=com
 - C. OU=Sales
 - D. /syngress.com/Sales/JaneD

Understanding How Active Directory Works

6. You are making changes to object classes and attributes used in Active Directory. On which of the following DCs will you make these changes?
- A. Schema Master
 - B. RID Master
 - C. Infrastructure Master
 - D. PDC Emulator

7. Your network consists of two forests, with two domains in one forest and three domains in the other. Based on this information, how many of the following master roles will be in the forests and domains?
 - A. There will be five Schema Masters, Domain Naming Masters, RID Masters, PDC Emulators, and Infrastructure Masters.
 - B. There will be two Schema Masters, Domain Naming Masters, RID Masters, PDC Emulators, and Infrastructure Masters.
 - C. There will be five Schema Masters and Domain Naming Masters, and two RID Masters, PDC Emulators, and Infrastructure Masters.
 - D. There will be two Schema Masters and Domain Naming Masters, and five RID Masters, PDC Emulators, and Infrastructure Masters.

8. A user recently changed her last name, and you make changes to the user object in the directory to reflect this. Just before the change, inter-site replication has taken place using the default schedule. Just after the change, a link between the DC on which the changes were made and the DC in the other site fails. It will be another hour until the link is back up again. There are four DCs in each site. Which of the following will occur?
 - A. Replication between the DCs will occur normally, because at least two connections to each DC are created by the Knowledge Consistency Checker (KCC). Because one has failed, the other connection will be used.
 - B. Replication between the DCs won't occur. After 15 seconds, a notification of the change will be sent out, and replication partners will then request updated data.
 - C. Replication will occur normally, because the information won't be replicated until three hours after the last replication.
 - D. Another link will be used to replicate the data, based on the information gathered by the topology generator.

Using Active Directory Administrative Tools

9. You are using the Microsoft Management Console (MMC) to administer objects in Active Directory. You decide to view information about a DC. Which of the following snap-ins will you use to view this information?
 - A. Active Directory Users and Computers
 - B. Active Directory Domains and Trusts
 - C. Active Directory Sites and Services
 - D. Dcgpofix

10. Your company has merged with another company that uses UNIX machines as their servers. Users in your Windows Server 2003 domain need to access information on these UNIX machines, but you don't want to have information accessed by clients outside your domain. Which of the following types of trusts will you create to make it possible to share information in this way?
- A. One-way transitive forest trust
 - B. Two-way transitive realm trust
 - C. One-way nontransitive realm trust
 - D. External trust

Implementing Active Directory Security and Access Control

11. There is a concern about someone accessing objects in Active Directory using someone else's account. You want to audit an object to view how users and groups are accessing it. In setting auditing on this object, where will information on what to track be stored?
- A. Discretionary access control list
 - B. Security access control list
 - C. X.509
 - D. Auditing isn't provided on objects
12. You are configuring permissions on Active Directory so that managers can modify the user objects in the OU representing the department each manager is in charge of. In configuring these permissions, you also want each manager to have the ability to create new OUs within the OU representing his or her department. You want to give the most restrictive permissions to achieve these tasks. What permissions will you give these managers? (Choose all that apply.)
- A. Read
 - B. Write
 - C. Create All Child Objects
 - D. Delete All Child Objects
13. You have set permissions on a parent container, and want to prevent these permissions from being applied to a child container within it. How will you achieve this?

- A. In **Active Directory Users and Computers**, open the properties of the parent OU, and select the **Security** tab. Click the **Advanced** button, and when the dialog box appears, ensure that the **Allow Inheritable Permissions From Parent To Propagate To This Object** check box is checked.
- B. In **Active Directory Users and Computers**, open the properties of the parent OU, and select the **Security** tab. Click the **Advanced** button, and when the dialog box appears, ensure that the **Allow Inheritable Permissions From Parent To Propagate To This Object** check box is cleared.
- C. In **Active Directory Sites and Services**, open the properties of the parent OU, and select the **Security** tab. Click the **Advanced** button, and when the dialog box appears, ensure that the **Allow Inheritable Permissions From Parent To Propagate To This Object** check box is cleared.
- D. The objective cannot be achieved. Permissions will always be inherited by child objects. You must move the OU so it is at the same level in the hierarchy as the parent container.

What's New in Windows Server 2003 Active Directory?

- 14. You are upgrading your existing network to use Windows Server 2003. The network has Windows NT 4.0 domain controllers and the Windows Server 2003 server you're adding to the domain. After adding the first Windows Server 2003 DC to the network, you want to raise the domain functional levels to the highest level available for your network. To which level will you raise the domain?
 - A. Windows 2000 mixed
 - B. Windows 2000 native
 - C. Windows 2003 interim
 - D. Windows 2003
- 15. You are upgrading your domain to use a mix of Windows 2000 and Windows Server 2003 DCs. After installing the first Windows Server 2003 DC on the domain, you want to raise the domain functional level to the highest level possible. Which of the following will you choose?
 - A. Windows 2000 mixed
 - B. Windows 2000 native
 - C. Windows 2003 interim
 - D. Windows 2003

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|-------------|--------------------|
| 1. A | 9. A |
| 2. C | 10. C |
| 3. C | 11. B |
| 4. C | 12. A, B, C |
| 5. D | 13. B |
| 6. A | 14. C |
| 7. D | 15. B |
| 8. C | |

Chapter 2

MCSA/MCSE 70-294

Working with User, Group, and Computer Accounts

Exam Objectives in this Chapter:

- 3 Planning and Implementing User, Computer, and Group Strategies
 - ☑ Summary of Exam Objectives
 - ☑ Exam Objectives Fast Track
 - ☑ Exam Objectives Frequently Asked Questions
 - ☑ Self Test
 - ☑ Self Test Quick Answer Key


Introduction

An important part of the network administrator's job involves management of the network's users and computers. Windows Server 2003 assigns *accounts* to both users and computers for security and management purposes. User accounts can be further managed by placing them in groups so that tasks—such as assigning permissions—can be applied to an entire group of users simultaneously rather than having to do so for each individual user account.

This chapter introduces you to the concept of security principles—users, groups, and computers—and the Security Identifiers (SIDs) that are used to represent them. You'll learn about the conventions and limitations for naming these objects.

We show you how to work with Active Directory user accounts, including the built-in accounts and those you create. You'll also learn to work with group accounts, and you'll learn about group types and scopes. You'll learn to work with computer accounts, and how to manage multiple accounts. We'll show you how to implement user principal name (UPN) suffixes, and we'll discuss how to move objects within Active Directory.

You'll learn to use the built-in tools—both graphical and command line—to perform the common administrative tasks associated with the management of users, groups, and computers; and the exercises will walk you through the steps of creating and managing all three types of accounts.



Understanding Active Directory Security Principal Accounts

Active Directory is made up of a wide variety of different directory service objects. Among these objects are security principal accounts, which consist of the following:

- User accounts
- Computer accounts
- Groups

Security principal accounts are used in authentication and access control, and provide a means to manage what can be accessed on the network. Based on the security settings associated with a security principal account, you can control whether a user, group, or computer has access to Active Directory, printer, and file system objects, as well as domain controllers (DCs), member servers, client computers, applications, and other elements of the network. They are a major factor in keeping your network protected and controlling what users and computers are authorized to access.

Because security principals represent people, services, computers, and others who individually and collectively access the directory, there are a number of different ways to identify them. Some names might be familiar to you, such as the username you use to log on to

Windows Server 2003 or the name identifying your computer on the network, while others are used for specific technologies (such as Web browsers) or backward compatibility to older operating systems. As we'll see later, the different names provide friendly and unique methods of identifying users, groups, and computers.

Security Principals and Security Identifiers

Security principals get their name because they are Active Directory objects that are assigned SIDs when they are created. The SID is used to control access to resources and by internal processes to identify security principals. Because each SID is unique, unless security is breached, there is no way for accounts to mistakenly gain access to restricted resources when the system is properly configured by an administrator.

SIDs are able to remain unique because of the way they are issued. In each domain, there is a DC that acts as a Relative ID (RID) Master. The RID Master is responsible for generating relative identifiers, which are used in creating SIDs. The SID is a number that contains a domain security identifier and relative identifier. The domain ID is the same for all objects in the domain, but the relative identifier is unique. A pool of these numbers is issued to each DC within the domain, so they can be assigned to security principals that are created on the DC. When 80 percent of the numbers in the pool have been assigned to objects, DCs will then request a new pool from the RID Master.

SIDs are used because unlike the names associated with objects, SIDs don't change. When the object is created, a unique alphanumeric value is associated with it, and this stays with the object until it is deleted. Such things as changing the object's name or other attributes don't affect the SID. For example, if you created a user account called "Jane Doe," a SID would also be generated for that account. If you later changed the account's name to "Jane Smith," the object's name would be altered, but the SID would remain the same. If the name were used to determine access, it would appear that a completely different user was attempting to access resources in the domain. Because the SID is used to determine access, the user's identity remains constant, and any access the user has will be unaffected.




TEST DAY TIP

Don't forget that the only accounts that are security principals are user accounts, computer accounts, and group accounts. These are the only objects that are given SIDs at the time they are created. If one of these accounts were deleted, and then recreated with the same information, it would be given a new SID and appear to internal processes and access control lists (ACLs) as a completely different account.

To better understand a SID, you could compare it to an employee ID. When an employer issues you one of these numbers, it doesn't change as situations change in your life. You could change your name, address, office location, title in the company, or other identifying factors, but this number will always be your number. The SID is used in the same manner.

As shown in Figure 2.1, the SID is used as part of the authentication process. When a user logs on to a domain, the Local Security Authority (LSA) is used to authenticate to Active Directory, and create an access token. The access token is used for controlling a user's access to resources, and contains the user's logon name and SID, the names and SIDs for any groups the user is a member of, and privileges assigned to the user. The token is created each time the user logs on, and holds all of the information needed for access control.

Figure 2.1 How Security Identifiers Are Used in Access Control



When a user attempts to access a resource, Windows Server 2003 compares the SID with the resource's security descriptor. A security descriptor contains two components, the discretionary access control list (DACL) and the system access control list (SACL). An ACL contains access control entries (ACEs), which are used to control or monitor access to a resource. An ACE determines whether a user associated with a particular SID is to be allowed or denied access, or whether the user is to be audited.

The SACL is used for auditing access to a resource. An ACE in a SACL contains information on whether logging should be generated on attempts to access a resource. This logging can be generated when a specified user or group attempts to access a resource and is successful, fails, or both.

The DACL is used for a different purpose. DACLs determine whether a security principal is granted or denied access to a resource. The DACL catalogs who has access to the

resource and what level of access they have. When a user tries to access an object, the user's SID is compared to entries in the DACL. If the user's SID or the SID of a group he or she belongs to matches an entry in the DACL, that user can be either explicitly permitted or denied access to use the resource.

When users access a resource, a process begins to determine the level of access they have, based on the permissions they have to a resource. The system will first determine whether a DACL exists on a resource. If it does not, then there is no access control for the resource, so access is granted. If a DACL does exist, then the system will go through the ACEs until one or more matches are found, or until it finds an ACE that specifies that access is denied.



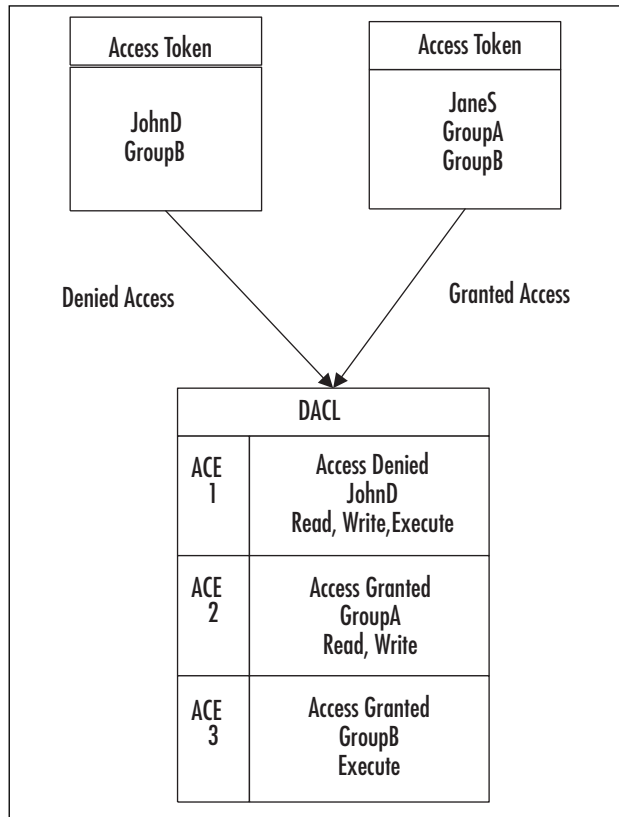
NOTE

On a Windows Server 2003 computer using the NTFS file system, all file and folder objects will have DACLs that control access to local resources and allow auditing. When using the FAT file system, they will not.

When a security principal attempts to access a resource that is protected by a DACL, each ACE in the DACL is analyzed in sequence to determine if access should be allowed or denied. As shown in Figure 2.2, the SID of the user and any groups he or she belongs to is compared to the ACEs in the DACL. Windows Server 2003 will look at each ACE until one of the following occurs:

- An entry is found that explicitly denies access to the resource.
- One or more entries are found that explicitly grants access to the resource.
- The entire DACL is searched but no ACE is found that explicitly grants or denies access. Since no entry is found, the security principal is implicitly denied access.

In Figure 2.2, one user is granted access while the other is denied access. When the SIDs associated with the access token of the JaneS user is compared with the entries in the DACL, the system will find that she is a member of GroupA (which has Read and Write access) and GroupB (which has Execute access). Because of her membership in these groups, she will be granted Read, Write, and Execute permissions for the resource. When the SIDs associated with the access token of the JohnD user is compared with the DACL, the system will find that he is a member of GroupB, which has Execute permission for the resource. However, there is also an ACE that explicitly denies JohnD Read, Write, and Execute access. When the user's SID is compared with this entry, he will be denied access. In general, the most permissive combination of permissions will be allowed when a user accesses a resource, unless an explicit deny is assigned.

Figure 2.2 The User's SID Is Compared with the DACL's ACE to Determine Access

In addition to containing the user's SID and the SIDs of any groups the user is a member of, the access token might also contain additional SIDs that result from group membership the operating system assigns dynamically or other special logon scenarios. These SIDs are common to installations of Windows Server 2003 in a stand-alone and/or Active Directory environment, and for this reason are referred to as *well-known security identifiers*. Table 2.1 lists the different types of well-known SIDs.

Table 2.1 Well-Known Security Identifiers

Type of SID	SID	Description
Anonymous Logon	S-1-5-7	Used when a user logs on without supplying a username and password.
Authenticated Users	S-1-5-11	Used when users or computers have been authenticated with individual accounts. The exception is when someone logs on with the Guest account, as this isn't considered an authenticated user.

Continued

Table 2.1 Well-Known Security Identifiers

Type of SID	SID	Description
Batch	S-1-5-3	Used when users log on through a batch queue facility, such as when scheduled tasks are run under a specific account.
Creator Owner	S-1-3-0	Used as a placeholder in inheritable ACEs. When an ACE is inherited, this SID is replaced with the SID of the current owner of the object.
Creator Group	S-1-3-1	Used as a placeholder in inheritable ACEs. When an ACE is inherited, this SID is replaced with the SID of the primary group to which the object's current owner belongs.
Dialup	S-1-5-1	Used when the user logs on to the system using a dial-up connection.
Everyone	S-1-1-0	Used to specify the Everyone group. In Windows Server 2003, this includes authenticated users and the Guest account. Earlier versions of Windows include these accounts and anonymous logons.
Interactive	S-1-5-4	Used for users who are logged on to the local machine (as opposed to connecting via the network) or connected through Terminal Services.
Local System	S-1-5-18	Used for a service account that's run by Windows.
Network	S-1-5-2	Used when a user logs on through a network connection. Because of the methods in which interactive users log on to the system, this type of SID isn't used for their access tokens.
Other Organization	S-1-5-1000	Used to determine whether users from other domains or forests are permitted to authenticate to services.
Self/Principal Self	S-1-5-10	Used as a placeholder in ACEs. When permissions are granted to Principal Self, they are given access to the security principal represented by the object. The SID acting as a placeholder is replaced during access checks with the SID for the user, group, or computer represented by the object.
Service	S-1-5-6	Used for security principals that log on as a service.

Continued

Table 2.1 Well-Known Security Identifiers

Type of SID	SID	Description
Terminal Server Users	S-1-5-13	Used for users that log on to a Terminal Services server running in Terminal Services version 4.0 application compatibility mode.
This Organization	S-1-5-15	Used to add data to the authentication information of the user who's logged on by the authentication server. This is used if the Other Organization SID isn't used.

Tools to View and Manage Security Identifiers

Windows Server 2003 provides command-line utilities that can be used to view and manage SIDs. Because of the way in which SIDs are handled in a domain, it is rare to ever need this capability. However, should the need arise, it is useful to know these tools exist.

As we saw in Chapter 1, "Active Directory Infrastructure Overview," the WHOAMI tool allows you to display information about the user who is currently logged on. By using the */ALL* parameter, you can display all of the user's access token information, including information on the username, groups, associated SIDs, and privileges, for the user who is currently logged on. The syntax for viewing this information is:

```
WHOAMI /ALL
```

The NTDSUTIL tool is another utility we discussed in the previous chapter. This tool is a general-purpose tool that provides a variety of functionality. Using various menus, this text-based tool allows you to navigate through functions that (among other things) allow you to manage SIDs in rare instances where duplicate SIDs exist.



EXAM WARNING

The functionality of WHOAMI and NTDSUTIL are quite different. Although both are tools for working with SIDs, WHOAMI is only used to view SID information, while NTDSUTIL is a tool for managing duplicate SIDs that might exist in Active Directory.

To avoid potential conflicts of DCs issuing the same SID to an object, only one RID Master exists in a domain. While this will generally ensure that SIDs are unique within a domain, there is the remote possibility that duplicate SIDs can be still be issued. An administrator has the ability to assign the role of RID Master to another DC in the domain. If the RID Master was temporarily unavailable and an administrator seized the RID Master role on another DC, there would be two RID Masters when the original RID Master became available. This generally isn't a problem, because the RID Masters will recognize that there are two of them in the domain during the next replication cycle, and one of the

two will relinquish its role. However, in this situation, if two DCs requested a new pool of relative IDs while both RID Masters were active, then two sets of identical numbers could be issued to the DCs.

Duplicate SIDs can be found and dealt with by using the Security Account Management menu of NTDSUTIL. By typing **security account management** from the NTDSUTIL prompt, commands shown in Table 2.2 can be accessed. Using these commands, you can check the domain for any objects with duplicate SIDs and delete them to resolve the issue.

Table 2.2 Commands Available in NTDSUTIL from the Security Account Management Menu

Command	Description
Check duplicate SID	Checks the local SAM or domain database for objects with duplicate SIDs.
Cleanup duplicate SID	Checks the local SAM or domain database for objects with duplicate SIDs and deletes them.
Connect to server %s	Specifies the server or DC to connect to. %s is a variable denoting the server or DC to connect to.
Log file %s	Specifies where to create a log file. %s is a variable denoting a location. If this parameter isn't used, events will be logged to a file named dupsid.log that will be placed in the root folder of the profile for the user who executes the command.
Quit	Returns to previous menu. When the main menu is displayed again, this command will exit you from NTDSUTIL.
? or help	Displays help

Now that we've discussed these tools and the tasks you can perform with them that relate to SIDs, let's put this knowledge to practice in Exercise 2.01

EXERCISE 2.01

VIEWING AND MANAGING SECURITY IDENTIFIERS

1. From the Windows **Start** menu, click the **Command Prompt** menu item.
2. When the command prompt opens, type **WHOAMI /ALL** and then press the **Enter** key. A list of information about the account you're currently logged on with will appear on the screen. In this listing, information on SIDs associated with your account will also be displayed.

3. When you have finished viewing the list of information, type **NTDSUTIL** at the prompt, and then press the **Enter** key.
 4. At the prompt, type **security account management**, and then press the **Enter** key.
 5. At the Security Account Maintenance prompt, type **connect to server localhost** to establish a connection to your local computer. You can also use an IP address, NetBIOS, or DNS name to make this connection to your local computer or a remote one.
 6. At the Security Account Maintenance prompt, type one of the following:
 - To check for duplicate SIDs but not remove them, type **check duplicate sid**, and then press the **Enter** key.
 - To both check for duplicate SIDs and remove any that are found, type **cleanup duplicate sid** and then press the **Enter** key.
 7. Use Notepad to open the `dupsid.log` file in the root folder of your user profile to determine if any duplicates were found or removed.
 8. At the command prompt, type **Quit** and then press the **Enter** key to return to the previous menu. Type **Quit** and then press **Enter** again to exit NTDSUTIL.
-

Naming Conventions and Limitations

In looking at the relationship between security principals and SIDs, it becomes apparent that it would be difficult to use SIDs as the sole method of identifying an account. While SIDs uniquely identify users, computers, and groups, trying to remember the SID of users and computers you commonly access through the directory would be almost impossible. For this reason, various naming conventions are used to distinguish objects in Active Directory.


Every object in Active Directory has a name to make it easier to find and use. When naming these objects, each object must be unique so that multiple objects can't be confused with one another. This doesn't mean that two objects can't be named the same within Active Directory; just that two objects with the same name can't exist within the same administrative boundary under certain circumstances.

In Windows Server 2003, a number of different boundaries exist to logically group objects and manage them. For example, domains provide a boundary that allows two computers with the same name to exist in different domains. If you had a computer named `server1` in one domain and another computer with this same name in a different domain,

there would be no conflict. Similarly, UPNs in Windows Server 2003 must be unique in each forest, and Pre-Windows 2000 logon names must be unique in each domain. As you can see, there are a number of complex naming rules that you'll need to familiarize yourself with for the exam. Even if you're new to Windows Server 2003 networks, you're probably familiar with how domains are named on the Internet, which is the largest DNS namespace in the world. DNS is used to resolve domain names to IP addresses (and vice versa), and also provides a standard method of naming domain structures. Active Directory domains follow the same naming conventions as DNS. For example, `syngress.com` can be both an Active Directory domain name and a DNS domain name. Active Directory and DNS are tightly integrated, which is why the domain names in Active Directory require corresponding DNS namespaces.

As shown in Figure 2.3, an Active Directory namespace is arranged in a hierarchy. In the hierarchy, objects branch off from a *root* object, and other objects can branch off from these. The object that has branched off is called a child object, and the object it branches from is called a parent. If a domain branches off from another domain, it is called a *child domain*. The domain it branches off from is called the *parent domain*. The child domain retains information about the domain above it, and includes the full name of the parent. This contiguous namespace provides a way of conveying how the hierarchical structure of the domains is designed in a network.

Figure 2.3 Hierarchical Structure of an Active Directory Namespace



In looking at Figure 2.3, you can see that `syngress.com` is the parent domain, and that it has two child domains (`sales.syngress.com` and `dev.syngress.com`). Because a child domain can also have child domains beneath it, we also see that `sales.syngress.com` has its own child domain called `marketing.sales.syngress.com`.

In addition to providing a method to uniquely identify domains within a namespace, individual computers are also provided with unique names that relate to the hierarchy. Each computer in a domain has a host name that is unique to the domain. Because of the boundary provided by domains in the namespace, a computer with the host name `server1` could exist in each domain. To show the place of these computers in a DNS domain structure, fully qualified domain names (FQDNs) are used, which combine the host name with the domain name. For example, a computer with the host name `server1` in the domain `syngress.com` would have the FQDN of `server1.syngress.com`. In Active Directory, the computer name is also appended to the name of the domain of which it is a member. This combination can be up to 255 characters in length. When creating computer accounts in Active Directory, the name of the computer can only consist of letters (a to z, A to Z), numbers (0 to 9), and hyphens (-). The name cannot consist of all numbers. Because pre-Windows 2000 computers didn't use a DNS-like naming scheme, Active Directory also allows NetBIOS names to be used for backward compatibility. In older network operating systems such as Windows NT, a hierarchical structure of domains wasn't used. Computers used a flat namespace in which all computers were given NetBIOS names that were unique to the domain. In a single domain environment, these were single word names that allowed clients and services to identify computers on the network, and were used by users to log on to systems and domains. In a multiple domain environment, the domain name was typically appended to the computer name using the backslash (\) character; for example, `domain\computername`. Similarly, logons between domains took the form of: `domain\username`. By providing support for NetBIOS names and pre-Windows 2000 domain naming schemes, older clients can access Windows Server 2003 computers and domains.



NOTE

Backward compatibility to previous versions of Windows operating systems is of major importance to a Microsoft network. Although Microsoft will try to recommend that newer operating systems are used on client machines and servers, a network might consist of a variety of Windows versions and editions. Even in situations where clients and servers are all being upgraded to the latest versions, older versions will reside on the network during the upgrade process. To accommodate these environments, Windows Server 2003 supports a number of schemes and technologies for backward compatibility.

The NetBIOS name of a computer can be up to 15 characters in length. When installing a new Windows Server 2003 server on the network, the operating system will

suggest a NetBIOS name that is based on the first 15 bytes of the relative distinguished name (RDN). As we'll see later in this section, the RDN is another method of identifying security principals in Active Directory. In an FQDN, the host name (which always appears left of the leftmost dot in the FQDN) is also the RDN.

In addition to names used by computers and domains, user accounts also have distinct methods of being named. User accounts have a UPN that can be used to log on from Windows 2000, XP Professional, and Server 2003 machines. They also have a backward-compatible login name known as the pre-Windows 2000 name. During the process of creating a new user account, Active Directory will suggest a pre-Windows 2000 name that is based on the first 20 characters of the UPN that you type in. Either name can be changed at any time. The pre-Windows 2000 logon name is limited to 20 characters. Although the pre-Windows 2000 name can still be used to log on to domains on newer operating systems, the UPN logon name is preferred when logging on from Windows 2000 or later. UPNs consist of a logon account name and a UPN suffix. By default, the UPN suffix is the domain name that contains the user account. For example, if the user account was located in a domain called syngress.com, then this domain name would also be the UPN suffix. The logon name and UPN suffix is connected together using the at (@) sign. This makes it appear like an Internet e-mail address (username@domain).

While the default UPN suffix for a user account is the domain containing the account, other UPN suffixes can also be used. To make it simpler to log on to the network, Active Directory allows you to implement a single UPN suffix for all users. In such a case, you could have all users have the forest root domain name as their UPN suffix, or even create alternate UPN suffixes. For example, let's say that users logged on to an internal Active Directory domain name called syngress.net, and knew that the external DNS name used by the Web site was syngress.com. This could be confusing to users, because the public Web site is popular and well known. To make it less confusing, you could implement an alternate UPN suffix so that all users logged on to the UPN suffix syngress.com. In using alternate UPN suffixes, you aren't limited to using valid DNS or Active Directory domain names. You can use any UPN suffix you choose.

Because Windows Server 2003 allows users to have the same UPN suffix regardless of which domain contains their account, UPN logon names must be unique within a forest. In Windows Server 2003, Active Directory won't permit two UPNs with the same name to exist in the same forest.



NOTE

The concept of UPNs can be confusing initially. Keep in mind that a UPN has no relation whatsoever to a DNS or Active Directory domain name. Despite the fact that it can have the same name as a valid domain within your forest, it is just a label or alias that is used as part of a user account name. UPNs are stored in the Global Catalog (GC), so a GC server must be available to log on using them. Remember, it is the SID associated with your user account that determines what domain it belongs to; it can have any valid name associated with it. The purpose of

using a UPN is to allow users to continue using the same logon name, even if the domain their account is associated with changes. Because the UPN is not mapped to a specific domain, it does not need to change when you move a user account from one domain to another within your forest.

In creating a name that's used by the user to log on to Active Directory, certain limitations exist in what you can use in the logon name. The name cannot contain all spaces, or must not contain any leading or trailing spaces, or any of the following characters: " / \ [] : ; | = , + * ? < >



NOTE

Several of Microsoft's instructional materials state that the logon portion of the UPN name cannot exceed 20 characters, despite the fact that the text box allows many more to be typed in. This is incorrect. If you test it, you'll learn that you can exceed the 20 character limit. The pre-Windows 2000 logon name does recognize this limit.

In addition, many of the restricted characters listed in the preceding paragraph will not raise an error during account creation if typed in as part of the User logon name. However, when you attempt to log on you will receive an error. Strangely, if you type these restricted characters in as part of the pre-Windows 2000 logon name, it will notify you that they will be changed to the underscore character (an allowed character) and proceed with account creation.

Groups also have certain requirements that must be adhered to when they're created. As we'll see later in this chapter, groups allow you to add users as members. This collection of user accounts can then be managed as a unit, making it easy to assign many users permission to access a resource with a single assignment.

When creating groups in Active Directory, the name cannot be longer than 64 characters in length. In addition, the name must not contain leading or trailing spaces, trailing periods, or any of the following characters: / [] : | = ? * , + " \ < > ;

Security principals also make use of three other types of naming conventions, which are common to all objects in Active Directory:

- RDN
- Distinguished name (DN)
- Canonical name

Through these naming conventions, security principals can be located in the directory using the Lightweight Directory Access Protocol (LDAP). Each of these is used to identify

objects in Active Directory, and provide methods of locating objects within the directory, regardless of how many layers of organizational units (OUs) it is stored under.

As we've discussed, every object is given a name when it is created, and this name is used to easily identify the object within the directory. As we'll see later in this chapter, the name of a user object generally corresponds to the user's name, while the computer's name can correspond to either the Active Directory object name you specify or the NetBIOS name. These names are RDNs.

RDNs refer to the name of the object in relation to where it is located in the directory. In other words, it doesn't show the path to where you can find the object in the directory structure. For example, a user object named John Smith would be a valid RDN. In looking at this name, however, you wouldn't know whether it was stored in the Users container or another location. The RDN doesn't show exactly where the object can be found, but refers only to the name of the object.

Because the RDN identifies the object within a container, this name must be unique within the container in which it is stored. In other words, you can't have two users named John Smith in the same OU. If two objects did exist with the same name, confusion could occur as to which object you really wanted to access.


To provide more specific information concerning the exact location of an object within the directory's hierarchy, a DN is used. The DN is used to show the path to an object. It says that the object is located in a particular domain, and possibly even within a specific OU. This path identifies the name of the object, and the hierarchy to the container in which it is stored. To provide this information, the DN uses the following notations:

- **CN** The common name of the object.
- **OU** An organizational unit that contains the object, or contains another OU in the hierarchical path to the OU that contains the object.
- **DC** A domain component that specifies a DNS name in the hierarchy to the object. Just as with OU objects, there may be multiple domain components in the hierarchical path to the object.

As shown in Figure 2.4, the DN uses these notations to provide a map to how you can find a single object within the structure of Active Directory. In Figure 2.4, we see that there are a number of users located in the Accounts Receivable OU, which resides in the Accounting OU of knightware.ca. If you wanted to find a user object named John Doe within this structure, you would use the following DN:

```
CN=John Doe,OU=Accounts Receivable,OU=Accounting,DC=knightware,DC=ca
```

In comparing the DN to Figure 2.4, you can see that it starts with the object, and works its way up to the highest level of the structure.

Figure 2.4 Distinguished Names Are Used to Show Location of Objects


As mentioned earlier, the CN within the DN is also the RDN of the object. The RDN uniquely identifies an object within a parent container, and in the case of our previous example, refers to the John Smith object. In this case, we could refer to the RDN as follows:

CN=John Smith

Another way of looking at this same information is with the canonical name. Canonical names provide the same information as the DN, but in reverse. Rather than beginning with the object and working its way up the highest level, the canonical name starts at the highest level and works its way down. In looking at the previous example, the DN “CN=John Doe, OU=Accounts Receivable, OU=Accounting, DC=knightware, DC=ca” would be translated to the following canonical name:

knightware.ca/Accounting/Accounts Receivable/John Doe

By providing different options for identifying objects and their location in the directory, there is greater versatility in the methods tools that you can use to access these objects. Before using DNs, RDNs, and canonical names, you need to give the object a unique name. As we’ll see in the sections that follow, the names for users, computers, and groups are specified when the security principal itself is created.



Working with Active Directory User Accounts

User accounts are used by people and services so they can be authenticated and access resources. Each user account contains information about the person or service that uses it, and provides a means to grant permissions, apply scripts, assign profiles, and control what actions the user can perform and what he or she can access. Through the account, a set of credentials is created that protects from unauthorized access.

In Windows Server 2003, two different types of user accounts can be created: local and domain-based user accounts. Local user accounts are used to control access to the computer on which you are working. They are created on Windows Server 2003 by using the Local Users and Groups snap-in, or the Users node under the Local Users and Groups node in the Computer Management utility. Once created, the account information is stored in a local database called the Security Accounts Manager (SAM). The account information only applies to the local computer, and isn't replicated to other machines within the domain. When a user logs on to the computer, Windows Server 2003 authenticates the user with this information, and either permits or denies access to the machine.

Domain accounts are created in Active Directory and are considerably different from local user accounts. Rather than storing information on the local machine, account information is stored in the directory and replicated to other DCs. As we discussed earlier in this chapter, when the user logs on to a DC, the account information is used to build an access token. This access token is used for the duration of time that the user is logged on to the network, and determines what the user is allowed to access on the network, and actions he or she can perform.

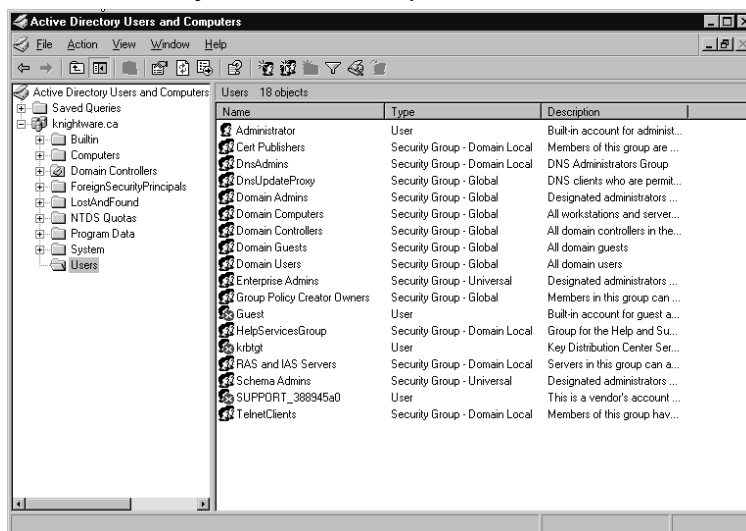


TEST DAY TIP

Don't confuse local accounts with domain accounts. Local accounts are stored on computers and only apply to the security of those machines. Domain accounts are stored in Active Directory, and security settings for the account can apply to accessing resources and services across the network.

Active Directory user accounts are created and managed using the Active Directory Users and Computers snap-in. As shown in the Figure 2.5, this snap-in provides a graphical user interface (GUI) that allows you to point-and-click through the various tasks related to administering user objects. The left pane of this tool is the console tree, which contains nodes representing your domain and the container objects within your domain such as OUs. Expanding the node of a domain displays the containers, which can be selected to view objects stored within them. These objects within the container are displayed in the right pane of the console.

Figure 2.5 Active Directory Users and Computers



As mentioned in Chapter 1, a number of containers are automatically created when Active Directory is first installed. Each stores different types of objects, some of which are used in managing users and computers on the network. These containers are:

- **Builtin** The default location for most domain local groups that are created during the installation of Windows Server 2003 and Active Directory. A few service-specific domain local groups, such as the DnsAdmins group, are created in the Users container.
- **Computers** The default location that is used to store computer objects for members of the domain. This container does not contain objects for Active Directory DCs.
- **Domain Controllers** The default location that is used to store Active Directory DC objects.
- **ForeignSecurityPrincipals** The Active Directory location used to store foreign SIDs for user accounts in external trusted domains.
- **Users** The default location for user accounts, global groups, and universal groups that are created during the installation of Active Directory. This container often contains additional domain local groups that are used by services such as RRAS and DNS.

In addition to these containers, others also exist that are hidden. You must enable Advanced Features to display these additional containers. To do this, select **View | Advanced Features**. Once enabled, the following containers can then be seen:

- **LostAndFound** Used to store objects whose parent containers no longer exist. If an object is created on one DC close to the time that its parent container is deleted on another DC (or if it is moved to a location that's missing after replication), the object is considered orphaned and is placed in this container.
- **System** Contains information about the domain, objects used by Active Directory, and the underlying Windows Server 2003 operating system. Unlike most of the other containers, the objects in this container generally cannot be modified by the administrator.

While these containers are created by Active Directory, objects can also be stored in OUs that are created by the administrator. By using OUs, you can arrange user accounts, computer accounts, and other objects into containers that reflect the department or location of these objects. For example, you could create an OU for a branch office, and then store accounts for users at that location within the OU. This makes it easier to delegate administrative control, and manage users using Group Policy.



NOTE

By looking carefully at Figure 2.5, you'll see a difference in the icons between the Users node and the Domain Controllers node. The Users node is a default container, while Domain Controllers is an OU. A key difference between a default container and an OU is that you cannot apply group policy to a default container. Therefore, if you create all of your users in the Users container, you will have to move them into an OU to apply group policy to them, or apply the group policy at the domain level and allow it to inherit down to the Users container. Administrators can create OUs, but they cannot create default containers.

Built-In Domain User Accounts

As we'll see later in this chapter, you can create user objects for accounts used by users and services within your organization. In addition to those you create, several user accounts are automatically created when Active Directory is first installed. These built-in accounts are stored in the Users container of Active Directory Users and Computers, and are named:

- Administrator
- Guest
- HelpAssistant
- SUPPORT_388945a0

While we'll discuss each of these accounts in the sections that follow, it is important to realize that each of the accounts created by Active Directory is assigned group memberships

and user rights that provide different levels of access. The rights and permissions they have vary, because each is designed to be used for a different purpose. You wouldn't want everyone in a domain to be able to view, modify, or delete anything they want on the network. This would be a major security issue, and accidents and malicious actions could have potentially devastating consequences.

Administrator

The Administrator account is the first account that's created when Active Directory is installed. As we saw in the previous chapter, when you use the Active Directory Installation Wizard and set up a new domain, this account is created to give you the access to perform domain configuration. Once created, it can be used to create and manage security principals and other objects, administer policies, assign permissions, and other tasks needed in the design and administration of Active Directory.

The Administrator account has the highest level of access of any default account created in Active Directory. It is a member of the Administrators, Domain Admins, Domain Users, Enterprise Admins, Group Policy Creator Owners, and Schema Admins groups. While we'll discuss groups later in this chapter, the default membership in these groups allows the Administrator account to have full control over Active Directory. Because these groups provide the security access attributed to this account, you can create additional accounts for people who perform administrative tasks and make them members of these groups to give them the same abilities as the Administrator account.

Due to the importance of the Administrator account, it cannot be deleted from Active Directory, or removed from the Administrators group. It can, however, be disabled or renamed to make it more difficult for unauthorized or malicious users to use this account by guessing its password.

Securing the Administrator Account

Anyone familiar with Windows network operating systems will recognize the Administrator account as one that's existed in previous versions. An account called Administrator has been used in Windows NT, Windows 2000, and Windows Server 2003. Because it is known, it can be a target for hackers, who want to access the account and obtain full control over the domain.

To protect your system, you should rename and/or disable the Administrator account after creating other accounts with administrator-level access. By renaming the account, others won't know what the account is called, and thereby won't know which account to access to obtain this high level of control. By disabling the account, they won't be able to use it at all.

To prevent unauthorized users from accessing accounts with administrative credentials, any accounts that are members of the groups that provide administrator access should use strong passwords. Strong passwords are more difficult

Continued

to guess, and use a combination of three or more of the following keyboard characters:

- Lowercase letters (a–z)
- Uppercase letters (A–Z)
- Numbers (0–9)
- Special characters ({}[].,<>:;'"?/|\`~!@#\$%^&*()_+ =)

While strong passwords can't be hacked using dictionary attacks that look for specific words and character combinations in the password, they can still be cracked using brute-force hacking programs, which try to determine the password by using all possible combinations of characters in a password. Brute-force hacks can take a considerable amount of time, but are commonly used to obtain unauthorized access to an account. To make it more difficult for hackers to access an account using a brute-force attack, you should use longer passwords that consist of a minimum of eight characters. In a Windows Server 2003 domain, strong passwords are enabled in Group Policy as a requirement for all accounts. This means that all passwords that are created after the installation of Active Directory must use at least one uppercase character, lowercase character, and numeric character.

Guest

The Guest account is another built-in account, but provides the lowest level of access. It is designed to be used by occasional users who need minimal access and don't want to log on with their own account, or users who don't have an account of their own in the domain. For example, you could give the Guest account permissions to view files on an intranet Web site, or have read-only access to unclassified material, such as forms new employees must fill out to apply for a job. People applying for a job could then access and fill out these forms even though they don't have their own accounts.

When Active Directory creates this account, it makes the Guest account a member of the Guests group and Domain Guests global group. Membership in these groups allows a person using this account to log on to the domain. Just as with other accounts, you can control what rights and permissions this account has, and add or remove this account from group memberships. This allows you to increase or decrease the level of access a person using this account has to resources in the domain.

Because it is better for users to have their own accounts when logging on to the domain, the Guest account is disabled by default. Having this account disabled prevents unauthorized persons from using this account to access the domain, and potentially use it to obtain additional levels of access. As we saw with the Administrator account, the Guest account can't be deleted, but can be renamed. Doing so will make it more difficult for hackers to take advantage of this account, which has also existed in previous versions of Windows.

HelpAssistant

The HelpAssistant account is automatically created in Active Directory when a Remote Assistance session is established. Remote Assistance allows a user to connect to a machine and assist them, such as by taking control of the computer remotely. For example, a person working Help Desk could take over a user's computer remotely, and show the user how to perform a particular task.

To prevent others from indiscriminately taking over a computer and performing tasks while a person is logged on, this connection is established with the permission of the person using the computer. In other words, the person using the computer must allow another person to take control. A person couldn't simply hijack another person's computer using Remote Assistance, and then begin accessing resources under their logged-on account.

The HelpAssistant account is used to establish the session between the person assisting a user and his or her machine, and provides limited access to the computer being accessed remotely. This account is managed by the Remote Desktop Help Session Manager service, and is deleted automatically when there are no pending Remote Assistance requests. Because it is removed when no longer needed, it doesn't always appear in the Users container of Active Directory Users and Computers.

SUPPORT_388945a0

The Support_388945a0 account is used by the Help and Support Service to provide interoperability with and allow access to signed scripts that are made available within Help and Support Services. An administrator can delegate the ability for a normal user to run these scripts from links in Help and Support Services. The scripts can be programmed to use the Support_388945a0 account instead of the logged-on user's credentials to perform administrative tasks on the local system that the user would not typically be allowed to perform without administrative-level access.

InterOrgPerson

Unlike the previous accounts we've discussed, InetOrgPerson accounts don't actually refer to an account named InetOrgPerson, but an object class used to create accounts. Object classes are sets of attributes used to determine what attributes an object might have when it is created. In Active Directory, the InetOrgPerson class is used to define user accounts. Because it exists as a type of user class, accounts created with this class are security principals.

InetOrgPerson accounts are used to represent users in non-Microsoft directory services. While Active Directory is the only directory service used by Windows 2000 and Windows Server 2003, it isn't the only directory service in existence. Other network operating systems, such as Novell NetWare, use their own implementations of a directory service, which aren't always compatible with Active Directory. The InetOrgPerson is used to assist applications written for other directories, or when migrating from these directory services to Active Directory.

As we'll see in the next section, InetOrgPerson accounts are created in the same way that user accounts are created. In **Active Directory Users and Computers**, you would select the container in which you want to create the new InetOrgPerson account. After doing this, click the **InetOrgPerson** menu item under **Action | New**. This will start the same wizard that allows you to create a new user account, but will use the InetOrgPerson object class to create it.

New & Noteworthy...

InetOrgPerson Object Class Is Part of the Windows Server 2003 Active Directory Schema

The InetOrgPerson object class and the attributes it contains originates from RFC 2798. RFC is an acronym for Requests for Comments, and is a document that is used to specify information and/or technical specifications. RFC 2798 was created by the Internet Engineering Task Force (IETF) to address the need for a class of user that accessed directory services over the intranet or Internet. This class of user was designed to hold attributes about people who accessed the directory using LDAP in this way.

Because of the need for this type of user class, Microsoft provided a kit that added an InetOrgPerson object class to the schema in Windows 2000. In Windows Server 2003, an InetOrgPerson is included in the Active Directory schema as a type of user class that can be used by LDAP applications requiring this type of object and when migrating to Active Directory from other directory services. This saves administrators from needing to extend the schema to create a new InetOrgPerson object class.

Creating User Accounts

Windows Server 2003 provides multiple ways of creating user accounts in Active Directory. As mentioned, Active Directory Users and Computers provides a GUI that allows you to create new accounts quickly and efficiently. As a new method of adding user accounts to Active Directory, you can also use the DSADD command. In the paragraphs that follow, we will look at each of these tools.

Creating Accounts Using Active Directory Users and Computers

Active Directory Users and Computers is a tool that is installed on DCs, and is used by those with the appropriate access to create domain accounts. Only members of the Administrators group, Account Operators group, Domain Admins group, Enterprise Admins group, or someone who's been delegated authority can create a user account. When someone is delegated authority to perform a task, it means that he or she been given

administrative credentials to carry out the action. Responsibility can be delegated through the Delegation of Control Wizard, Group Policy, or security groups (which we'll discuss later in this chapter).

Active Directory Users and Computers is started in a number of ways. As we saw in Chapter 1, the Active **Directory Users and Computers** snap-in can be loaded into **Microsoft Management Console** (MMC). Using the Windows **Start** menu can also start this tool by clicking on **Start | Administrative Tools | Active Directory Users and Computers**. The final method of starting it is through the Control Panel. In **Control Panel**, open **Performance and Maintenance | Administrative Tools | Active Directory Users and Computers**.

Once Active Directory Users and Computers is open, expand the domain in which you have access, and want to create the account. Within the domain, you select the container in which you want to create the user object. You can create accounts at the domain level, use an existing default container (such as Users), or use an OU that you've created. Once you've selected the container, click **Action | New | User**. InetOrgPerson accounts can also be created this way, by selecting **InetOrgPerson** instead of User from **Action | New**.

Once the **User** menu command has been clicked in the **Action | New** menu, a wizard will start that will take you through the steps of creating a new account. As shown in Figure 2.6, the first screen provides a number of fields in which you can enter information relating to the new user:

- **First name** The given name of the user.
- **Initials** The middle initial(s) in the user's name.
- **Last name** The surname of the user.
- **Full name** The entire name of the user. As you enter the user's first name, initials, and last name, this field will automatically be filled. It can, however, be modified. This name must be unique to the container in which the account is being created.
- **User logon name** The UPN that the user will enter to log on to the domain. The UPN (remember, this is comprised of the user logon name and the UPN suffix) must be unique to the enterprise. In other words, two logon names that are the same cannot exist within the same Active Directory forest. The drop-down list in this field allows you to select a UPN suffix.
- **User logon name (pre-Windows 2000)** The NetBIOS name that the user will use when logging on to the domain from pre-Windows 2000 operating systems.

Once you have entered this information, click the **Next** button to continue providing information to setup the account.

Figure 2.6 Screen Used to Enter New User Information When Creating a New User Object

The screenshot shows a dialog box titled "New Object - User". At the top, it says "Create in: knightware.ca/Users". Below this, there are several input fields: "First name:" with "Michael" entered, "Initials:" which is empty, "Last name:" with "Cross" entered, and "Full name:" with "Michael Cross" entered. There are also fields for "User logon name:" with "mcross" and a dropdown menu showing "@knightware.ca". Below that is a field for "User logon name (pre-Windows 2000):" with "KNIGHTWARE\" and "mcross" entered. At the bottom, there are three buttons: "< Back", "Next >", and "Cancel".

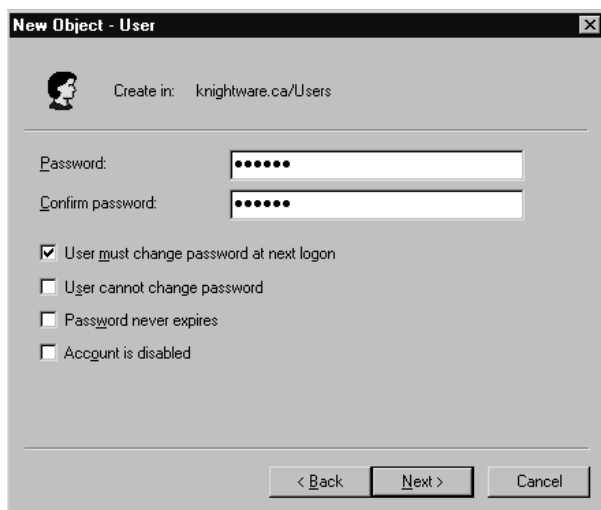
The screen that follows allows you to configure password information for the user's account. As with the previous screen, a number of fields can be used to manage the initial setup of the account. These fields are depicted in Figure 2.7, and consist of:

- **Password** The password that the user will enter, so he or she can be authenticated.
- **Confirm Password** The same password that's entered in the Password field. Confirming the password a second time ensures that it has been entered correctly.
- **User must change password at next logon** Checking this box requires the user to change his or her password at first logon. Having the user change the password, ensures that only the user knows the new password and is a good security practice. This setting is enabled by default.
- **User cannot change password** When selected, this ensures that only administrators can change the user's password. This can be used if the account is being used by more than one person (such as a Guest account), services, or programs. This setting is also used in more secure environments that require a guarantee of password strength. In these environments, administrators are often responsible for maintaining and changing passwords to ensure they meet security standards.
- **Password never expires** This setting exempts the user's password from the maximum password age Group Policy setting that is enabled for the domain. This is helpful if a service or program is using the account. If the **User must change password at next logon** option is set, this option is overridden.

- **Account is disabled** This check box is used to prevent a user from logging on with this account. This can be used if an employee has been terminated, is on extended leave or vacation, or hasn't started in his or her job yet.

After entering this information, click the **Next** button to reach the final screen for creating a new user account.

Figure 2.7 Screen Used to Enter Password Options When Creating a New User Object



As shown in Figure 2.8, the final screen used in creating a new user object provides summary information. The information that appears on this screen provides a recap of the user information and password options that will be used for creating the account. After the **Finish** button is clicked, the account is then created in the container you initially chose to store the object.

Figure 2.8 Summary Screen of New Object Dialog Box



Now that we've looked at how to create a user account, let's put that knowledge to practice and create one in the Exercise 2.02.

EXERCISE 2.02

CREATING A USER OBJECT IN ACTIVE DIRECTORY USERS AND COMPUTERS

1. Open **Active Directory Users and Computers** from **Start | Administrative Tools Active Directory Users and Computers**.
2. When the utility opens, expand the console tree so that your domain and the containers within it are visible.
3. Select the **TestOU** OU that you created in Chapter 1 from the console tree. From the **Action** menu, select **New | User**.
4. When the **New Object - User** dialog box appears, enter the following information in the corresponding fields:

Field	Data to Enter
First name	John
Initials	Q
Last name	Public
Full name	John Public
User logon name	Jpublic
User logon name (pre-Windows 2000)	Jpublic

5. After entering this information, click the **Next** button to continue.
6. Enter a password of your choosing in the **Password** field, and then reenter it in the **Confirm password** field.
7. Clear the **User must change password at next logon** check box.
8. Click **Next** to continue. When the summary screen appears, review the settings you have entered and click **Finish** to create the account.
9. From the **Action** menu, select **New | User**.
10. When the **New Object - User** dialog box appears, enter the following information in the corresponding fields:

Field	Data to Enter
First name	Jane
Last name	Doe
Full name	Jane Doe
User logon name	Jdoe
User logon name (pre-Windows 2000)	Jdoe

11. After entering this information, click the **Next** button to continue.
12. Enter a password of your choosing in the **Password** field, and then reenter it in the **Confirm password** field.
13. Click **Next** to continue. When the summary screen appears, review the settings you have entered and click **Finish** to create the account.
14. Log off and then log back on as the jdoe user. Notice that you are required to change the password.
15. Log off and then log back on as the jpublic user. Notice that you aren't required to change the password.

Creating Accounts Using the DSADD Command

As we saw in Chapter 1, Windows Server 2003 includes a number of command-line tools that allow you to perform common administrative tasks from a command prompt. Using the DSADD command, you can create new objects in Active Directory, including user objects. As is the case when using Active Directory Users and Computers, only members of the Administrators group, Account Operators group, Domain Admins group, Enterprise Admins group, or someone who's been delegated authority can create a user account. This means that the DSADD command can't be used as a workaround to creating an account without authorization.

Create a new user with DSADD by entering the following syntax:

```
DSADD USER UserDN [-samid SAMName] -pwd {Password|*}
```

In entering this command, the following parameters must be entered:

- **UserDN** This is the DN of the user object you are adding. This provides information on where the account will be created.
- **SAMName** This is a NetBIOS name, which is used when logging on from pre-Windows 2000 computers. If this parameter isn't added, DSADD will create one,

based on the first 20 characters of the common name you entered for the UserDN parameter.

- **Password** This is the password that will be used for this account. If an asterisk (*) is entered for this parameter, you will be prompted to enter a password.

In addition to these parameters, additional settings can be applied when creating a user account by using the following syntax. Note that this is all one long line.

```
dsadd user UserDN [-samid SAMName] [-upn UPN] [-fn FirstName] [-mi
Initial] [-ln LastName] [-display DisplayName] [-empid EmployeeID]
[-pwd {Password | *}] [-desc Description] [-memberof Group;...]
[-office Office] [-tel PhoneNumber] [-email Email] [-hometel
HomePhoneNumber] [-pager PagerNumber] [-mobile CellPhoneNumber]
[-fax FaxNumber] [-iptel IPPhoneNumber] [-webpg WebPage] [-title Title]
[-dept Department] [-company Company] [-mgr ManagerDN] [-hmdir
HomeDirectory] [-hmdirv DriveLetter:] [-profile ProfilePath] [-loscr
ScriptPath] [-mustchpwd {yes | no}] [-canchpwd {yes | no}]
[-reversiblepwd {yes | no}] [-pwdneverexpires {yes | no}] [-acctexpires
NumberOfDays] [-disabled {yes | no}] [{-s Server | -d Domain}]
[-u UserName] [-p {Password | *}] [-q] [{-uc | -uco | -uci}]
```

As you can see, a considerable number of options can be set in using the DSADD command, which are not available when initially creating an account with Active Directory Users and Computers. We'll explain how such information can be added to an account with Active Directory Users and Computers in the next section. First, let's examine the various parameters that can be used in association with the DSADD command. The parameters in this syntax are explained in Table 2.3.

Table 2.3 DSADD Parameters for Creating Users

Parameter	Description
-upn <i>UPN</i>	Specifies the UPN for the account.
-fn <i>FirstName</i>	Specifies the first name of the user.
-mi <i>Initial</i>	Specifies the initial(s) of the user.
-ln <i>LastName</i>	Specifies the last name of the user.
-display <i>DisplayName</i>	Specifies the display name of the account.
-empid <i>EmployeeID</i>	Specifies the user's employee ID.
-desc <i>Description</i>	Information that describes the account.
-memberof <i>Group</i>	Specifies the DNs of groups of which this account will be a member.
-office <i>Office</i>	Specifies the office location of the user.

Continued

Table 2.3 DSADD Parameters for Creating Users

Parameter	Description
-tel <i>PhoneNumber</i>	Specifies the telephone number of the user.
-email <i>Email</i>	Specifies the user's e-mail address.
-hometel <i>HomePhoneNumber</i>	Specifies the user's home telephone number.
-pager <i>PagerNumber</i>	Specifies the pager number of the user.
-mobile <i>CellPhoneNumber</i>	Specifies the cellular telephone number of the user.
-fax <i>FaxNumber</i>	Specifies the user's fax number.
-iptel <i>IPPhoneNumber</i>	Specifies the user's IP phone number.
-webpg <i>WebPage</i>	Specifies the URL of the user's Web page.
-title <i>Title</i>	Specifies the title of the user.
-dept <i>Department</i>	Specifies the user's department.
-company <i>Company</i>	Specifies company information.
-mgr <i>ManagerDN</i>	Specifies the DN of the user's manager.
-hmdir <i>HomeDirectory</i>	Specifies the home directory of the user.
-hmdrv <i>DriveLetter</i> :	Specifies the drive letter used by the user to access his or her home directory. This parameter is used if the <i>HomeDirectory</i> is specified using the universal naming convention.
-profile <i>ProfilePath</i>	Specifies the profile path for the account.
-loscr <i>ScriptPath</i>	Specifies the logon script path for the account.
-mustchpwd {yes no}	Specifies whether the user needs to change his or her password the next time he or she logs on. By default, the user doesn't need to change the password, so this would be the same as specifying no for this parameter.
-canchpwd {yes no}	Specifies whether the user is allowed to change his or her password. By default, the user can change his or her password, so this would be the same as specifying yes for this parameter.
-reversiblepwd {yes no}	Specifies whether the password is stored using reversible encryption, which is used by Macintosh computers and some forms of Windows-based authentication. By default, reversible encryption isn't used, so this is the same as this parameter being set to no.

Continued

Table 2.3 DSADD Parameters for Creating Users

Parameter	Description
-pwdneverexpires {yes no}	Specifies whether the password expires. By default, a password will expire, so this is the same as this parameter being set to no.
-acctexpires <i>NumberOfDays</i>	Specifies the number of days before the account expires. If the value of <i>NumberOfDays</i> is set to 0, the account will expire at the end of the day. If the value of <i>NumberOfDays</i> is set to a negative value, it will set that the account has already expired that many days ago. If set to a positive value, it will expire that many days in the future. If the value of <i>NumberOfDays</i> is set to Never, the account will never expire.
-disabled {yes no}	Specifies whether the account has been disabled. By default, the account is enabled, so this is the same as this parameter being set to no.
{-s <i>Server</i> -d <i>Domain</i> }	Specifies to connect to a remote server or domain. By default, the computer is connected to the DC in the logon domain.
-u <i>UserName</i>	Specifies the username to log on to a remote server. By default, the username that the user is logged on to the local system with is used. The following formats can be used for the <i>UserName</i> variable: Username Domain\username User principal name
-p { <i>Password</i> *}	Specifies the password to log on to a remote server. If an asterisk (*) is used, you will be prompted for a password.
-q	Specifies quiet mode, and suppresses output.
{-uc -uco -uci}	Specifies Unicode to be used for input or output. If -uc is used, then input or output is to a pipe (). If -uco is used, then output is to a pipe or file. If -uci is used, then input is from a pipe or file.

Managing User Accounts

Managing user accounts is done through the properties of the object, which is accessible by using Active Directory Users and Computers. You can access the properties of a user object

by selecting the object, and then clicking on **Action | Properties**. You can also right-click on the object and select **Properties** from the context menu.

Upon opening the Properties of the user, you will see a number of tabs that allow you to set various options and provide information dealing with the account:

- **General** This tab contains basic information about the user, including name, telephone number, and other information.
- **Address** This tab contains personal properties about the user, including home address.
- **Account** This tab contains settings that control account expiration, logon hours, and other account options.
- **Profile** This tab contains settings for the user's logon script, user profile, and home directory.
- **Telephones** This tab contains telephone numbers and other contact information about the user.
- **Organization** This tab contains information about the user's organization, department, and manager.
- **Environment** This tab contains settings for configuring the startup environment for Terminal Services.
- **Sessions** This tab contains timeout and reconnection settings for Terminal Services.
- **Remote Control** This tab configures remote control settings for Terminal Services.
- **Terminal Services Profile** This tab contains configuration information for the user's Terminal Services profile and home directory.
- **COM+** This tab allows you to set whether the user is a member of a particular COM+ partition set.
- **Published Certificates** This tab lists X.509 certificates and their use.
- **Member Of** This tab lists the user's group membership.
- **Dial-in** This tab contains settings used by the user for dialing in to or using a virtual private network (VPN) connection to the network.
- **Object** This tab shows the fully qualified name of the user.
- **Security** This tab displays an ACL for the user account object.

Individually, each of these tabs allows you to manage different settings related to the user account. However, a number of these tabs are related, in that they deal with particular aspects of user account management. As we'll see in the sections that follow, by using them together, you can configure how the account can be used.

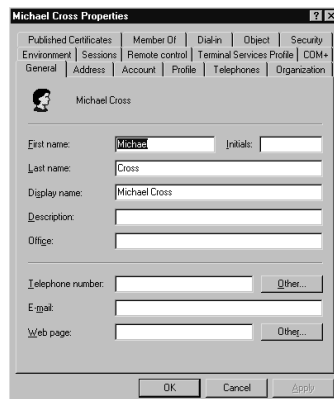
Personal Information Tabs

In looking at what properties can be set with these tabs, you will see that there are four tabs that contain personal information about the user: General, Address, Telephones, and Organization. As shown in Figure 2.9, the General tab contains a number of fields that contain information provided when the account was initially created:

- **First name** This field contains the given name of the user.
- **Initials** This field contains the middle initial(s) of the user.
- **Last name** This field contains the user's surname.
- **Display name** This field contains the full name of the user.
- **Description** This field allows you to provide a description of the user account. For example, if this were an account that was to be used for a service or application, you could provide information on which service or program will use this account.
- **Office** This field allows you to provide information on the office in which the user works.
- **Telephone number** This field allows you to provide the user's telephone number.
- **Email** This field allows you to enter the e-mail address of the user.
- **Web page** This field allows you to enter the Web page belonging to the user.

In looking at this tab, notice that Telephone and Web page fields have a button beside them named **Other**. When this button is clicked, a dialog box will open that allows you to enter additional entries. As you might guess, this is because many users might have more than one Web page or telephone number associated with them. If additional entries exist, you can also click the **Other** button to view these entries in the dialog box that appears.

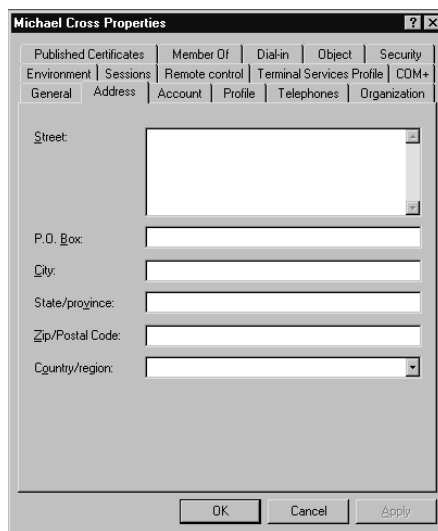
Figure 2.9 General Tab of User's Properties



The Address tab is used to store contact information dealing with a user's physical or mailing address. As shown in Figure 2.10, information on this tab includes:

- **Street** This text box allows you to enter the full street address of the user. It allows multiple lines, to account for apartment numbers and other information.
- **P.O. Box** This text box allows you to enter the P.O. Box number.
- **City** This text box is the city in which the user lives or has a mailing address.
- **State/province** This text box is the state or province in which the user lives or has a mailing address.
- **Zip/Postal Code** This text box is the zip code or postal code in which the user lives or has a mailing address.
- **Country/region** This drop-down box contains list of countries so that you can select the one in which the user lives or has a mailing address.

Figure 2.10 Address Tab of a User's Properties



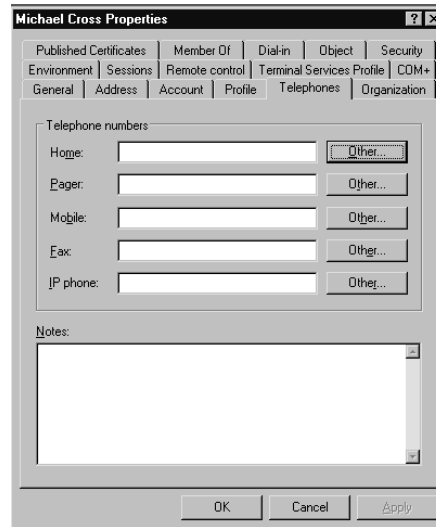
The Telephones tab is another tab that contains personal properties related to the user. As shown in Figure 2.11, this tab provides contact information relating to various methods of verbal or digital communication:

- **Home** This text box is used to enter the home telephone number of the user.
- **Pager** This text box is used to enter the pager number of the user.
- **Mobile** This text box is used to enter the cellular phone number of the user.
- **Fax** This text box is used to enter the fax number of the user.

- **IP phone** This text box is used to enter the IP phone number of the user.
- **Notes** This text box is used to enter any additional notes pertaining to the user.

Because users might have multiple telephone numbers, pagers, and other methods of communication, each of these fields (except for Notes) also includes an Other button. As discussed earlier, clicking this button allows you to enter and/or view additional entries for that particular field.

Figure 2.11 Telephones Tab of User's Properties

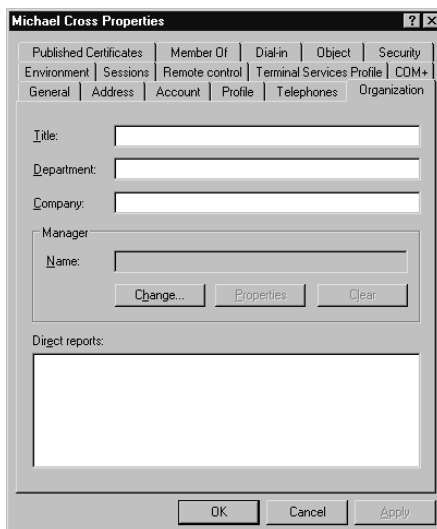


The Organization tab is the final tab that contains personal properties for the user. This tab allows you to enter information relating to the organization in which the user works. As seen in Figure 2.12, the fields on this tab include:

- **Title** This text box is used to specify the title of the user.
- **Department** This text box is used to specify the department in which the user works.
- **Company** This text box is used to specify the name of the company the user works for.
- **Manager** This field allows you to specify the DN pointing to the user's manager or supervisor within the company. With this field, you can click the **Change** button to change the DN pointing to the manager's user object, **Properties** to view that object's properties, or **Clear** to erase the entry pointing to the manager's user object.

- **Direct reports** This field shows any other user accounts that have this user's account specified as their manager.

Figure 2.12 Organization Tab of User's Properties



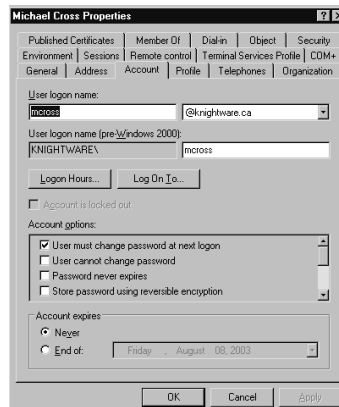
Account Settings

Not all of the tabs in the user's Properties deal with personal information. As seen in Figure 2.13, the Account tab is used to store information relating to the domain user account, including password options. The fields on this tab include:

- **User logon name** This text box is used to specify the UPN that the user will use when logging on to the domain. This field also includes a drop-down list for specifying the UPN suffix.
- **User logon name (pre-Windows 2000)** This text box is used to specify the logon name that is used when logging on from pre-Windows 2000 computers.
- **Account is locked out** This check box specifies that the account is locked out, preventing the user from logging in.
- **User must change password at next logon** This check box specifies that the user must change his or her password when he or she first logs on.
- **User cannot change password** This check box prevents the user from changing his or her password. This is used to restrict password changes so only administrators have the ability to manage them.

- **Password never expires** This check box prevents the password from expiring after a specific time. If the **User must change password at next logon** option is set, this option is overridden.
- **Store password using reversible encryption** This check box requires users to use reversible encryption. Users who are logging on from Macintosh computers require this setting. Some forms of Windows authentication, such as CHAP and Digest, also require this setting to be enabled.
- **Account is disabled** This check box prevents users from logging on with this account.
- **Smart card is required for interactive logon** This check box allows the user to log on using a smart card.
- **Account is trusted for delegation** This allows the account to be used to run as an identity of a service, so that it can impersonate the account and acquire necessary access.
- **Account is sensitive and cannot be delegated** This check box allows a user to assign responsibility over a portion of the namespace to another user, group, or computer.
- **Use DES encryption types for this account** This check box requires Data Encryption Standard (DES) to be used with this account.
- **Do not require Kerberos preauthentication** This check box removes the need for preauthentication for accounts that are using another version of Kerberos that doesn't require it.
- **Account expires** This option button is used to set the expiration date of the account. Options for this field are Never, which means the account never expires; or End of, which requires a date to be set specifying when the account will expire.


Figure 2.13 Account Tab of User's Properties



In addition to the fields we've discussed, the Account tab also includes a Logon Hours button, which opens a dialog box that allows you to control when this user can log on or remain logged on to the network. By default, users are able to log on and remain logged on to the network 24 hours a day, 7 days a week. However, in secure environments, you might want to control when a user is able to log on. To provide a maintenance window, you might want to limit users' ability to log on or remain logged on after regular hours of work, or during weekends.

As shown in Figure 2.14, the Logon Hours dialog box contains a series of boxes that determine the times and days when a user can log on. After selecting the boxes representing the times and dates to log on, click the **Logon Permitted** or **Logon Denied** option buttons to respectively permit or deny access during those times. If all of the boxes are selected and **Logon Permitted** is selected, then there are no restrictions set for the user.

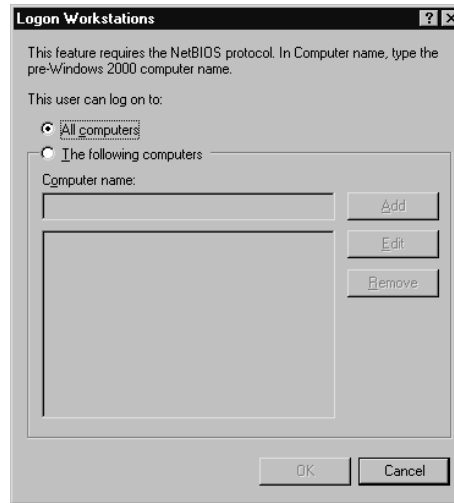
Figure 2.14 Logon Hours Dialog Box



The other button that appears on the Account tab is the Log On To button. When this button is clicked, the Logon Workstations dialog box shown in Figure 2.15 appears. On this dialog box, you can control what computers the user can use when logging on to the domain. By default, users can log on from any computer. However, by using the fields on this tab you can heighten security by limiting users to working on the machine at their desk, or a group of computers within their department. For example, you might want to prevent users from logging on to the domain from a specific machine so that they cannot access another user's data that is stored on that computer.

On this dialog box, there are two options: **All computers** and **The following computers**. When **All computers** is selected, there are no restrictions regarding which machines a user can log on from. When **The following computers** is selected, you can then enter the name of the computer(s) you want to restrict a user to using. After entering the NetBIOS name of the computer, click the **Add** button to add that computer to the list. You can then select computer names from the listing, and click **Edit** to modify the entry or **Remove** to delete it from the list.

Figure 2.15 Logon Workstations Dialog Box



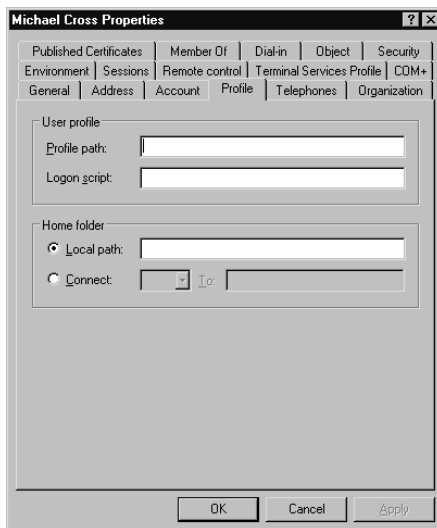
The Profile tab is also used to configure elements of the user's account, relating to profiles, logon scripts, and home folders. Roaming profiles can be used to provide consistency across the network, by ensuring that a user has the same desktop environment, application settings, drive mappings, and personal data regardless of which computer he or she uses on the network. The **Profile path** field on this tab is used to specify the path to the user's profile. Similarly, logon scripts are also used to apply settings to a user's account, by running a script when the user logs on to the network. The **Logon script** field is used to set where this script is located, so it will automatically run each time the user logs on to this account. Through these, the user's environment is configured each time he or she logs on to a DC.



NOTE

The **Logon script** field on this tab is a hold over from Windows NT days. Group Policy allows you to specify multiple logon, logoff, startup, and shutdown scripts. However, if the computers that your users are logging on from does not support Group Policy, this field comes in handy to ensure that they can still get a logon script applied to them. Windows 2000 and later computers are capable of using Group Policy.

Finally, as shown in Figure 2.16, the **Home folder** section of this tab is used to specify the location of a home directory that will contain the user's personal files. The **Local path** text box is used to specify a path to the directory on the local system. Alternatively, you can specify a network location by using the **Connect** drop-down box to specify a drive letter that the path will be mapped to, and then enter a UNC path to the directory in the **To** text box.

Figure 2.16 Profile Tab of User's Properties

Terminal Services Tabs

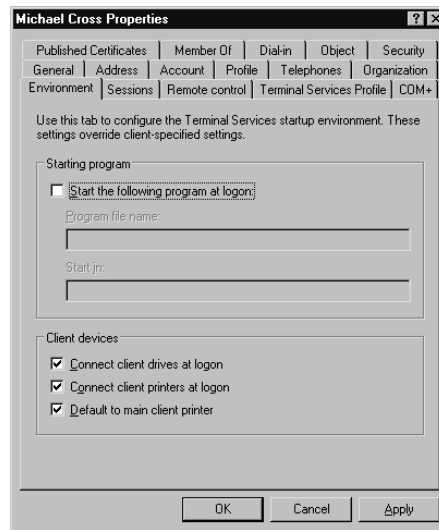
Terminal Services allows users to access applications that are run on the server. Rather than running as independent computers, the clients that access Terminal Services act as dumb terminals. Throughout their session with the server, users run an environment that allows them to view and interact with programs that appear to run on the local computer but are actually running on the Terminal Services server located on the network.

The Properties dialog box of a user provides four tabs that specifically deal with Terminal Services: Environment, Sessions, Remote Control, and Terminal Services Profile. As seen in Figure 2.17, the **Environment** tab is used to configure settings for Terminal Service's startup environment. By default, users receive a Windows Server 2003 desktop when connecting using Terminal Services. The **Starting program** section contains fields for specifying a particular program to run when logging on to Terminal Services. If this option is enabled, users will receive the program instead of a desktop. When the **Start the following program at logon** check box is selected, you can enter the path and executable name for the program.

The **Client devices** section also allows you to configure how devices on the computer you're working on will be dealt with. The check boxes you can check to enable these settings are:

- **Connect client drives at logon** This check box causes drive letters to be mapped to local client drives in the user's Terminal Services session. This enables these local drives to be accessed from the session. Remember, when you are working in a session, you are actually working on the server, not your client. Therefore, the drives you see in Explorer are the drives on the server, not the client. When this box is checked, you see both.

Figure 2.17 Environment Tab of User's Properties



- **Connect client printers at logon** This check box causes printers that are configured on the client computer to be connected and available while you are working in a Terminal Services session.
- **Default to main client printer** This check box causes the default printer on your client computer to be the default printer in your Terminal Services session.

In addition to the settings on the Environment tab, the Sessions tab is also used for configuring Terminal Services. As seen in Figure 2.18, this tab includes numerous settings for configuring timeout and reconnection settings for Terminal Services sessions:

- **End a disconnected session** This drop-down box allows you to choose how long a disconnected session should remain on the server. If set to **Never**, which is the default setting, a disconnected session will remain on the server indefinitely. If set to a specific time limit, the session will be terminated on the server when this time limit is reached.
- **Active session limit** This drop-down box is used to specify how long a user's session should remain active. The default setting is **Never**, meaning that the user's session can remain active indefinitely. If set to a limit by selecting a time from the drop-down list, the user's session will remain active until that time is reached. At this time, the user is either disconnected from the session or the session is terminated, depending on what is set for the **When a session limit is reached or connection is broken** option.
- **Idle session limit** This drop-down box controls how long a user can remain connected without any activity before the session ends or the user is disconnected

from it, depending on what is set for the **When a session limit is reached or connection is broken** option.

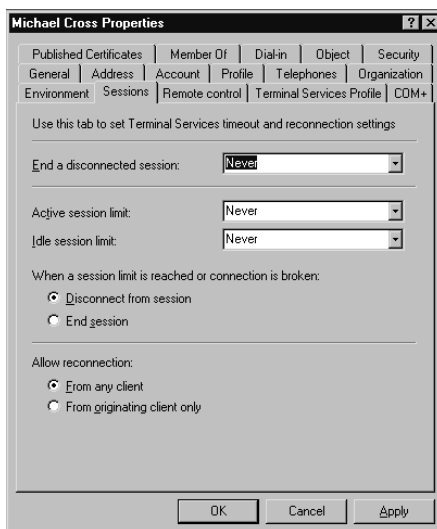
- **When a session limit is reached or connection is broken** This option button allows you to set what will happen when either of these situations occurs. You can choose either **Disconnect from session** or **End session**.
- **Allow reconnection** This option button allows you to control how a user can reconnect to a disconnected session. You can choose to either reconnect **From any client** or **From originating client only**.



NOTE

There are two ways to end a Terminal Services session. The first is to log off. This completely destroys the session that you were using on the server and frees up the resources it was consuming. The second option is to disconnect. This simply breaks the communication from your client to the Terminal Services server. It does not terminate your session, which means that it continues to consume resources on the server. When you reconnect to a disconnected session, it appears exactly as it did when you disconnected from it.

Figure 2.18 Sessions Tab of User's Properties

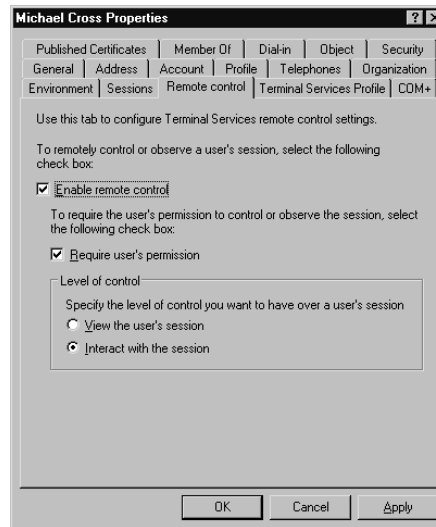


The Remote Control tab allows you to configure remote control settings for the user, which enables others to take over a session. By taking over the computer, the other person can then perform actions on the remote computer, enabling that person to perform various

actions and show the user how to do certain tasks. As shown in Figure 2.19, the fields available to configure these settings are:

- **Enable remote control** This check box determines whether others can control or view a user's session. This turns on or off remote control for a session belonging to a user.
- **Require user's permission** This check box determines whether the user needs to give permission to allow another user to control or view a session.
- **Level of control** This option button determines how someone can use remote control with the user's session. The options include **View the user's session**, which only allows someone else to view what the user is doing, and **Interact with the session**, which allows another user to take control of the session and perform tasks as if he or she were sitting at the computer of the user that initiated the session.

Figure 2.19 Remote Control Tab of User's Properties



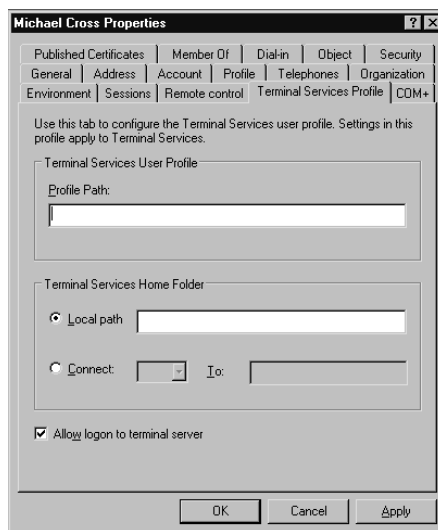
The Terminal Services Profile tab is similar to the Profile tab discussed earlier, except that settings on this tab exclusively relate to a user's Terminal Services session. As shown in Figure 2.20, this tab provides the following fields:

- **Profile Path** This text box specifies the path to the user's profile that should be used with Terminal Services. If no path is specified, the setting from the Profile tab is used.
- **Terminal Services Home Folder** This text box allows you to specify where the user's home directory that should be used during a Terminal Services session is

located. As with the Profiles tab, this section provides two options. The **Local path** field is used to specify a local path to a location on the Terminal Services computer. The **Connect** drop-down box allows you to specify a drive letter for a mapped drive that will be available within the session, and a UNC path to the network location of the home directory. If no home folder is specified, the settings on the Profile tab are used.

- **Allow logon to terminal server** This check box is used to enable users to establish a Terminal Services session. It is enabled by default.

Figure 2.20 Terminal Services Profile

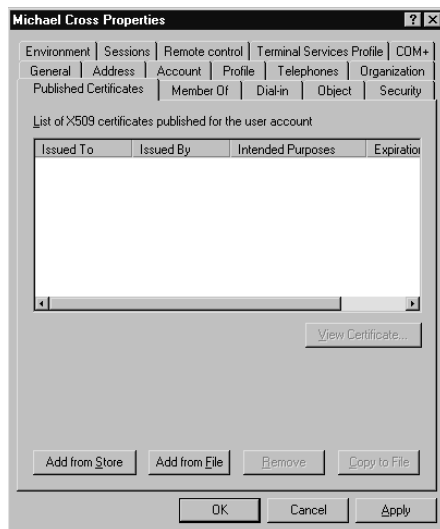


Security-Related Tabs

Several tabs are available through the user object's properties that control security settings associated with the account. These tabs are Published Certificates, Dial-in, Security, and Member Of. Together, they allow you to manage issues related to access control and authentication.

The Published Certificates tab provides a listing of certificates that are used by the account, and allows you to add others. As shown in Figure 2.21, this tab allows you to view any X.509 certificates that have been published for the user account, and includes fields that explain who it was issued by, who it was issued to, the intended purpose of the certificate, and its expiration date. The **Add from Store** button can be used to add additional certificates to the listing from the computer's local certificate store. The **Add from File** button can also be used to add a certificate from a file. If a certificate is no longer needed, you can select the one you no longer want to be applied to the account and click the **Remove** button. Finally, the **Copy to File** button will export the certificate that is selected in the list to a file.

Figure 2.21 Published Certificates



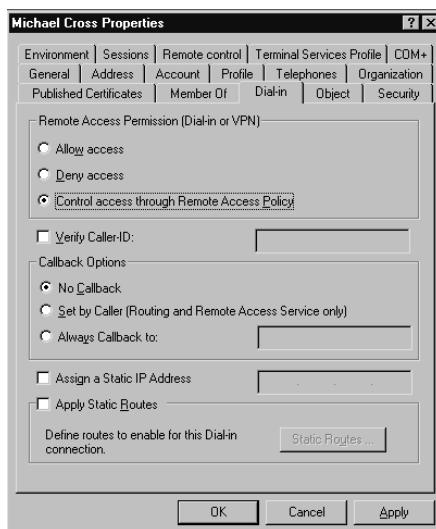
The Dial-in tab allows you to configure settings that are used when the user attempts to connect to the network remotely using a dial-up or VPN connection. These settings are applied when the user dials in to a Windows Server 2003 remote access server or attempts to use a VPN connection. As shown in Figure 2.22, this tab is divided into several sections with multiple options, including:

- **Remote Access Permission (Dial-in or VPN)** This option button specifies whether the user can connect to the network via a dial-up or VPN connection. The options in this section include **Allow access**, which enables dial-in or VPN remote access; **Deny access**, which prohibits dial-in or VPN remote access; and **Control access through a Remote Access Policy**, which is the default option and specifies that a remote access policy is used to control permission for remote access.
- **Verify Caller-ID** This check box allows you to specify the telephone number that the user must be calling from in order to establish a successful connection. It requires hardware capable of detecting the number that the user is calling from.
- **Callback Options** The configuration settings in this section are **No Callback**, **Set by Caller (Routing and Remote Access Service Only)**, and **Always Callback To**. **No Callback** is the default option. It enables users to connect remotely and without the use of callback. When this option is set, the user will pay for any long distance charges. **Set by Caller (Routing and Remote Access Service Only)** allows the caller to specify a telephone number that the server will call back. When a remote connection is made, the user is prompted for a username and password. If successfully authenticated, the settings on this tab are checked and the user is prompted for a telephone number to be called back at.

The server then disconnects and calls the user back at that number. This allows the company to pay for any long distance fees, which typically results in cost savings. **Always Callback To** is the final option. This is a security, not a cost savings, option that forces the server to call the user back at a preconfigured telephone number. Because this setting requires the user to be at that telephone number, the risk of unauthorized users attempting to connect remotely is reduced.

- **Assign a Static IP Address** This check box assigns a specific IP address to the user when the user connects remotely.
- **Apply Static Routes** This check box places additional routes in the routing table upon connection.
- **Static Routes** This button is used to define the additional routes that will be placed in the routing table upon connection.

Figure 2.22 Dial-In Tab



As we saw in Chapter 1, the Security tab (Figure 2.23) is used to configure what permissions other users and groups have to an object. This tab consists of two panes. The top pane lists users and groups that have been added to the DACL for the account. It also allows you to add or remove users and groups from the DACL. In the lower pane, you can enable or disable specific permissions by checking a check box in the Allow or Deny column. Special permissions can also be set for objects by clicking the **Advanced** button, which displays a dialog box (seen in Figure 2.24) where additional permissions can be applied.

Figure 2.23 Security Tab

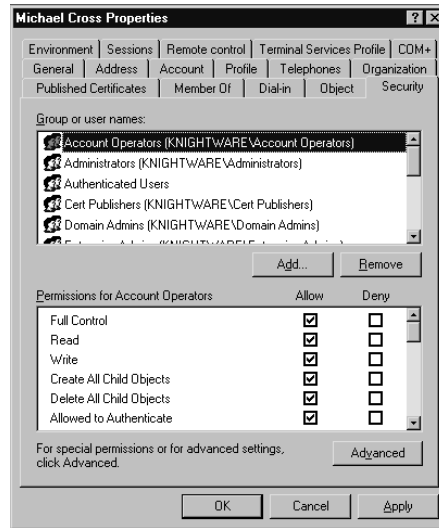
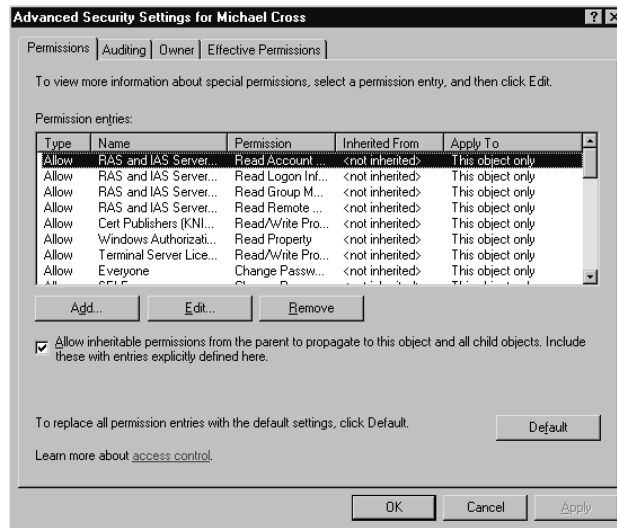


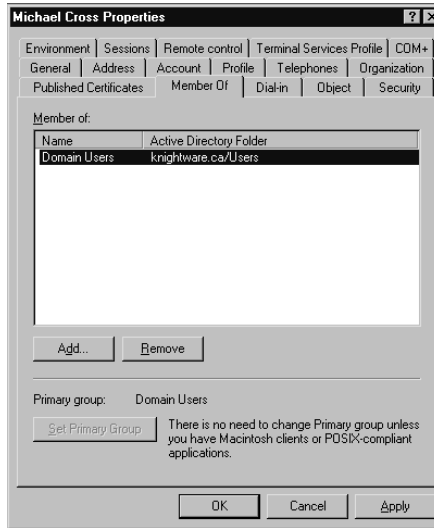
Figure 2.24 Special Permissions Dialog Box



As seen in Figure 2.24, the Special Permissions dialog box that's accessed through the **Advanced** button of the Security tab allows you to configure advanced settings and apply additional permissions to an account. As seen in this dialog, the Permissions tab also provides an option labeled **Allow inheritable permissions from the parent to propagate to this object and all child objects**. When this check box is checked, any permissions applied to the parent object (which in this case would be an OU) are also applied to this account. If this check box is unchecked, then any permissions applied at the higher level will not be applied, and the object will only have the permissions that have been explicitly set for it.

The final tab we'll discuss is the **Member Of** tab. As seen in Figure 2.25, this tab provides a listing of the user's group membership(s). By clicking the **Add** button, a dialog box will appear with a list of available groups of which the user can become a member. Selecting a group from the list on the **Member Of** tab and clicking the **Remove** button will remove that user from the group's membership.

Figure 2.25 Member Of Tab




At the bottom of this tab is a button called **Set Primary Group**, which only applies to a limited number of users. A primary group is needed by users who use Macintosh computers, and log on to the network through File or Print Services for Macintosh. The other users who require a primary group are users who are running POSIX-compliant applications.



EXAM WARNING

The default primary group of a new account generally doesn't need to be changed. The primary group for new user accounts is Domain Users, while the primary group for new computer accounts is Domain Computers.

To fully understand how the **Member Of** tab affects a user's level of security, we must look at how groups impact a user's access. In the section that follows, we will look at the various groups that users can become members of, and see what each group offers.



Working with Active Directory Group Accounts

Groups are collections of accounts that are grouped together, so that they can be granted rights and permissions as a single unit. By using groups, you aren't limited to granting rights and permissions to each account individually. You can make user accounts, computer accounts, and other groups members of a group, and grant rights and permissions to all of them by making changes only at the group level. When a group is given permissions, privileges, or rights, what is granted to it applies to all accounts that are members of the group.

Using groups, you can perform a variety of tasks that will affect the accounts and groups that are members. These include:

- Assigning rights to a group account to authorize them to perform a certain task
- Assigning permissions on shared resources to a group, so that all members can access the resource in the same manner
- Distributing bulk e-mail to all members of the group

As we'll see in the sections that follow, group accounts are a powerful tool for managing large numbers of users as if they used a single account. In associating accounts with groups, you will find that some groups will have a much larger membership than others, and some will be used for purposes other than dealing with security issues.

Group Types

The first step in working with group accounts is deciding on the type of group you want to create and work with. In Active Directory, there are two different types, which are used for two different purposes:

- Security groups
- Distribution groups

The difference between these groups resides in how they are used. Security groups are designed to be used for security purposes, while distribution groups are designed to be used for sending bulk e-mail to collections of users. As we'll see in the paragraphs that follow, the type of group being used not only relates to its intended purpose, but how it will be handled in terms of security.

Once you create a particular type of group account, it is possible to switch its type at any time. If you create a security group and later decide to convert it into a distribution group (or vice versa), Active Directory will allow it depending on the domain functional level that's been set. If the domain functional level is set to Windows 2000 native or higher, the conversion can take place. However, it might not be allowed if the domain is running at the Windows 2000 mixed level.

Security Groups

Security groups are what most people think of when discussing groups. A security group is a collection of users who have specific rights and permissions to resources. Rather than giving rights to perform certain tasks to individual users, and then setting permissions as to what resources that user can access, the rights and permissions are applied to the group. Any users who are members of the group then acquire this same level of security access. In doing so, collections of users are handled as a single unit, rather than as individuals.

Although both can be applied to a group account, rights and permissions are different from one another. Rights are assigned to users and groups, and control the actions a user or member of a group can take. For example, a member of the Backup Operator's group has the ability to back up and restore data, while a member of the Administrators group has the ability to perform almost any action. As we'll see later in this chapter, there are a number of security groups to which users can be added, and each of these groups provides differing levels of access. In Windows Server 2003, rights are also sometimes called *privileges*. You might have noticed this earlier when viewing the output of the command *WHOAMI /ALL*.

Permissions are used to control access to resources. When permissions are assigned to a group, it determines what the members of the group can do with a particular resource. For example, one group might only be given Read permissions to a file (so they can view but not modify it), while another group might be given Full Control (allowing them to do anything to the file). Through permissions, you can control the level of access a user or group receives to a shared resource.

Security groups are able to obtain such access because they are given a SID when the group account is first created. Because it has a SID, it can be part of a DACL, which lists the permissions users and groups have to a resource. When the user logs on, an access token is created that includes their SID and those of any groups of which they're a part. When they try to access a resource, this access token is compared to the DACL to see what permissions should be given to the user. It is through this process and the use of groups that the user obtains more (and in some cases, less) access than has been explicitly given to his or her account.

Another benefit of a security group is that you can send e-mail to it. When e-mail is sent to a group, every member of the group receives the e-mail. In doing so, this saves having to send an e-mail message to each individual user.

Distribution Groups

While security groups are used for access control, distribution groups are used for sharing information. This type of group has nothing to do with security. It is used for distributing e-mail messages to groups of users. Rather than sending the same message to one user after another, distribution groups allow applications such as Microsoft Exchange to send e-mails to collections of users.

The reason why distribution groups can't be used for security purposes is because they can't be listed in DACLs. When a new distribution group is created, it isn't given a SID,

preventing it from being listed in the DACL. Although users who are members of different security groups can be added to a distribution group, it has no effect on the permissions and rights associated with their accounts.

Head of the Class...

Security Groups vs. Distribution Groups

In looking at the differences between security groups and distribution groups, it might seem confusing as to why two such groups are used. After all, both can be used as e-mail entities where e-mail can be sent to the group with the message being distributed to members. The reason why two types of groups are used lies in the fact that only security groups have SIDs.

Because security groups have SIDs, they can be included in DACLs and be given permissions to resources. Distribution groups don't use SIDs, so they can't be given permissions. For this reason, you can add users to as many distribution groups as needed, and never have to worry that they've accidentally been given greater access because of their membership. If distribution groups didn't exist, you might add a user to a security group so he or she could receive e-mails meant for that group, but erroneously give that user greater access than desired.

The other reason why two types of groups are used is for performance. When security groups are used, the SID of that group is added to the access token of its members when they log on. The more security groups a user belongs to, the longer it will take to log on. Each security group that's added to the access token increases the token's size, making the logon and subsequent resource access slower and more resource intensive. When distribution groups are used, there is no SID for the group, so the number of distribution groups a user is a member of has no effect on the token's size. Because of these issues, it is always better to use distribution groups for e-mail needs, and security groups for collections of people who need access.

Group Scopes in Active Directory

Scope is the range that a group will extend over a domain, tree, and forest. The scope is used to determine the level of security that will apply to a group, which users can be added to its membership, and the resources that they will have permission to access. As we'll discuss in the sections that follow, Active Directory provides three different scopes for groups:

- Universal
- Global
- Domain Local

Universal

Universal groups have the widest scope of any of the different group scopes. Members of this group are able to contain accounts and groups from any domain in the forest, and can be assigned permissions to resources in any domain in the forest. In other words, it is all encompassing within any part of the forest.

Whether a universal security group can be used depends on the functional level that the domain has been set to. Domains that have the functional level set to Windows 2000 mixed won't allow universal security groups to be created. However, if the domain functional level is Windows 2000 native or Windows Server 2003, then universal security groups can be created. In this situation, the group can contain user accounts, global groups, and universal groups from any domain in the forest, and be assigned permissions to resources in any domain. Universal distribution groups can be used at any functional level, including Windows 2000 mixed.

Universal groups can be converted to groups with a lesser scope. Providing the group doesn't contain any universal groups as members, a universal group can be converted to a global group or a domain local group. If universal groups are members of the universal group that's being converted, you won't be able to perform the conversion until these members are removed.



NOTE

Universal groups are stored in the GC, along with their membership lists. Because of this, any change in membership triggers forestwide replication. To limit the impact of this type of replication, Microsoft recommends using relatively static members (such as global groups) in these groups.

Global

Global groups have a narrower scope than universal groups. A global group can contain accounts and groups from the domain in which it is created, and be assigned permissions to resources in any domain in a tree or forest. Because it only applies to the domain in which it's created, this type of group is commonly used to organize accounts that have similar access requirements.

As we saw with universal groups, however, the members that can be part of a global group depend on the domain functional level. If the functional level of the domain is set to Windows 2000 mixed, then the membership of a global group can only consist of user accounts from the same domain. If the functional level of the domain is set to Windows 2000 native or Windows Server 2003, then the global group can have user accounts and other global groups from the same domain as members. User accounts and global groups from other domains cannot become members of a global group.

Global groups can also be converted into a universal group, provided that the global group isn't a member of any other global groups. If other global groups are members of the global group, then these must be removed before the conversion can take place. The domain functional level must be Windows 2000 native or Windows Server 2003 to convert to a universal security group.

Domain Local

Domain local groups also have a scope that extends to the local domain, and are used to assign permissions to local resources. The difference between domain local and global groups is that user accounts, global groups, and universal groups from any domain can be added to a domain local group. Because of its limited scope, however, members can only be assigned permissions within the domain in which this group is created.

As you might expect from the two previous scopes, the abilities of a domain local group depends on the domain functional level. If the functional level is set to Windows 2000 mixed, then the domain local group can only contain user accounts and global groups from any domain. It cannot contain universal groups when Windows Server 2003 is using this level of functionality. If the functional level is set to Windows 2000 native or Windows Server 2003, then the domain local group can contain user accounts and global groups from any domain, as well as universal groups. In addition, it can contain other domain local groups from the same domain. These abilities, however, have no impact on permissions. In all cases, permissions can only be assigned to resources in the local domain.

Domain local groups can be converted to a universal group, provided that there are no other domain local groups in its membership. If the domain local group does have other domain local groups as members, then these must be removed from the membership before a conversion is made.

Built-In Group Accounts

As we saw when we discussed user objects, a number of built-in accounts are automatically created when you install Active Directory. This not only applies to user accounts, but group accounts as well. Many of these groups have preconfigured rights, which allow members to perform specific tasks. When users are added to these groups, they are given these rights in addition to any assigned permissions to access resources.

The groups that are created when Active Directory is installed can be accessed through Active Directory Users and Computers, and are located in two containers: Builtin and Users. Although they are stored in these containers, they can be moved to other OUs within the domain. Those in the Built-in container have a domain local scope, while those in the Users container have either a domain local, global, or universal scope. In the paragraphs that follow, we will look at the individual groups located in each of these containers, and see what rights they have to perform network-related tasks.

Default Groups in Builtin Container

Up to 14 different built-in groups that might be located by default in the Builtin container, including:

- **Account Operators**, which allows members to manage accounts
- **Administrators**, which gives members full control
- **Backup Operators**, which allows members to back up and restore files
- **Guests**, which gives members minimal access
- **Incoming Forest Trust Builders**, which is only available in forest root domains, and gives members permission to Create Inbound Forest Trusts
- **Network Configuration Operators**, which allows members to manage network settings
- **Performance Monitor Users**, which allows users to manage performance counters and use System Monitor
- **Performance Log Users**, which allows users to manage performance counters and use Performance Logs and Alerts
- **Pre-Windows 2000 Compatible Access**, which is used for backward compatibility
- **Print Operators**, which allows members to manage printers
- **Remote Desktop Users**, which allows members to connect to servers using Remote Desktop
- **Replicator**, which is used for replication purposes
- **Server Operators**, which allows members to manage servers
- **Users**, which contains every user account created in the domain

The Account Operators group is used to allow members to perform group management. Users who are part of its membership have the ability to create, modify, and delete many of the accounts that are stored in Active Directory. They can manage accounts in any OU except the Domain Controllers OU, or those located in the Users or Computers containers. To prevent members of this group from affecting administrator accounts, members of the Account Operators group cannot modify the Administrators and Domain Admins groups, or any accounts that are members of these groups.

Members of the Account Operators group also have certain abilities when dealing with DCs in the domain in which this group is located. They can log on locally to a DC, which means that they can physically sit at a DC and log on to it. In doing so, they could then make modifications to the DC. They also have the ability to shut down the DC, which is useful if there is a problem with the DC and no one else is available to restart the system.

The Administrators group is the most powerful of the groups in the Builtin container, and has full control over the domain. This account can access DCs over the network, back up files and directories, change system time, adjust memory quotes, create page files, load and unload device drivers, delegate responsibility to users and computers, shut down the system, and perform other tasks relating to accounts and DCs. By default, Domain Admins and Enterprise Admins groups and the Administrator account are members of the Administrators group.

The Backup Operators group is used to give members the ability to back up and restore files on DCs. It doesn't matter what the member's permissions on different files are, as they can back up and restore any file on the system. In addition, they have the ability to log on locally to DCs and shut down the system. Due to the level of abilities attributed to members of this group, by default there are no members when it is first created.

The Guests group is the least powerful group in the Builtin container, and has a membership that consists of accounts and groups for people who require minimal access, or haven't logged on using their own accounts. The Guest account and Domains Guests group are members of this group. As you'll recall, the Guest account is disabled by default, meaning that when this group is initially created it has no active users.

Because of its purpose, the Incoming Forest Trust Builders group is only available in forest root domains. Members of this group have the permission to Create Inbound Forest Trust. This permission gives them the ability to create one-way, incoming forest trusts, which can only be made between the root domains of two forests. A one-way trust means that users from one forest can access resources in another forest, but not vice versa. Because of the ability to create trusts between two domains, there are no default members in this group when it is initially created.

As its name states, the Network Configuration Operators group is used to manage changes to the network settings. The members in this group have the ability to renew and release IP addresses on servers in the domain, and modify TCP/IP settings. Because this can possibly make the server inaccessible if done incorrectly, this group has no default members, and new members should be added with caution.

Members of the Performance Monitor Users and Performance Log Users groups are used for managing performance counters on servers within the domain. Performance counters are used to monitor and measure elements of the DC, such as memory, hard disk, processor, network activity, and so on. These utilities are used by two related utilities in Windows 2000 and Windows Server 2003: System Monitor, and Performance Logs and Alerts. Both of these utilities can be accessed through the **Performance** console that is available under **Administrative Tools** in the Windows **Start** menu.

Members of the Performance Monitor Users group can use System Monitor to monitor performance counters. They can view counters locally or remotely, viewing them in a graphical or textual format. By doing so, they can determine if performance issues exist on servers within a domain.

Members of the Performance Log Users group also have the ability to manage performance counters, but can use the Performance Logs and Alerts utility to create and view

logs, and configure alerts that will notify specific users (such as administrators) if a problem exists. For example, if the amount of free hard disk space drops below a certain level, a message can be sent to a network administrator advising of the potential problem. Members of this group can also configure certain programs to run if the values of performance counters exceed or fall below a specific setting.

The Pre-Windows 2000 Compatible Access group is used for backward compatibility for older versions of Windows. Members of this group have Read access for viewing all users and groups within the domain. Depending on the security settings chosen during the installation of Active Directory, the Everyone group might be a member of this group; however, additional members can be added that are running Windows NT 4.0 or earlier if needed.

The Print Operators group allows members to perform tasks that are necessary in the administration of printers. Users who are members of this group can manage printer objects in Active Directory, and create, share, manage, and delete printers that are connected to DCs within the domain. Because adding new printers to a server might require performing certain actions like rebooting the computer, this group also has the ability to load and unload device drivers, and shut down the system. As with other groups discussed in this section, the Printer Operators group has no members added to it when initially created.

The Remote Desktop Users group allows members to connect remotely to servers in the domain. Being able to remotely log on to the DC allows them to perform actions as if they were physically sitting at the server and working on it. Because of the power this group gives members, it has no default members.

The Replicator group is one that should never have users added to it. This group is used by the File Replication Service (FRS) and provides support for replicating data; therefore, it isn't meant to have users as members.

The Server Operators group provides a great deal of power to its membership, which is why there are no default members when it is initially created. Members of this group can perform a number of administrative tasks on servers within the domain, including creating and deleting shared resources, backing up and restoring files, starting and stopping services, shutting down the system, and even formatting hard drives. Because members have the potential to cause significant damage to a DC, users should be added with caution to this group.

The Users group includes every user account that's created in the domain as part of its membership. By default, the Domain Users, Authenticated Users, and Interactive groups are members of this group. By being part of this group, members are able to run applications, access local and network printers, and perform other common tasks that are necessary for normal job functions.

Default Groups in Users Container

In addition to the groups we've discussed, up to 13 built-in groups can be located by default in the Users container, including:

- **Cert Publishers**, which gives members the ability to publish certificates
- **DnsAdmins**, which provides administrative access to the DNS Server service

- **DnsUpdateProxy**, which provides members with the ability to perform dynamic updates for other clients
- **Domain Admins**, which gives members full control of the domain
- **Domain Computers**, which includes computers that are part of the domain
- **Domain Controllers**, which includes DCs
- **Domain Guests**, which includes guests of the domain
- **Domain Users**, which includes users of the domain
- **Enterprise Admins**, which gives full control over every domain in the forest
- **Group Policy Creator Owners**, which allows members to manage group policies in the domain
- **IIS_WPG**, which is used by Internet Information Service (IIS)
- **RAS and IAS Servers**, which allows members to manage remote access
- **Schema Admins**, which allows members to modify the schema
- **Telnet Clients**, which is used for clients to connect using Telnet

The Cert Publishers group is used for digital certificates, which we discussed in Chapter 1. Although this group has no default members, when members are added to it they have the ability to publish certificates for users and computers. This allows data to be encrypted and decrypted when sent across the network.

The DnsAdmins and DnsUpdateProxy groups are installed when DNS is installed. Both of these groups have no default members, but when members are added they have abilities relating to the DNS Server service. The DnsAdmins group allows members to have administrative access to the DNS Server service. The DnsUpdateProxy group allows members to perform dynamic DNS updates on behalf of other clients, and circumvent the DACLs that typically accompany Secure Dynamic Updates.

The Domain Admins group has full control in a domain. This group becomes a member of the Administrators group on each DC, workstation, and member server when they join a domain. Because of this membership, group members have all of the rights associated with the Administrators group, including the ability to back up and restore files, change the system time, create page files, enable accounts for delegation, shut down a computer remotely, load and unload device drivers, and perform other tasks relating to administration of Active Directory and servers.

The Domain Computers and Domain Controllers groups have memberships consisting of computers in the domain. The Domain Computers group contains all workstations and servers that have joined a domain, except for DCs. When a computer account is created, the computer object automatically becomes a part of this group. Similarly, the Domain Controllers group contains all DCs that are part of the domain. Using these groups, you can set permissions and rights that apply to the computer accounts that exist within a domain.

The next two groups we'll discuss are for users who have their own accounts, or log on using a guest account. The Domain Guests group has a membership consisting of any domain guests, while the Domain Users group consists of all domain users, by default. Any user account that is created in a domain automatically becomes a member of the Domain Users group.

Enterprise Admins is a group that appears in the forest root domain, and allows members to have full control over every domain in the forest. Members of this group are automatically added to the Administrators group on every DC in every domain of the forest. As discussed earlier in this chapter, the Administrator account is a member of this group. Because of the power it gives a user, additional members should be added with caution.

The Group Policy Creator Owners group is used to manage group policy within a domain. Group policies allow you to control a user's environment. Using policies, you can control such things as the appearance and behavior of a user's desktop, and limit the user's control over his or her computer. Members of the Group Policy Creator Owners group can modify these policies. Due to the power these members have over users within a domain, the Administrator account is the only default member of this group.

The IIS_WPG group is installed when IIS is installed. IIS version 6.0 uses worker processes to serve individual DNS namespaces, and allow them to run under other identities. For example, a worker process might serve the namespace `www.syngress.com`, but could also run under another identity in the IIS_WPG group called Syngress. Because these identities need configuration to apply them to a particular namespace, there are no default members in this group.

The RAS and IAS Servers group is used for the Remote Access Service (RAS) and Internet Authentication Service (IAS), which provide remote access to a network. The members of this group have the ability to access the remote access properties of users in a domain. This allows them to assist in the management of accounts that need this access.

The Schema Admins group is another group that only appears in the forest root domain. This group allows members to modify the schema. The schema is used to define the user classes and attributes that form the backbone of the Active Directory database. As mentioned previously, the Administrator account is a default member of this group. Additional users should be added with caution, due to the widespread effect this group can have on a forest.

Creating Group Accounts

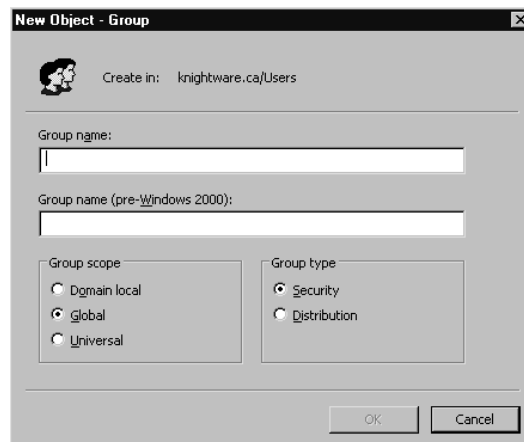
In addition to the built-in groups that are created when Active Directory and other services are installed on DCs, you can also create group accounts to suit the needs of your organization. To create group accounts, you can use either Active Directory Users and Computers or the DSADD command-line tool. Regardless of the method you use, only members of the Administrators group, Account Operators group, Domain Admins group, Enterprise Admins group, or another user or group that's been delegated authority can create a new group.

Creating Groups Using Active Directory Users and Computers

Creating new groups in Active Directory Users and Computers begins by selecting the container or OU in which you want the group to be stored. Once this is done, click **Action | New | Group**. Alternatively, you can right-click on the container, and select **New | Group**. In either case, this will open the **New Object – Group** dialog box.

The **New Object – Group** dialog box requires a minimal amount of information to create the new group. As shown in Figure 2.26, the **Group name** text box is where you enter the Active Directory name of the group. As you enter information into this field, it will also fill out the **Group name (pre-Windows 2000)** text box. This is the name that older operating systems will use to refer to the group. By default, it is the same as the **Group name**, but can be modified to any name you want within the naming rules covered previously in the chapter.

Figure 2.26 New Object Dialog Box for Creating New Groups



Below the fields designating the group's name is a section that allows you to control the scope. As discussed previously in this chapter, there are three different scopes for groups: *Domain local*, *Global*, and *Universal*. A Security group type can only be given a universal scope if the functionality level has been raised to Windows 2000 native or higher. If the functionality level is Windows 2000 mixed, then the Universal option on this dialog box will be disabled when creating a Security type group, and the only available options will be Domain local and Global.

To the right of this section is another one that allows you to specify the type of group you are creating. Two different types of groups can be created: *Security* and *Distribution*. As mentioned earlier in this chapter, security groups are used to control access, while distribution groups are used by applications for sending bulk e-mail to collections of users.

Once you have provided the information about the new group, click the **OK** button to create the group. After clicking this button, this new object will appear in the container that you initially selected to store the group. As we'll see later in this chapter, you can then modify the properties of this object to provide additional information, such as membership, descriptions, and other factors.

Creating Groups Using the DSADD Command

As we saw earlier in this chapter, the DSADD command is a useful tool for creating accounts from the command line. In addition to creating user accounts, you can also use it to create groups. Creating a new group with DSADD is done by entering the following syntax:

```
DSADD GROUP GroupDN -samid SAMName -secgrp yes | no -scope l | g | u
```

When using this command, the following parameters must be entered:

- **GroupDN** This parameter is used to specify the DN of the object being added to Active Directory and where the object will be created.
- **SAMName** This parameter is the NetBIOS name that will be used by pre-Windows 2000 computers.
- **yes | no** This *parameter* is used to specify whether the account will be created as a security or distribution group. If a security group is being created, then you would enter **yes**. If you were going to create a distribution group, then you would enter **no**.
- **l | g | u** This parameter is used to specify the scope of the group. If you were creating a domain local group, you would enter **l**. If you were creating a global group, you would enter **g**. If you were creating a universal group, you would enter **u**.

In addition to these parameters, you can also specify others by using the following syntax:

```
DSADD GROUP GroupDN [-secgrp {yes | no}] [-scope {l | g | u}] [-samid
SAMName] [-desc Description] [-memberof Group ...] [-members
Member ...] [{-s Server | -d Domain}] [-u UserName] [-p {Password
| *}] [-q] [{-uc | -uco | -uci}]
```

These options provide a variety of settings that can be applied to the group when creating it. In addition to the ones already mentioned, the meanings of these different parameters are explained in Table 2.4.

Table 2.4 DSADD Parameters for Creating Groups

Parameter	Description
-desc <i>Description</i>	Specifies the description you want to add for the group.
-memberof <i>Group ...</i>	Specifies the groups to which this new group should be added.
-members <i>Member ...</i>	Specifies the members that should be made a part of this group.
{-s Server -d Domain}	Specifies to connect to a remote server or domain. By default, the computer is connected to the DC in the logon domain.
-u <i>UserName</i>	Specifies the username to use when logging on to a remote server. By default, the username that the user is logged on to their local system is used. The following formats can be used for the <i>UserName</i> variable: Username Domain\username User principal name
-p <i>{Password *}</i>	Specifies the password to use when logging on to a remote server. If an asterisk (*) is used, you will be prompted for a password.
-q	Specifies quiet mode, and suppresses output.
{-uc -uco -uci}	Specifies Unicode to be used for input or output. If -uc is used, then input or output is to a pipe (). If -uco is used, then output is to a pipe or file. If -uci is used, then input is from a pipe or file.

Managing Group Accounts

As we've seen, the DSADD command provides a number of options for configuring new groups, while there are only a minimal number of options available when creating them through Active Directory Users and Computers. However, most of these options can be configured and reconfigured at any time by using the object's properties. By modifying the group's properties, you can perform a variety of administrative tasks related to managing group accounts.

Accessing the properties of a group account is done through **Active Directory Users and Computers**. Select the object and click **Action | Properties**. You can also right-click on the object, and select **Properties** in the context menu. Regardless of the method used to display the properties, a dialog box similar to that shown in Figure 2.27 will appear.

The dialog box contains a great deal of information about the group, and a number of options that can be configured. As seen in this figure, the title bar states the group's name followed by the word "Properties." In the case of this figure, the properties being viewed are those of a group called "Accounting Users." The dialog also provides six different tabs, which can be used for managing different facets of the account.

The **General** tab, shown in Figure 2.27, allows you to modify much of the information you provided when creating the account in Active Directory Users and Computers. On this tab, the **Group name (pre-Windows 2000)** field contains the NetBIOS name that older operating systems use to access the group. As you'll notice, this name can be modified, so it is different from the Active Directory group name. A group can have the name "Accounting Users," but have the name "Accounting" for its pre-Windows 2000 name.

Figure 2.27 General Tab in the Properties of a Group



The **Description** and **Notes** fields allow you to enter comments about this group, which can be referred to as needed. The value of the **Description** field will appear in Active Directory Users and Computers, and should describe what the group's purpose is. For example, if you were creating a special group for backing up files on a server, you could enter a description that states this purpose. The **Notes** field also allows you to enter comments, but is used for notations about the group. This can include such information as changes that were made to the account, members that were added, and so forth.


The **Group scope** section of the dialog box contains options that are used to change the scope of the group. **Domain local** groups can be converted to universal groups, if there are no other domain local groups in the membership. **Global groups** can also be converted to universal groups, providing this group isn't a member of any other global groups. Finally, **Universal** groups can be converted to global groups, if there are no universal groups that are part of this group's membership.

The **Group type** section is used to convert the group's type from being a security group to a distribution group, or vice versa. As stated previously, the **Security** option is used to create a group that controls access to resources and rights to perform certain tasks, while the **Distribution** option is used to create a group that is used for sending e-mail to collections of users. Remember that whether the group is a security or distribution group,

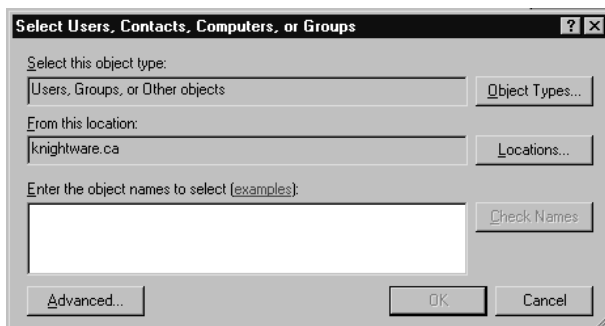
e-mail can be sent to either group type. To enable users to send e-mail to the group, you enter an e-mail address in the **E-mail** field. When a message is sent to this e-mail address, all members in the group receive a copy.

The **Members** tab is used to view current group members and add new ones. As shown in Figure 2.28, this tab provides a field that shows all current members of the group. To add new members, you click the **Add** button, which opens a dialog box that allows you to enter the names of accounts to add. Clicking **OK** in this dialog adds the name of the user, computer, or group to the list on the **Members** tab. Removing accounts from membership is also simple. Just select the account to remove from the list, and then click the **Remove** button.

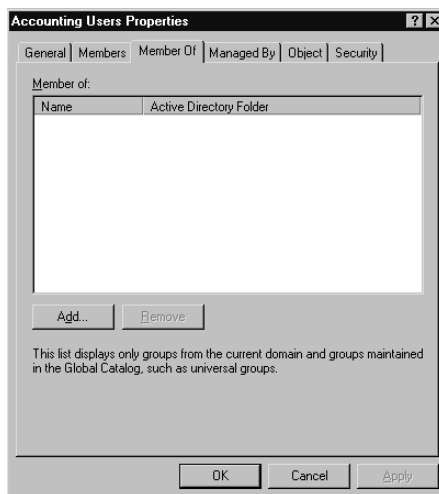
Figure 2.28 Members Tab in the Properties of a Group



By clicking the **Add** button, the dialog box shown in Figure 2.29 appears. In this dialog, you can search for the objects you want to add to the Members list. By clicking the **Object Types** button, a dialog will appear allowing you to specify the object types you want to find. In this dialog, you can click check boxes to specify whether to search for Contacts, Computers, Groups, Users, or Other objects. To limit the search to only start from a specific point in the directory structure, you can click the **Locations** button to open a dialog box showing the directory tree, where you can select the point to begin the search. Finally, the **Enter the object names to select** is where you would enter the name of the object. Upon clicking **OK**, Active Directory will use these parameters to find the object to add to the Membership list.

Figure 2.29 Select Users, Contacts, Computers, or Groups Dialog Box

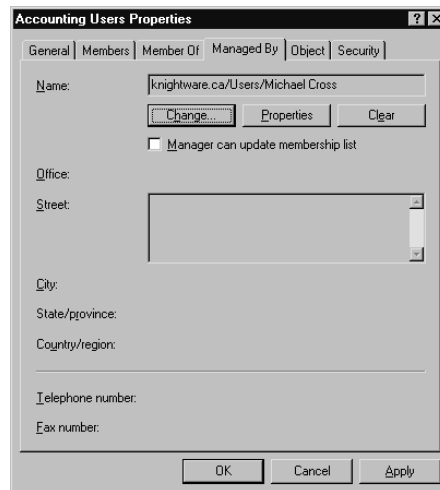
The Member Of tab, shown in Figure 2.30, is used to add this group to other existing groups in Active Directory. This tab provides a field that lists all groups to which this group belongs. To add this group to other groups, click the **Add** button to open a dialog box where you can enter the names of the groups you'd like this one to be a member of. Upon clicking **OK**, the name of the group is added to the listing on the **Member Of** tab. Removing this group from membership in another group is done by selecting that group from the list, and then clicking the **Remove** button.

Figure 2.30 Member Of Tab in the Properties of a Group

The Managed By tab is used to designate an account that is responsible for managing this group. This makes it easy for users to determine who they have to contact to request membership in the group, and how to establish contact. Checking the **Manager can update membership list** check box also allows the account listed on this tab to add and remove members from the group. To designate a manager, click the **Change** button and

specify the account. Once added, it will be displayed in the **Name** field on this tab. The properties of this account can then be viewed by clicking the **Properties** button; however, many of the commonly viewed elements of this account will automatically appear on the tab. As shown in Figure 2.31, information such as the **Office**, **Street**, **City**, **State/province**, **Country/region**, **Telephone number**, and **Fax number** will appear. To remove this account from a managerial role, click the **Clear** button.

Figure 2.31 Managed By Tab in the Properties of a Group



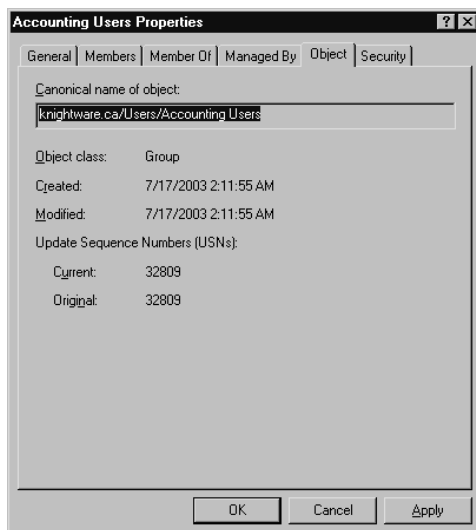
To view information about the group, you can use the **Object** tab. As shown in Figure 2.32, this tab allows you to view information about this Active Directory object. The **Canonical name of object** field displays the canonical name of the group, while the fields below this provide other data that can't be modified through the tab. The **Object class** field informs you that this is a Group, and information below this tells you when it was **Created** and last **Modified**. The **Update Sequence Numbers (USNs)** fields below this shows you what the original and current update sequence numbers for this object are, which are used by replication to ensure that all DCs have an updated copy of object information.



EXAM WARNING

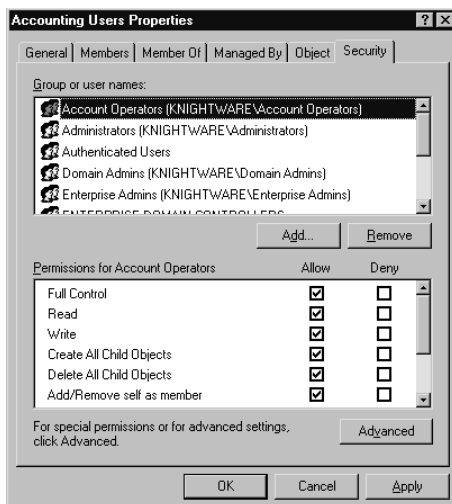
USNs are an important part of replication, and are used to indicate that changes have occurred in an object. When changes occur in an account, its USN is incremented to indicate a change has occurred.

Figure 2.32 Object Tab of Group Properties



The Security tab is used to configure the permissions that other accounts have over the group. As shown in Figure 2.33, the top pane of this tab lists users and groups with permissions over the account, while the lower pane shows the permissions of an account that's selected in the top pane. New accounts can be given access by clicking the **Add** button. Once an account is added and selected in the top pane, you enable or disable specific permissions by selecting the check box in the **Allow** or **Deny** column. Special permissions can also be set for objects by clicking the **Advanced** button. To remove an account, select the account in the top pane and click the **Remove** button.

Figure 2.33 Security Tab of Group Properties




Now that we've seen how group accounts are created and can later be managed and modified, let's put this knowledge into practice in Exercise 2.03.

EXERCISE 2.03

CREATING AND MODIFYING GROUP ACCOUNTS

1. Open **Active Directory Users and Computers** by clicking selecting **Start | Administrative Tools | Active Directory Users and Computers**.
2. When **Active Directory Users and Computers** opens, expand the console tree so that your domain and the containers within it are visible.
3. Select the **TestOU OU** from the console tree. From the **Action** menu, select **New | Group**.
4. When the **New Object – Group** dialog box appears, enter **Accounting Users** into the **Group name** text box.
5. Edit the **Group name (pre-Windows 2000)** text box so it contains the value **Accounting**.
6. Select the **Global** option under **Group scope**.
7. Select the **Security** option under **Group type**.
8. Click **OK** to create the group.
9. Right-click on the newly created **Accounting Users** group, and select **Properties** from the context menu.
10. On the **General** tab, click in the **Description** field and then enter **Group account for users in the Accounting department**.
11. On the **Members** tab, click the **Add** button.
12. When the **Select Users, Contacts, Computers, or Groups** dialog box appears, enter **John Public; Jane Doe** in the **Enter the object names to select** text box. These are the two users you created in Exercise 2.02 separated by a semicolon.
13. Click **OK** to add these users. When the **Members** tab appears again, the two users should now appear in the list of Members.
14. On the **Member Of** tab, click the **Add** button.
15. When the **Select Groups** dialog box appears, enter **Backup Operators** into the **Enter the object names to select** text box.
16. Click **OK** to make this group a member of the Backup Operators group.

17. Click **OK** to confirm these changes and exit the group Properties dialog box.
-



Working with Active Directory Computer Accounts

Computer accounts are objects that are stored in Active Directory and used to uniquely identify computers in a domain. With computer accounts, data on the computer is stored within Active Directory, allowing you to view information about the machine and use the account to set privileges on resources, install applications, and perform other actions related to its usability on the network.

Creating Computer Accounts

Computer accounts can be created in the Computers container or OUs that have been created in Active Directory. To create a new computer account, you need the same privileges as when creating user and group accounts. Only members of the Administrators group, Account Operators group, Domain Admins group, Enterprise Admins group, or a user or group that has been delegated authority can create a new account. If a user has been issued the **Add workstations to a domain** right, then he or she can create up to 10 computer accounts in a domain.



NOTE

By default, normal domain users have been delegated permission to add up to 10 computers to the domain. This default limit can be changed. For more information, see Microsoft Knowledge Base article Q251335. If the administrator has already added the computer account to Active Directory, a user can join his or her computer to the domain without using any of the 10 delegated instances mentioned previously.

There are three different methods in which a new computer account can be created:

- Joining a workstation to a domain using a user account that has the right to create a new computer account in the domain
- Creating a computer account in Active Directory Users and Computers and then joining the workstation to the domain
- Creating the computer account using DSADD and then joining the workstation to the domain

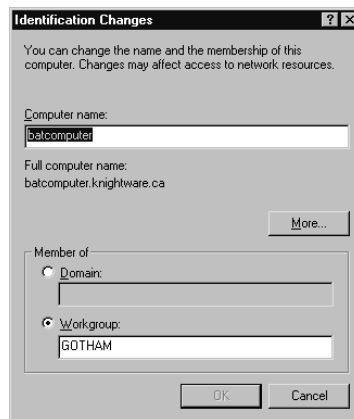
While accounts can be created before a workstation is added to the domain, only minimal information about the computer will be included in the account. Once the workstation is added to the domain, data is retrieved from the computer that is added to the account. This includes such facts as the operating system installed on the machine, the version of the operating systems, and other relevant information.

Creating Computer Accounts by Adding a Computer to a Domain

Computer accounts can be created when adding a computer to a domain. Computers can be added to a domain by using the same dialog box you use to change the computer's name. On a Windows 2000 Professional machine, this is done on the Network Identification tab of the System Properties dialog. To access this dialog, you can right-click the **My Computer** icon located on the desktop, and select **Properties** on the context menu. You can also access this dialog by double-clicking the **System** icon in **Control Panel**. Once the System Properties dialog appears, click the **Properties** button on the **Network Identification** tab.

As shown in Figure 2.34, the dialog box that appears after clicking the **Properties** button allows you to modify the name of the computer, and choose whether the computer is part of a workgroup or domain. The **Member Of** section provides two options. The **Domain** option enables a text box that allows you to provide the name of a domain this computer will join. The **Workgroup** option enables a text box that allows you to provide the name of a workgroup this computer will join. At any time, the computer can be switched from being a member of a workgroup or domain. If the computer is joining a domain where a computer account doesn't exist for this machine, then the **Computer name** field is used to specify the new Active Directory account's name.

Figure 2.34 Identification Changes Dialog Box

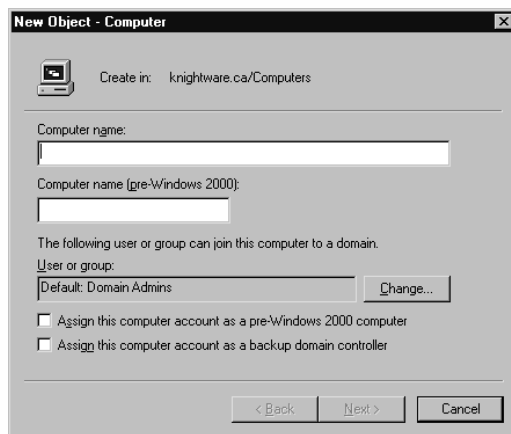


After entering the name of a domain this computer will join, click the **OK** button. The computer then proceeds to connect to a DC for the domain you are attempting to join, and if it finds one, a dialog box will be displayed asking you for the username and password of an account permitted to add workstations to the domain. Once this information is provided and you click **OK**, the username and password you provided will be authenticated and (if the user account has the necessary privileges) the workstation will be joined to the domain. If a computer account already exists for the computer, then data is retrieved and the account is updated. If no account exists, the account is created.

Creating Computer Accounts Using Active Directory Users and Computers

Computers can also be created using Active Directory Users and Computers. Right-click on the container or OU that you want to create the object in, and select **New | Computer**. Alternatively, you can select the container or OU in which you want to create the computer account, and then click **Action | New | Computer**. A dialog box similar to the one shown in Figure 2.35 will appear.

Figure 2.35 New Object – Computer Dialog Box



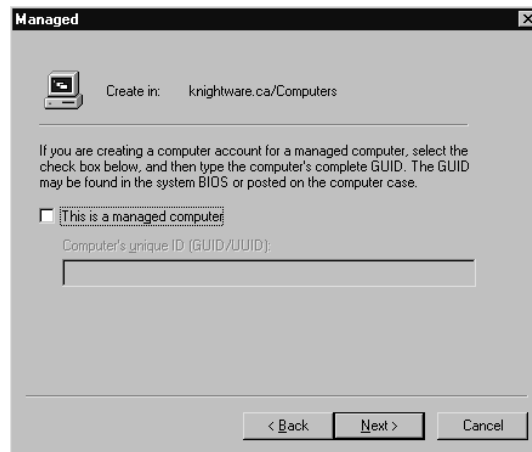
The first field on this screen is used to identify the computer. The **Computer name** text box is used to specify the name that you want this computer account to be called in Active Directory. This will be the RDN of the computer. The **Computer name (pre-Windows 2000)** text box is where you would enter the NetBIOS name of this computer, which older operating systems will use when connecting to this computer. As mentioned before, the NetBIOS name of a computer can be up to 15 characters in length. When you enter a value in the **Computer name** text box, a NetBIOS name will be suggested based on the first 15 characters of the **Computer name** field. However, this can be changed to another name.

Below this is a field that states which user or group can join the computer to the domain. As we saw in the previous section, when the computer is added to a domain, a username and password of a user account with the necessary rights is required. By default, the Domain Admins group has this ability, but this can be changed. To specify another user or group, click the **Change** button and enter the name of the user or group that should be given this privilege. The selected user or group will appear in the **User or group** field of this screen. The final options on this screen deal with older machines in a domain. The **Assign this computer account as a pre-Windows 2000 computer** designates that this machine is running an older operating system, such as Windows NT. The **Assign this computer account as a backup domain controller** specifies that this is a Windows NT BDC. Only Windows NT and newer operating systems can have accounts in Active Directory.

The remaining screens require little input. Click the **Next** button to continue to the screen that allows you to specify whether the computer is managed. A managed computer is a Remote Installation Services (RIS) client. If the **This is a managed computer** check box is checked, you must then enter the client computer's globally unique identifier (GUID). After providing this information and clicking **Next**, a screen will appear that offers the following options:

- **Any available Remote Installation Services (RIS) server**, which specifies that any RIS server can provide remote installation services to this computer.
- **The following RIS server**, which specifies that only designated RIS servers can service this computer

Figure 2.36 Managed Screen of New Object – Computer



While the screen with these RIS options will appear if the computer is managed, this will not occur if the **This is a managed computer** check box isn't checked. Upon clicking **Next**, you proceed to the final summary screen, which you can review before creating the computer account. As shown in Figure 2.37, this screen informs you of what the computer

will be called in Active Directory, and other information on options you chose during setup. Click the **Finish** button on this screen to close the wizard and create the account.

Figure 2.37 Final Screen of New Object – Computer



EXERCISE 2.04

CREATING A NEW COMPUTER ACCOUNT USING ACTIVE DIRECTORY USERS AND COMPUTERS

1. Open **Active Directory Users and Computers** by going to **Start | Administrative Tools | Active Directory Users and Computers**.
2. When Active Directory Users and Computers opens, expand the console tree so that your domain and the containers within it are visible.
3. Select the **Computers** container from the console tree. On the **Action** menu, select **New | Computer**.
4. When the **New Object-Computer** dialog box appears, enter the name of the computer you will be adding to the domain in the **Computer name** text box. Click **Next** to continue.
5. Click **Next** to go to the final screen, and then click **Finish**.

Creating Computer Accounts Using the DSADD Command

As was the case with users and groups, computer accounts can also be created using the DSADD command. The command-line method can be used in scripts to automate the addition computer objects to Active Directory. You can use the DSADD command to create computer objects using the following syntax:

```
DSADD COMPUTER ComputerDN
```

In using this command, *ComputerDN* specifies the DN of the computer that's being added. This provides information on where in the directory structure this account will be created. However, this isn't the only parameter that's available for DSADD. As shown in Table 2.5, each of these parameters provides different information that is used to set up the account. To use additional options, the following syntax can be used:

```
dsadd computer ComputerDN [-samid SAMName] [-desc Description] [-loc
Location] [-memberof GroupDN ...] [{-s Server | -d Domain}] [-u
UserName] [-p {Password | *}] [-q] [{-uc | -uco | -uci}]
```



TEST DAY TIP

Prior to Windows Server 2003, DSADD wasn't available to use with Active Directory. It is a new tool for creating user accounts, computer accounts, and group accounts in Active Directory. Depending on the type of account being created, the parameters for this tool will vary. It is important to understand how this tool works prior to taking the exam.

Table 2.5 DSADD Parameters for Creating Computers

Parameter	Description
- samid <i>SAMName</i>	Specifies the NetBIOS name used by pre-Windows 2000 computers.
- desc <i>Description</i>	Specifies a description to be used for the account.
- loc <i>Location</i>	Specifies the location of the computer.
- memberof <i>GroupDN</i>	Specifies the groups that this new computer account will be a member of.
{- s <i>Server</i> - d <i>Domain</i> }	Specifies a connection to a remote server or domain. By default, the computer is connected to the DC in the domain that the local user is logged on to.

Continued

Table 2.5 DSADD Parameters for Creating Computers

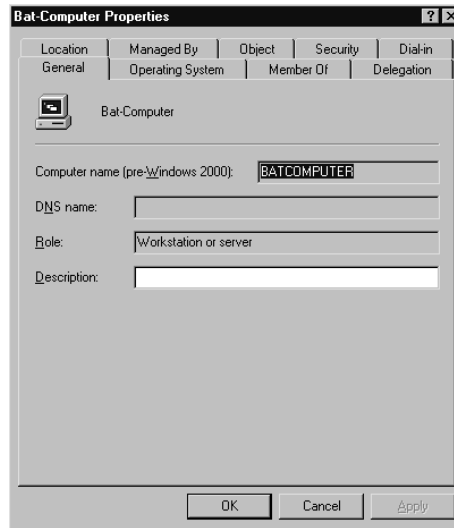
Parameter	Description
-u <i>UserName</i>	Specifies the username to use when logging on to a remote server. By default, the username that the user logged on to the local system with is used. The following formats can be used for the <i>UserName</i> variable: Username Domain\username User principal name
-p { <i>Password</i> *}	Specifies the password to use when logging on to a remote server. If an asterisk (*) is used, you will be prompted for a password.
-q	Specifies quiet mode, and suppresses output
{-uc -uco -uci}	Specifies Unicode to be used for input or output. If -uc is used, then input or output is to a pipe (). If -uco is used, then output is to a pipe or file. If -uci is used, then input is from a pipe or file.

Managing Computer Accounts

As seen previously, accounts can be administered through the properties of the object, which can be accessed using Active Directory Users and Computers. To view the properties, select the object and click **Action | Properties**. You can also right-click on the object, and select **Properties** from the context menu. Using either method, a dialog box with nine tabs will be displayed.

The **General** tab of a computer account's properties allows you to view common information about the computer. As seen in Figure 2.38, the top of the tab displays the name of the computer, which is also displayed in the title bar of the Properties dialog box. Below this, the **Computer name (pre-Windows 2000)** field displays the NetBIOS name of the computer, which is used by older computers to access this machine. The **DNS name** field supplies information on the name used by DNS to access the computer, while the **Role** field identifies the role this computer plays on the network. Finally, the **Description** field allows you to enter information that describes this computer. For example, you could specify whether it is a computer used for training purposes, development, or a particular server that provides application services (such as a Web server).

Figure 2.38 General Tab in the Properties of a Computer Account



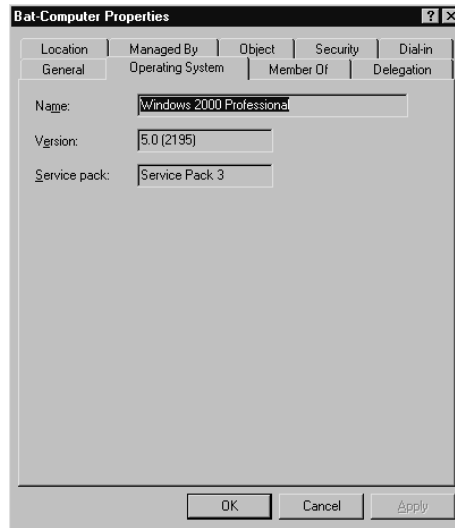
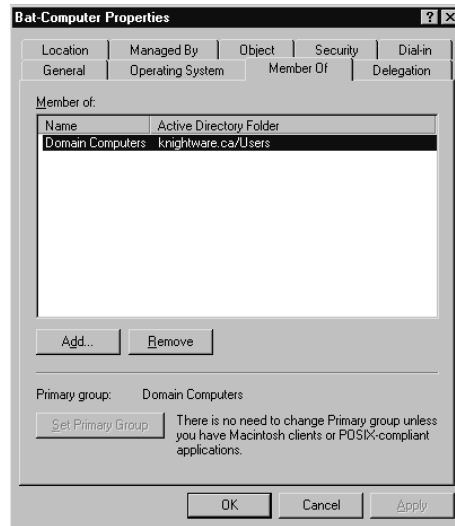
As shown in Figure 2.39, the **Operating System** tab provides information about the operating system running on the computer that has joined the domain. The **Name** field provides the name of the operating system, **Version** provides the version of the operating system, and **Service pack** displays the service pack level that has been applied to the operating system. These values are retrieved from the computer and can't be modified.



EXAM WARNING

Information on the Operating System tab, and some of the other data that appears in a computer account is retrieved from the computer when it joins the domain and is refreshed periodically thereafter. Because this information is acquired from the machine itself, it can't be manually modified through the account's properties.

The Member Of tab shown in Figure 2.40 displays existing group memberships for this computer and allows you to add the computer to groups in Active Directory. By default, it will be a member of the Domain Computers or Domain Controllers group depending on its network role. The computer account can be made a member of other groups by clicking the **Add** button. To remove the computer from a group, select the group in the list and click the **Remove** button.

Figure 2.39 Operating System Tab in the Properties of a Computer Account**Figure 2.40** Member Of Tab in the Properties of a Computer Account

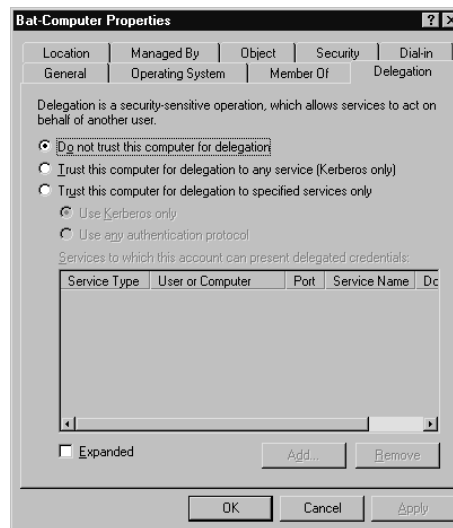
At the bottom of this tab is a section that allows you to set the primary group to which the computer belongs. By default, computers are made a member of the Domain Computers group, which is displayed in the **Primary group** field on this tab. To change the primary group, you could use the **Set Primary Group** button, but this generally isn't required. Primary groups are used by Macintosh computers and POSIX-compliant applications, and aren't required by other operating systems or applications.

The Delegation tab shown in Figure 2.41 is used to control whether services can act on behalf of another user from this computer. Using this tab, you can specify that the account can be used by specific services. By using the account's credentials, they are able to impersonate the account. This tab has three options relating to delegation:

- **Do not trust this computer for delegation** The default value, and doesn't allow the computer to be used for delegation.
- **Trust this computer for delegation for any service (Kerberos only)** Allows any service to use the computer providing Kerberos is used.
- **Trust this computer for delegation to specified services only** Only allows the services you specify to use the computer for delegation.

When the final option is selected, two additional options become available: **Use Kerberos only** and **Use any authentication protocol**. **Use Kerberos only** specifies that delegation can only be performed if Kerberos is used for authentication, while **Use any authentication protocol** allows any protocol to be used.

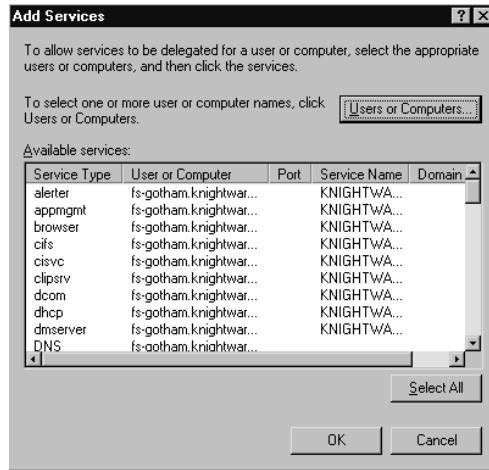
Figure 2.41 Delegation Tab in the Properties of a Computer Account



In addition to these options, the two buttons at the bottom will also be enabled. The **Add** button can be clicked to open a dialog that allows you to specify the services that can use the computer for delegation. This dialog is shown in Figure 2.42. By clicking the **Users or Computers** button, another dialog box will open, allowing you to specify the user or server that has these services associated with them. This will populate the **Available Services** field on this screen. By selecting services in this listing or alternatively clicking **Select All**, the selected services are delegated for the user or computer accounts selected.

Clicking **OK** returns you to the **Delegation** tab, where the services you selected will appear in the Delegation tab's **Services to which this account can present delegated credentials** listing. By selecting a service from this list and clicking the **Remove** button, a selected service is removed from being able to use this computer.

Figure 2.42 Add Services Dialog



The Location tab of Computer Properties allows you to provide information on the location of the computer within the organization. This tab has a single text box that allows you to enter a location name, and a button labeled **Browse**. If no locations are available to select using browse, the **Browse** button will be grayed out.

The Managed By tab is similar to the tab we saw earlier in Figure 2.28 when we discussed group accounts. This tab designates the user account of the contact person who is responsible for managing the computer object. To designate a manager, click the **Change** button. The specified account will be displayed in the **Name** field on this tab. The properties of the account in this field can be viewed by clicking the **Properties** button, and the most pertinent contact information can be viewed on the **Managed By** tab itself. To remove this account from a managerial role, click the **Clear** button.

The Object tab provides information about the object, and is similar to the tab we saw in Figure 2.32 when discussing groups. The **Canonical name of object** field on this tab shows the computer's canonical name, while the **Object class** field informs you that this is a Computer object. Below this is information on when this object was **Created** and last **Modified**. The **Update Sequence Numbers (USNs)** fields below this show you the **Original** and **Current** update sequence numbers for this object, which are used by replication to ensure that all DCs have an updated version of this object.


The **Security** tab is similar to the one in Figure 2.33 that we saw when discussing group accounts. This tab is used to configure the permissions that other accounts have in Active Directory for this computer object. As discussed previously, the top pane of this tab

lists users and groups that can be granted permissions to the account, while the lower pane shows the permissions of an account that's selected in the top pane. The **Add** button on this tab allows you to add additional accounts for which permissions can be configured. By selecting one of these accounts, you can then enable or disable specific permissions by selecting a check box in the **Allow** or **Deny** column in the lower pane. Special permissions can also be set for objects by clicking the **Advanced** button. To remove an account, the **Security** tab also provides a **Remove** button, which will remove the account that is selected in the top pane.

The final tab in a computer's properties is the **Dial-in** tab. This tab is similar to the one we saw in Figure 2.22 when we discussed user accounts. It allows you to configure settings that are used when the computer attempts to connect to the network remotely using a dial-up or VPN connection. The options that appear on this tab include:

- **Remote Access Permission (Dial-in or VPN)** This option button specifies whether the user can connect to the network via a dial-up or VPN connection. The options in this section include **Allow access**, which enables dial-in or VPN remote access; **Deny access**, which prohibits dial-in or VPN remote access; and **Control access through a Remote Access Policy**, which is the default option and specifies that a remote access policy is used to control permission for remote access.
- **Verify Caller-ID** This check box allows you to specify the telephone number that the user must be calling from in order to establish a successful connection. It requires hardware capable of detecting the number from which the user is calling.
- **Callback Options** The configuration settings in this section are **No Callback, Set by Caller (Routing and Remote Access Service Only)**, and **Always Callback To**. **No Callback** is the default option. It enables users to connect remotely and without the use of callback. When this option is set, the user will pay for any long distance charges. **Set by Caller (Routing and Remote Access Service Only)** allows the caller to specify a telephone number that the server will call back. When a remote connection is made, the user is prompted for a username and password. If successfully authenticated, the settings on this tab are checked and the user is prompted for a telephone number to be called back at. The server then disconnects and calls the user back at that number. This allows the company to pay for any long distance fees, which typically results in cost savings. **Always Callback To** is the final option. This is a security, not a cost savings, option that forces the server to call the user back at a preconfigured telephone number. Because this setting requires the user to be at that telephone number, the risk of unauthorized users attempting to connect remotely is reduced.
- **Assign a Static IP Address** This check box assigns a specific IP address to the user when he or she connects remotely.
- **Apply Static Routes** This check box places additional routes in the routing table upon connection.

- **Static Routes** This button is used to define the additional routes that will be placed in the routing table upon connection.



Managing Multiple Accounts

In the previous sections, we discussed how you can use tools for Active Directory to create and manage individual objects. In addition to creating and modifying user accounts, computer accounts, and group accounts, you can also perform actions that affect large numbers of accounts at once. In the sections that follow, we'll look at how you can manage UPNs, move objects, and how to troubleshoot problems that might result when working with accounts in Active Directory.

Implementing User Principal Name Suffixes

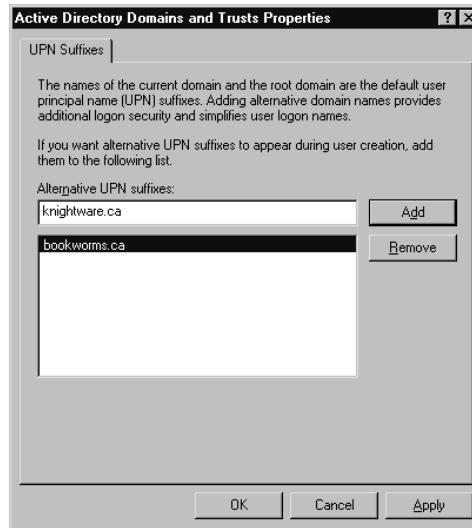
As discussed earlier in this chapter, UPNs consist of a logon account name and UPN suffix, which is connected together with an @ symbol. When combined they often look just like an e-mail address, and can in fact be used by programs to send messages to Active Directory accounts. The UPN is used when logging on to Windows 2000 and Windows Server 2003 domains from Windows 2000 or later clients.

In Active Directory, alternative UPN suffixes can be created, so the user can log on using a UPN suffix that is different from the name of the domain in which their user account resides. For example, if a user had to log on to a domain with an exceptionally long name, you could provide an alternate UPN suffix as part of the user's UPN. In doing so, the UPN is simplified, making it easier for users to enter it when logging on.

To add a UPN suffix, you must have the appropriate rights. UPN suffixes can only be added by a member of the Domains Admins group in the forest root domain, a member of the Enterprise Admins group, or a user or group that has been delegated the proper authority.

Adding UPN suffixes is done with the Active Directory Domains and Trusts console. This console is accessed from **Start | Administrative Tools | Active Directory Domains and Trusts**. As we saw in Chapter 1, it can also be started through MMC, by adding the **Active Directory Domains and Trusts** snap-in.

Once the console has opened, right-click on the **Active Directory Domains and Trusts** node in the console tree, and click **Properties** on the context menu. The properties can also be displayed by selecting the **Active Directory Domains and Trusts** node and clicking **Action | Properties**. Figure 2.43 shows the Active Directory Domains and Trusts Properties dialog box.

Figure 2.43 Active Directory Domains and Trusts Properties Dialog Box

As seen in Figure 2.43, the **UPN Suffixes** tab has a field called **Alternative UPN suffixes** where you can enter a new UPN suffix. This doesn't need to be a legitimate DNS name, which has been registered or is the name of a domain in the forest. You can create whatever name you want. Clicking the **Add** button after specifying a suffix adds the domain name you entered into the field below, which lists all alternative UPN suffixes that have been created to date. Selecting a UPN suffix from this list and clicking the **Remove** button will remove a previously created UPN suffix from the list.

EXERCISE 2.05

ADDING AND USING ALTERNATIVE UPN SUFFIXES

1. From the Windows **Start** menu, select **Administrative Tools | Active Directory Domains and Trusts**.
2. When the Active Directory Domains and Trusts console appears, select **Active Directory Domains and Trusts** from the console tree.
3. From the **Action** menu, select the **Properties** menu item.
4. When the **Active Directory Domains and Trusts Properties** dialog box appears, click in the **Alternative UPN suffixes** text box, and, enter the alternative UPN suffix you want to use (for example, eu.syngress.com).
5. Click the **Add** button. The listing should now appear in the lower pane.
6. Click **OK** to finish and close the Active Directory Domains and Trusts utility.

7. From the Windows **Start** menu, select **Administrative Tools | Active Directory Users and Computers**.
 8. When Active Directory Users and Computers opens, expand the console tree and then expand your domain. Once this is done, select the **TestOU** container.
 9. In the right pane, select the **Jane Doe** user that you created previously.
 10. From the **Action** menu, select the **Properties** menu item.
 11. When the Properties dialog box for the Jane Doe user account opens, select the **Account** tab.
 12. In the **User logon name** field, use the drop-down list to select the new UPN suffix for this user.
 13. Click **OK** to save the change and exit.
-

Moving Account Objects in Active Directory

Windows Server 2003 provides a number of tools that allow you to move objects within domains and between them. The tools that can be used for moving objects include **Active Directory Users and Computers**, and two command-line utilities. As we've seen, Active Directory Users and Computers is an MMC snap-in that allows you to interact with Active Directory through a graphical interface. The **DSMOVE** and **MOVETREE** are command-line tools that allow you to move objects by entering textual commands at the command prompt. In the sections that follow, we will look at these tools, and see how they can be used to move objects within and between domains.

Moving Objects with Active Directory Users and Computers

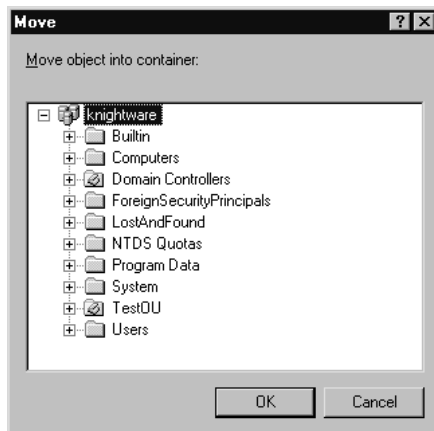
Active Directory Users and Computers can be used to move user, computer, and group accounts to other locations of the directory. With this tool, objects can be moved within a domain. It can't, however, be used to move objects to other domains.

Active Directory Users and Computers is the only tool that allows you to move accounts using a GUI. Because it's a graphical tool, you can move Active Directory objects using your mouse. Select an object by holding down your left mouse button, drag the object to a different container or OU, and release the left mouse button to drop it into the new location.

In addition, you can also move objects within the directory by right-clicking on the object, and selecting **Move** from the context menu. A dialog box will appear asking you to choose the container or OU the object should be moved to. As seen in Figure 2.44, the

Move dialog box displays a tree that represents the directory tree. By browsing the folders in this tree, you can select the container you want the object moved to, and then click **OK** to bring the move.

Figure 2.44 Move Dialog Box



When using Active Directory Users and Computers, multiple objects can be selected and moved to other locations. You can select these objects as you would files in Windows Explorer, by dragging your mouse over the objects to be moved. You can also select a series of objects by holding down the **Shift** key as you click on objects, or select a number of individual objects by holding down the **Ctrl** key as you click on them. After selecting the objects to be moved, perform the actions we just discussed to move them to another container or OU.

Moving Objects with the DSMOVE Command

As we saw in Chapter 1, DSMOVE is used to move objects within a domain, and can be used to rename objects. DSMOVE is a command-line utility that is used from the command prompt. Providing you don't need to move an object to another domain, you can use this tool to move an object to other locations in the directory tree. The syntax for using this tool is as follows:

```
DSMOVE UserDN [-newparent ParentDN] -pwd {Password/*}
```

In using this syntax, several different parameters must be entered for moving the object. The *UserDN* parameter specifies the DN of the object being moved. The **-newparent** switch indicates that you are using DSMOVE to move an object, and is used with the *ParentDN* variable to specify the DN of the new location.

To illustrate how this command is used, let's say you wanted to move an object called **BuddyJ** from the **Sales** OU in **knightware.ca** to the **Finance** OU in the same domain. To move this object, you would use the following command:


```
Dsmove CN=BuddyJ,OU=Sales,DC=knightware,DC=ca -newparent
OU=Finance,DC=knightware,DC=ca
```



TEST DAY TIP

DSMOVE is a new tool for managing Active Directory. This command-line tool will only allow objects to be moved within a domain. For moving objects to other domains, the MOVETREE command-line utility (which we'll discuss later in this chapter) must be used.

DSMOVE also provides additional parameters to perform actions such as renaming an object, or controlling the type of input and output for this command. To review these parameters, refer to the section on DSMOVE in Chapter 1.

Moving Objects with the MOVETREE Command

MOVETREE is the Active Directory Object Manager tool. In addition to other capabilities, it is a command-line tool that allows you to move objects to other domains in a forest. By using this tool, you have the freedom to move a user account, computer account, group, or OU to any location within the directory, regardless of the domain.

When an object is moved using this tool, it is first copied to the Lost and Found container before being moved to the destination domain. Objects that can't be moved remain in this container, so you can manage them as needed. Because orphaned data might reside in this domain after using MOVETREE, you should check this container after performing a move.



EXAM WARNING

The Active Directory Object Manager is MOVETREE.EXE. This tool isn't automatically installed with Active Directory and must be installed separately with the Active Directory Support Tools on the installation CD. This tool allows you to move objects from one domain to another in Active Directory.

A variety of information isn't moved with this tool. This includes data such as profiles, logon scripts, and personal information when moving user accounts. Local groups and global groups also aren't moved, but membership in these groups remains unaffected so that security involving the moved objects remains the same.

In addition to the limitations on data associated with accounts, there are also limitations when **MOVETREE** is used to move OUs between domains. When an OU is moved, group policies aren't affected, as clients will continue to receive these settings from a link to the policy in the original domain. In other words, although the OU is now in another domain, clients will connect to the Group Policy Object (GPO) that is located in the orig-

inal domain. Because this can cause performance issues, it is wise to recreate these policies in the domain where the OU has been moved, and then delete the GPO in the original domain (which is no longer needed).

As a command-line tool, **MOVETREE** requires that certain parameters be used to effectively complete operations. The syntax for **MOVETREE** is as follows, and the parameters are explained in Table 2.6.

```
MoveTree [/start | /continue | /check] [/s SrcDSA] [/d DstDSA]
        [/sdn SrcDN] [/ddn DstDN] [/u Domain\Username] [/p Password]
        [/quiet]
```

Table 2.6 Parameters for MOVETREE

Parameter	Description
/start	Specifies whether to start a move with a /check option, or with the /startnocheck option, which starts the operation without a check.
/continue	Specifies to continue the move after a failure.
/check	Specifies to check the entire tree before moving an object.
/s SrcDSA	The SrcDSA variable is used to specify the FQDN of the source server.
/d DstDSA	The DstDSA variable is used to specify the FQDN of the destination server.
/sdn SrcDN	The SrcDN variable is used to specify the source subtree's root DN.
/ddn DstDN	The DstDN variable is used to specify the destination subtree's root DN.
/u Domain\Username	Specifies the domain and user account to use for the operation.
/p Password	Specifies the password of the account to use for the operation.
/quiet	Specifies that quiet mode should be used, suppressing output.

The Active Directory Object Manager tool isn't installed with Active Directory, and thereby isn't initially available for use. **MOVETREE** is available as part of the Active Directory Support Tools on the installation CD, and can be installed through Windows Explorer. By accessing the **Support\Tools** folder on the installation CD, right-clicking on **SUPTOOLS.MSI**, and then choosing **Install** from the menu that appears, the **Windows Support Tools Setup Wizard** will start. By following the instructions in this wizard, which are detailed in Exercise 2.06, **MOVETREE** and the other support tools will be installed.

EXERCISE 2.06

INSTALLING MOVETREE WITH ACTIVE DIRECTORY SUPPORT TOOLS

1. Insert the Windows Server 2003 Server installation CD into your CD-ROM drive.
 2. From the Windows **Start** menu, select **Windows Explorer**.
 3. When Windows Explorer opens, expand the node representing your CD-ROM drive, and then expand the **Support | Tools** folder.
 4. When the contents of the **Tools** folder is displayed in the right pane, right-click on the **SUPTOOLS.MSI** file and click **Install** in the context menu.
 5. When the **Windows Support Tools Setup Wizard** appears, click **Next** to continue.
 6. On the **End User License Agreement** screen, click **I Agree** to install these tools, and then click **Next** to continue.
 7. On the **User Information** screen, enter your name in the **Name** field, and the company you work for in the **Organization** field. By default, these fields will already be completed from information acquired from Windows Server 2003 Server. Click **Next** to continue.
 8. On the **Destination Directory** screen, accept the default settings, and click **Install Now** to install the tools.
 9. A dialog box will appear showing that files are being copied to the folder specified in the **Destination Directory** screen, and being installed on Windows Server 2003. Once completed, the final screen of the wizard will appear, informing you that the tools were successfully installed. Click **Finish** to exit the wizard and complete the installation process.
-

Troubleshooting Problems with Accounts

Troubleshooting problems with accounts relies on the same methodologies and practices involved in troubleshooting other problems in Windows Server 2003. It requires an understanding of functions, configurations, and limitations. It also requires starting at the simplest possible solution for a problem and working up to the most complex. For example, if a user's account wasn't working, you wouldn't start by restoring Active Directory from a pre-

vious backup from when the user was able to log on. You might, however, check to see if the account was disabled or locked out.

It is important that you determine whether the problem exists with the user who's logging on from a computer, or with the machine itself. You'll remember that Active Directory uses both computer and user accounts. If a problem is resulting from the computer account, no user will be able to perform a certain action from the machine, regardless of what user account is used.

At times, the problems that exist in a computer account might require resetting it. If you want to reset a computer account, in **Active Directory Users and Computers**, you can right-click on the account you want to reset, and then click **Reset** from the menu that appears. After a moment, a message box will appear stating that the account was reset.

Another important part of troubleshooting is determining the scope of a problem. Is only one person experiencing a problem, or are a number of people experiencing the same difficulties? In doing so, you can determine whether the problem is with a user or computer account, or with a group of which these members are a part.

The problem might not exist in the user's account settings, but with DCs in the domain. For example, if you couldn't create security principals in Active Directory, the problem could stem from the fact that the RID Master is unavailable. The DC that has the RID operations master role allocates RIDs used for SIDs. Because SIDs can't be issued to new user accounts, computer accounts, and groups, these security principals can't be created.

You could use the command *netdom query fsmo* to identify which computers are holding single operation master roles. Once you've identified the DC serving in a particular master role, you could either repair the machine, or assign the operations master role to another machine. Before going through all this work, however, you should remember that the reason why others can't perform such actions might be because they don't have the proper rights, privileges, or permissions. In all cases, remember to start by looking at the simplest possible solution first.

Summary of Exam Objectives

In this chapter, we discussed topics relating to security principals, which are user accounts, computer accounts, and group accounts. Each security principal is assigned a security identifier (SID) when it is created. SIDs are used to uniquely identify the account, and allow the security principal to be used for authentication and access control.

In creating these accounts, we saw that there are a number of naming conventions and limitations. Each account name must be under a maximum length of characters, and refrain from including certain characters. In addition, each security principal has a relative distinguished name (RDN), distinguished name (DN), and canonical name.

User accounts, computer accounts, and group accounts can all be created using Active Directory Users and Computers, or by using the command-line utility DSADD. While these tools allow you to create new accounts, certain accounts are automatically created when Active Directory is installed. The Administrator, Guest, HelpAssistant, and SUPPORT_388945a0 user accounts are examples of these, as are the numerous built-in groups created by Active Directory upon installation.

Group accounts are collections of different accounts that are grouped together. There are two different types of groups: security groups and distribution groups. Security groups allow you to control the access permissions of users, while distribution groups are used by applications for sending e-mail to all users in the group. To further control the group, different scopes can be set to determine who can join the group and what they can access. By using groups, you can manage users as a single unit for the assignment of permissions, rights, and privileges.

Computer accounts represent workstations, DCs, and member servers. When this type of account is created and the computer joins the domain, information within the account is automatically filled in with data retrieved from the machine. Using these accounts, you can set rights, privileges, and permissions that apply to the machine, regardless of the user who is logged on at that machine.

Active Directory Trusts and Domains is a snap-in for Microsoft Management Console (MMC) that can be used to add alternate UPN suffixes. A UPN suffix combines with the user's logon name to form the user principal name (UPN). By providing an alternate UPN suffix, users can log on with a user-friendly name that is easier to remember and use.

Objects within Active Directory can be moved within the directory tree using different tools included with Windows Server 2003. Active Directory Users and Computers is a graphical tool, and DSMOVE is a command-line tool, both of which allow you to move objects within a domain. To move objects to other domains, the Active Directory Object Manager (also called MOVETREE) can be used.

Exam Objectives Fast Track

Understanding Active Directory Security Principal Accounts

- ☑ A security principal is a user account, computer account, or group account. Security principals are assigned security identifiers (SIDs) when they are created, which are used to control access to resources, and used by internal processes to identify security principals.
- ☑ WHOAMI and NTDSUTIL are tools that allow you to view and manage SIDs. WHOAMI displays information about the account, including data on SIDs for the account and groups it is a member of. NTDSUTIL is a tool used to manage SIDs, and can be used to locate and delete duplicate SIDs.
- ☑ Every security principal makes use of specific naming conventions, and has limits regarding the length and types of characters that can be part of the name. In addition to this, each security principal has a relative distinguished name, distinguished name, and canonical name .

Working with Active Directory User Accounts

- ☑ User accounts are objects that allow people and services to be authenticated and access resources.
- ☑ InetOrgPerson is a class of user account that is used when migrating to Active Directory from another directory service.
- ☑ The pre-Windows 2000 (NetBIOS) name of a user account can be up to 20 characters in length.

Working with Active Directory Group Accounts

- ☑ Group accounts are used to combine numerous accounts together as a single unit, and can be managed through Active Directory Users and Computers. With groups, you can assign rights to a group account to authorize its members to perform a certain task, assign permissions on shared resources so that all members can access the resources with the same level of permissions, or distribute bulk e-mail to all members of the group.
- ☑ Groups can be distribution groups or security groups.
- ☑ Group accounts in Active Directory can't have names that exceed 64 characters in length, and can't consist solely of numbers.

Working with Active Directory Computer Accounts

- ☑ Computer accounts can be created in Active Directory Users and Computers, by using DSADD, or by adding the workstation to a domain using a user account that has rights to create a new computer account in the domain.
- ☑ DSADD allows you to create computer accounts from the command line. DSADD can also be used to create user accounts and group accounts.
- ☑ The fully qualified domain name (FQDN) of a computer can be up to 255 characters in length. The pre-Windows 2000 (NetBIOS) name can be up to 15 characters in length.

Managing Multiple Accounts

- ☑ Active Directory Domains and Trusts can be used to create alternate UPN suffixes, which allow users to log on using a more appropriate or convenient name as part of their UPN than their domain name might have provided.
- ☑ DSMOVE is a command-line utility that allows users to move accounts from one Active Directory location within a domain to another.
- ☑ MOVETREE is a command-line tool that allows objects to be moved from one domain to another in Active Directory.

Exam Objectives

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

- Q:** I deleted a computer account by mistake, and recreated it with the same information. Now it isn't able to access the same resources it did before. Why is this?
- A:** The SID has changed. When an account is deleted and then recreated, it is given a new SID. The SID is compared to ACEs in the DACL that permits or denies access to a resource. Since the SID has changed, it now doesn't match the list, making it appear as if it is a completely different account.
- Q:** I want to view SIDs associated with the account I'm currently logged on to the computer with. Which tool can I use to do this?
- A:** WHOAMI is a command-line tool that allows you to display information about the user who is currently logged on. By typing **WHOAMI /ALL**, information about the account is displayed on the screen. Here, you can view information about the username, groups, privileges, and SIDs for the user who is currently logged on.
- Q:** I am converting a global group to another scope. When I access the options in the group's properties, I find that the option for Domain local is disabled. Do I need to change the domain functional level to have this option enabled?
- A:** No. This option is disabled because, regardless of the domain functional level, global groups can't be converted into domain local groups. Domain local and global groups can, however, be converted into universal security groups if the domain functional level is Windows 2000 native or higher.
- Q:** I want to add users to built-in groups so they can perform certain operations. In looking at the groups available in Active Directory, I see that multiple groups have the same rights and are able to perform some of the same tasks. How should I decide which groups users should be added to?
- A:** Add users to groups that will give them the necessary rights to perform tasks, but don't provide more rights than are needed to do their job. For example, if you wanted a user to perform backup and restore operations, you could add them to the Backup Operators group. Although other accounts such as the Administrators group will also allow this, it would give them considerably more rights than needed.

Q: I haven't created an account for a computer in Active Directory yet, but want to have one created when I join the computer to the domain. After changing the computer's properties to join the domain, I'm presented with a dialog box that asks for a username and password. I'm logged on to the machine using the local Administrator account for the computer. Why is this dialog box appearing?

A: The dialog box is requesting the username and password of a domain user account with appropriate rights to create the account. The Administrator account for the workstation is a local account that has nothing to do with Active Directory. It only provides administrator access to the local computer.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Understanding Active Directory Security Principal Accounts

1. You create a new user account and assign it permissions to resources. When this account is created, a SID is given to the account to uniquely identify it. When the user logs on and attempts to access one of these resources, which of the following will the SID be compared to when determining access?
 - A. Access token
 - B. SACL
 - C. DACL
 - D. SID
2. A user attempts to access a resource, and entries in the ACL are analyzed to match the SID. The system examines the entire ACL, but no match is found. Which of the following will occur?
 - A. The user will be denied access.
 - B. The user will be granted access.
 - C. The account will be disabled.
 - D. Each ACE in the ACL will be read until a match is eventually found.

3. A RID server has temporarily gone offline. During this time, you seize the RID Master role on another DC. After the original RID server becomes available again, you are concerned that duplicate SIDs might now exist for objects in Active Directory. Which of the following tools would you use to find and delete duplicates?
 - A. Active Directory Users and Computers
 - B. MOVETREE
 - C. WHOAMI
 - D. NTDSUTIL

Working with Active Directory User Accounts

4. You want to use Remote Assistance to help users with problems by connecting to their machine and taking control of it remotely. When this action is performed, which of the following accounts is automatically created and used?
 - A. HelpAssistant
 - B. Support_388945a0
 - C. Guest
 - D. InetOrgPerson
5. Your network consists of an Active Directory domain with DCs running Windows Server 2003 and another network running Novell NetWare. You are preparing to migrate the Novell NetWare network to Windows Server 2003, but want to use an account that will connect the directory services of these two networks together. Which of the following class of user account will you create?
 - A. HelpAssistant
 - B. Support_388945a0
 - C. InetOrgPerson
 - D. None. A regular user account should be created.
6. You are configuring a user account to use Terminal Services. Which of the following tabs on the user's account would you use to configure this user?
 - A. General, Address, Organization
 - B. Terminal Services Profile, Profile, Account
 - C. Environment, Sessions, Remote Control
 - D. Published Certificates, Member Of, Object

Working with Active Directory Group Accounts

7. You are creating a new group in Active Directory. In creating this group, you want users to be able to send e-mail to the group so that all members receive a copy of the message. Which type of group could be used for this purpose?
 - A. Security
 - B. Distribution
 - C. Both security and distribution
 - D. Neither security nor distribution

8. You created a new domain using DCs that are all running Windows Server 2003. The domain is part of a forest consisting of the domain you belong to, and three other domains. Each of these three other domains uses a Windows 2000 native functional level. The domain you belong to is running at the default domain functional level, and Active Directory has been configured so that all users in the domain have their own account. When adding users and groups to the groups you created, you decide that you want to change the scope of the Accounting and Sales groups. Which of the following must be taken into account when changes are made to these groups? (Choose all that apply.)
 - A. If the group has a domain local scope, it cannot contain universal groups.
 - B. Domain local groups can be converted to universal groups.
 - C. Global groups can be converted to universal groups.
 - D. None of the groups in the domain can be universal groups.

9. Your network consists of several domains in a forest that has been set to Windows Server 2003 forest functionality. You are preparing to create a group that will contain user accounts from this domain and other domains, and will be used to access resources located in several of these domains. What will be the scope of the group you create?
 - A. Universal
 - B. Global
 - C. Domain local
 - D. Distribution

10. You are an administrator in the domain dev.knightware.ca, which is a child domain beneath the forest root domain knightware.ca. You want to provide a user in this domain with the ability to create a forest trust between the dev.knightware.ca domain and the domain bookworms.ca. Which of the following built-in groups would you add this user to so he can create such a trust?

- A. Incoming Forest Trust Builders
- B. Administrator
- C. Account Operators
- D. None of the above

Working with Active Directory Computer Accounts

11. You want a new member of the IT staff to be able to create new computer accounts using Active Directory Users and Computers. Which of the following groups has the appropriate rights to create a computer account? (Choose all that apply.)
 - A. Backup Operators
 - B. Account Operators
 - C. Domain Admins
 - D. Domain Users
12. You have given a user the **Add workstations to a domain** right, so he can have his computer join the domain. In giving the user this right, how many computer accounts can the user create?
 - A. 1
 - B. 10
 - C. Unlimited
 - D. None
13. A new computer account is created in Active Directory Users and Computers for a workstation running Windows 2000 Professional. When viewing its properties, you check the Member Of tab and see that it is already included in the membership of a group. Which of the following groups is this account a member of?
 - A. Domain Users
 - B. Domain Computers
 - C. Domain Controllers
 - D. Enterprise Admins

Managing Multiple Accounts

14. Your company has an external DNS domain that is used for the company's Web site, and an internal DNS domain that is used for the network. The external DNS domain is hosted on a UNIX server that hosts the company's Web site. The Web site using this external domain name is well known, and due to its popularity, users are confused as to which domain to log on to. Which of the following can you do to allow users to log on to the internal network using the external domain name?
 - A. Use DSADD to add users to the external DNS domain.
 - B. Use MOVETREE to add users to the external DNS domain.
 - C. Use Active Directory Users and Computers to create an alternate UPN suffix.
 - D. Use Active Directory Domains and Trusts to create an alternate UPN suffix.

15. You want to move a user account from an OU located in one domain, to an OU located in another domain. Which of the following tools will you use to perform this task?
 - A. Active Directory Users and Computers
 - B. Active Directory Domains and Trusts
 - C. DSMOVE
 - D. MOVETREE

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|----------------|-----------------|
| 1. C | 9. A |
| 2. A | 10. D |
| 3. D | 11. B, C |
| 4. A | 12. B |
| 5. C | 13. B |
| 6. C | 14. D |
| 7. C | 15. D |
| 8. A, D | |

MCSE/MCSA 70-294

Creating User and Group Strategies

Exam Objectives in this Chapter:

- 3.1 Plan a security group strategy.
- 3.2 Plan a user authentication strategy.
 - 3.2.1 Plan a smart card authentication strategy.
 - 3.2.2 Create a password policy for domain users.

- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

Knowing how to create users and groups and the procedures for moving and managing them is only half the battle when it comes to effectively using these security objects on the network. The network administrator must also be able to develop strategies for authenticating the identity of anyone who uses network resources, and plan for how to use groups most effectively to provide the security and access needed.

In today's connected world, proof of your identity is often required to ensure that someone else is not trying to use your identity. It used to be that a username and password were sufficient to authenticate someone to a network. However, password authentication is only the first step in true authentication of a user's identity in today's environment. You must have a well-defined password policy, which includes account lockout, password rotation, and other options to ensure limited access to your network. In this chapter, we develop a password policy for your Windows Server 2003 network. However, sometimes passwords and password policies are not enough, and we have to take authentication to the next plateau.

Tools such as biometric devices, token devices, voice identification, and smart cards are becoming much more mainstream for user authentication as the price continues to drop and acceptance continues to rise. If you have ever seen a large data center, you have probably seen biometric tools such as thumbprint or palm scanners at entryways. Other sites use smart card readers for access to public computer kiosks. For example, Sun Microsystems requires the use of smart cards for students to sign into class each day. Each student is assigned a smart card and a four-digit personal identification number (PIN) that they must use to sign in.

In Windows Server 2003 and Windows XP, Microsoft has implemented smart card technology into the operating system as well as Active Directory to provide you with enhanced authentication abilities, and add security to your network. As a Windows Server 2003 MCSE, you are required to understand how to implement smart card technologies and manage resources using smart cards.

An effective authentication strategy works hand in hand with a security group strategy. A well-designed group strategy will ensure that users receive only the appropriate level of access to resources on the network. It will also reduce the workload of the administrator and make it easier to manage large numbers of users. Microsoft has a number of fairly complicated models that you will be expected to follow when designing group strategy in an Active Directory environment.

Creating a Password Policy for Domain Users

Since they are largely created and managed by end users, passwords have the potential to be the weakest link in any network security implementation. You can install all the high-powered firewall hardware and virtual private network (VPN) clients you like, but if your vice

president of sales uses the name of her pet St. Bernard as her password for the customer database system, all your preventative measures might be rendered useless. Since passwords are the “keys to the kingdom” of any computer system, the database that Windows Server 2003 uses to store password information will be a common attack vector for anyone attempting to hack your network. Luckily, Windows Server 2003 offers several means to secure passwords on your network. A combination of technical measures, along with a healthy dose of user training and awareness, will go a long way toward protecting the security of your network systems.

Creating an Extensive Defense Model

In modern computer security, a system administrator needs to create a security plan that uses many different mechanisms to protect a network from unauthorized access. Rather than relying solely on a hardware firewall and nothing else, *defense in depth* would also use strong passwords as well as other mechanisms on local client PCs, in the event that the firewall is compromised. The idea is to create a series of security mechanisms so that if one is circumvented, other systems and procedures are in place to help impede an attacker. Microsoft refers to this practice as an *extensive defense model*. The key points of this model are the following:

- A viable security plan needs to begin and end with user awareness, since a technical mechanism is only as effective as the extent to which the users on your network adhere to it. As an administrator, you need to educate your users about how to best protect their accounts from unauthorized attacks. This can include advice about not sharing passwords, not writing passwords down or leaving them otherwise accessible, and making sure to lock a workstation if the user needs to leave it unattended for any length of time. You can spread security awareness information via e-mail, posters in employee break areas, printed memos, or any other medium that will get your users’ attention.
- Use the system key utility (*syskey*) on all critical machines on your network. This utility, discussed later in this chapter, provides additional encryption for password information that is stored in the Security Accounts Manager (SAM) and Active Directory databases.
- Educate your users about the potential hazards of selecting the Save My Password feature or any similar feature on mission-critical applications, such as remote access or VPN clients. Make sure that users understand that the convenience of saving passwords on a local workstation is far outweighed by the potential security risk if the workstation becomes compromised.
- If you need to create one or more service accounts for applications to use, make sure that these accounts have different passwords. Otherwise, compromise of one account might leave multiple network applications open to attack.

- If you suspect that a user account has been compromised, change the password immediately. If possible, consider renaming the account entirely, since it is now a known attack vector.
- Create a password policy and/or account lockout policy that is appropriate to your organization's needs. It's important to strike a balance between security and usability in designing these types of account policies. For example, a 23-character minimum password length might seem like a good security measure on paper, but any security offered by such a decision will be rendered worthless when your users leave their impossible-to-remember 23-character passwords pasted to their monitors on sticky notes.

Strong Passwords

In discussing security awareness with your user community, one of the most critical issues to consider is that of password strength. A weak password will provide potential attackers with easy access to your users' computers, and consequently the rest of your company's network. Well-formed passwords will be significantly more difficult to decipher. Even though password-cracking utilities continue to evolve and improve, educating your users regarding the importance of strong passwords will provide additional security for your network's computing resources.

According to Microsoft, a weak password is one that contains any portion of your name, your company's name, or your network logon ID. For example, if a username was assigned as *JSmith*, and the user's password was *Smith12!@!*, that would be considered a weak password. A password that contains any complete dictionary word—*password*, *thunder*, *protocol*—is also considered weak. It should be understood that blank passwords are weak as well.

By comparison, a strong password will not contain any reference to your username, personal information, company name, or any word found in the dictionary. Strong passwords should also be at least seven characters long and contain characters from each of the following groups:

- **Uppercase letters** A, B, C ...
- **Lowercase letters** z, y, x ...
- **Numeric digits** 0, 1, 2, 3, 4, 5, 6, 7, 8, or 9
- **Non-alphanumeric characters** !, *, \$, }, etc.

Each strong password should be appreciably different from any previous passwords that the user has created. P!234abc, Q!234abc, and R!234abc, although each meeting the described password criteria, would not be considered strong passwords when viewed as a whole. To further complicate matters, an individual password can still be weak even though it meets the criteria. For example, IloveU123! would be a fairly simple password to crack, even though it possesses the length and character complexity requirements of a strong password.

System Key Utility

Most password-cracking software used in attacking computer networks attempts to target the SAM database or the Active Directory database in order to access passwords for user accounts. To secure your password information, you should use the system key utility (the `syskey.exe` file itself is located in the `%systemroot%\System32` directory by default) on every critical machine that you administer. This utility provides additional encryption for password information, which provides an extra line of defense against would-be attackers. To use this utility on a workstation or member server, you must be a member of the local Administrators group on the machine in question. If the machine is a member of a domain, remember that the Domain Admins group is a member of the local Administrators group by default. On a domain controller (DC), you need to be a member of the Domain Admins or Enterprise Admins group.



TEST DAY TIP

On workstations and member servers, password information is stored within the computer's Registry. Domain controllers integrate password information into the Active Directory database that is replicated between domain controllers.

In Exercise 3.01, we go through the steps in enabling the system key utility on a Windows Server 2003 server.

EXERCISE 3.01

CREATING A SYSTEM KEY

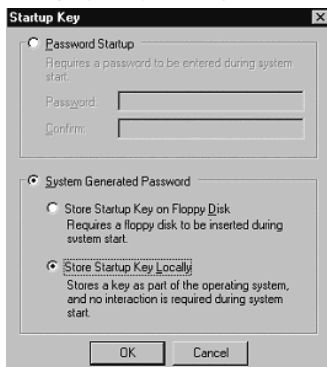
1. From the Windows Server 2003 server desktop, click **Start | Run**, then type `syskey` and click **OK**. You'll see the screen shown in Figure 3.1.

Figure 3.1 Enabling `syskey` Encryption



2. As shown in Figure 3.1, select **Encryption Enabled**, and then click **Update**.
3. Choose from the security options shown in Figure 3.2. The various options available to you are as follows:

Figure 3.2 Selecting syskey Encryption Options



- **Password Startup** This choice encrypts the account password information and stores the associated key on the local computer. You will also need to select a password that will be used to further protect the key. You'll need to enter this password during the computer's boot sequence. This is a more secure option than storing the startup key locally without requiring a password, because the password used to secure the system key isn't stored anywhere on the local computer. The drawback to this method is that an administrator must be present to enter the *syskey* password whenever the machine is rebooted, which might make this a less attractive option for a remote machine that requires frequent reboots.
- **System Generated Password, Store Startup Key on Floppy Disk** This option stores the system key on a separate diskette, which must be inserted during the system startup. This is the most secure of the three possible options, since the system key itself is not stored anywhere on the local computer and the machine will not be able to boot without the diskette that contains the system key.
- **System Generated Password, Store Startup Key Locally** This choice encrypts the SAM or Active Directory password information using a random key that's stored on the local computer. You can reboot the machine without being prompted for a password or a diskette. However, if the physical machine is compromised, the system key can be modified or destroyed. Of the three possible options when using *syskey*, this is the least secure.




EXAM WARNING

If you lose the diskette or forget the password that you created when you ran *syskey*, you won't be able to boot the computer in question without restoring the Registry or the Active Directory database from a point before you implemented *syskey*.

4. Once you have selected the option you want, click **OK** to finish encrypting the account information. You'll see the confirmation message shown in Figure 3.3.

Figure 3.3 Confirmation of *syskey* Success



Defining a Password Policy

Using Active Directory, you can create a policy to enforce consistent password standards across your entire organization. The options you can specify include: how often passwords must be changed, the number of unique passwords a user must use before being able to reuse one, and the complexity level of passwords that are acceptable on your network. Additionally, you can specify an account lockout policy that will prevent users from logging on after a specified number of incorrect logon attempts. In this section, we discuss the steps necessary to enforce password and account lockout policies on a Windows Server 2003 network.



TEST DAY TIP

To create or edit a password policy or an account lockout policy, you must be logged on as a member of the Domain Admins or Enterprise Admins group.

Applying a Password Policy

In Exercise 3.02, we discuss how to establish a password policy for your Windows Server 2003 domain.

EXERCISE 3.02

CREATING A DOMAIN PASSWORD POLICY

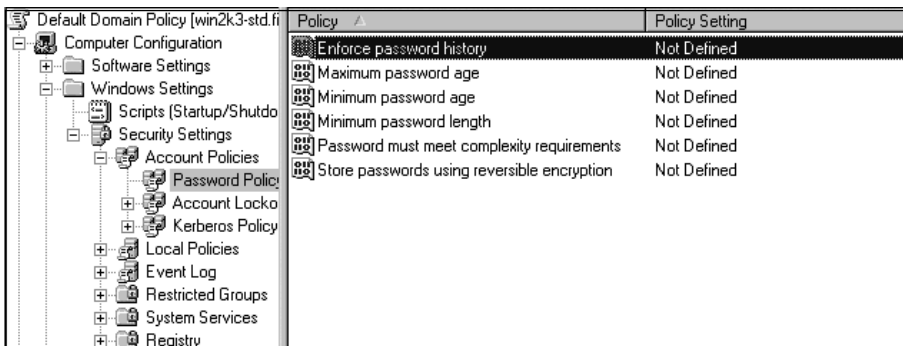
1. From the Windows Server 2003 desktop, open **Start | Administrative Tools | Active Directory Users and Computers**. Right-click the domain that you want to set a password policy for, and select **Properties**.
2. Select the **Group Policy** tab, followed by the **Default Domain Policy**, as shown in Figure 3.4. Click the **Edit** button.

Figure 3.4 The Group Policy Tab



3. Navigate to **Computer Configuration | Windows Settings | Security Settings | Account Policies | Password Policy**. You'll see the screen shown in Figure 3.5.

Figure 3.5 Configuring Password Policy Settings



Using password policies, you can configure the following settings:

- **Enforce password history** This option allows you to define the number of unique passwords that Windows will retain. This prevents users from using the same passwords again when their passwords expire. Setting this number to at least three prevents users from alternating between two passwords when they're prompted to change their passwords.
 - **Maximum password age** This setting defines how frequently Windows will prompt your users to change their passwords.
 - **Minimum password age** This setting ensures that passwords cannot be changed until they are more than a certain number of days old. This works in conjunction with the first setting by preventing users from repeatedly changing their passwords in an effort to circumvent the "Enforce password history" policy. For example, if you specify that password history is enforced and set the number of remembered passwords at 5, a user can simply change his or her password six times in rapid succession and reuse a favorite password.
 - **Minimum password length** This option dictates the shortest allowable length that a user's password can be. Remember that longer passwords are typically stronger than shorter ones. Enabling this setting also prevents users from setting a blank password.
 - **Password must meet complexity requirements** This policy setting, when activated, forces any new passwords created on your network to meet the following requirements: minimum of six characters in length, containing three of the following: uppercase letters, lowercase letters, numeric digits, and non-alphanumeric characters such as %, !, and [.
 - **Store passwords using reversible encryption** This option stores a copy of the user's password within the Active Directory database using reversible (cleartext) encryption. This is required for certain message digest functions and authentication protocols to work properly. This policy is disabled by default and should be enabled only if you are certain that your environment requires it.
4. For each item that you want to configure, right-click the item and select **Properties**. In this case, we're enforcing a password history of three passwords. In the screen shown in Figure 3.6, place a check mark next to **Define this policy setting**, and then enter the appropriate value.

Figure 3.6 Defining the Password History Policy

Modifying a Password Policy

You can modify an existing Windows Server 2003 password policy by navigating to the policy section listed in the previous exercise and making the changes. New and modified password policies are only enforced when passwords are changed. Therefore, altering password policy does not place an immediate burden on users. Typically, users won't notice the policy change until their passwords expire and they are forced to set new ones. If you need to ensure that all passwords are forced to comply with the new policy, you can set the **User must change password at next logon** option in the properties of the user accounts you administer.

Applying an Account Lockout Policy

In addition to setting password policies, you can configure your network so that user accounts will be locked out after a certain number of incorrect logon attempts. This can be a *soft lockout*, in which the account will be re-enabled after an administrator specified period of time. Alternatively, it can be a *hard lockout* in which user accounts can only be re-enabled by the manual intervention of an administrator. Before implementing an account lockout policy, you need to understand the potential implications for your network.

An account lockout policy will increase the likelihood of deterring a potential attack against your network, but you also run the risk of locking out authorized users. You need to set the lockout threshold high enough so that authorized users will not be locked out of their accounts due to simple human error, such as mistyping their passwords before they've had their morning coffee. Three to five is a common threshold. You should also remember that if a user changes his or her password on Computer A while already logged on to Computer B, the session on Computer B will continue to attempt to log on using the old (now incorrect) password. This will eventually lock out the user account and can be a common occurrence, especially in the case of service and administrative accounts. Exercise 3.03 details the necessary steps in configuring account lockout policy settings for your domain.



EXAM WARNING

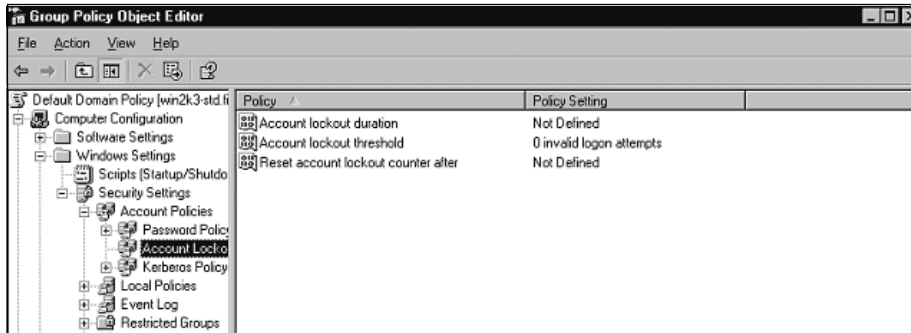
The issue of password synchronization described in the previous paragraph is not an issue for organizations that are only running Windows Server 2003 operating systems.

EXERCISE 3.03

CREATING AN ACCOUNT LOCKOUT POLICY

1. From the Windows Server 2003 desktop, click **Start | Administrative Tools | Active Directory Users and Computers**.
2. Right-click the domain you want to administer, and then select **Properties**.
3. Select the **Default Domain Policy**, and click the **Edit** button.
4. Navigate to the account lockout policy by clicking **Computer Configuration | Windows Settings | Security Settings | Account Policies | Account Lockout Policy**. You'll see the screen shown in Figure 3.7.

Figure 3.7 Account Lockout Policy Objects



Using Account Lockout Policy, you can configure the following settings:

- **Account lockout duration** This option determines the amount of time that a locked-out account will remain inaccessible. Setting this option to 0 means that the account will remain locked out until an administrator manually unlocks it. Select a lockout duration that will deter intruders without crippling your authorized users; 30 to 60 minutes is sufficient for most environments.

- **Account lockout threshold** This option determines the number of invalid logon attempts that can occur before an account will be locked out. Setting this option to 0 means that accounts on your network will never be locked out.
 - **Reset account lockout counter after** This option defines the amount of time in minutes after a bad logon attempt that the “counter” will reset. If this value is set to 45 minutes, and user *jsmith* types his password incorrectly two times before logging on successfully, his running tally of failed logon attempts will reset to 0 after 45 minutes have elapsed. Be careful not to set this option too high, or your users could lock themselves out through simple typographical errors.
5. For each item that you want to configure, right-click the item and select **Properties**. To illustrate, we create an Account lockout threshold of three invalid logon attempts. In the screen shown in Figure 3.8, place a check mark next to **Define this policy setting**, and then enter the appropriate value.

Figure 3.8 Configuring the Account Lockout Threshold



Creating User Authentication Strategies

Any well-formed security model needs to address the following three topics: authentication, authorization, and accounting (or *auditing*). Authentication deals with *who* a person is, authorization centers around *what* an authenticated user is permitted to do, and accounting/auditing is concerned with tracking *who did what* to a file, service, or other resource. Windows Server 2003 addresses all three facets of this security model, beginning with the user authentication strategies that we discuss in this chapter.

Regardless of which protocol or technical mechanism is used, all authentication schemes need to meet the same basic requirement of verifying that a user or other network object is in fact who or what it claims to be. This can include verifying a digital signature on a file or hard drive, or verifying the identity of a user or computer that is attempting to access a resource. Windows Server 2003 offers several protocols and mechanisms to perform this verification, including (but not limited to) the following:

- Kerberos
- NT LAN Manager (NTLM)
- Secure Sockets Layer/Transport Security Layer (SSL/TLS)
- Digest authentication
- Smart cards

The following sections cover the details of each authentication mechanism available with Windows Server 2003, and the appropriate use for each. The most common authentication mechanism dates back to mainframe computing, *password authentication*. This occurs when the user supplies a password to a server or host computer and the server compares the supplied password with the information that it has stored in association with the user or resource in question. In the case of users, if the two passwords match, the system permits the user to log on. Concerns regarding password authentication have largely been connected with ensuring that user passwords are not transmitted in an easily intercepted and decipherable form over a network connection. In fact, many modern password authentication schemes, such as NTLM and Kerberos, never transmit the actual user password at all.

An issue that is more difficult to address is user education. Even after a great deal of effort has been put into teaching users the importance of strong passwords, many still use relatively weak passwords for convenience. In a world of increasingly connected computing systems, the importance of creating strong password policies as part of your network's security plan cannot be overstated. To assist in this task, Windows Server 2003 allows you to establish password policies that mandate the use of strong, complex passwords, as discussed earlier in the chapter. You can also require that your users log on using smart cards, a topic that we cover in depth in a later section.

Need for Authentication

User authentication is a necessary first step within any network security infrastructure because it establishes the identity of the user. Without this key piece of information, Windows Server 2003 access control and auditing capabilities would not be able to function. Once you understand how the various authentication systems operate, you'll be able to use this information to create an effective user authentication strategy for your network. The location of your users, whether they are connected to the LAN via a high-speed network connection or a simple dial-up line, and the client and server operating systems in use throughout your organization will dictate the appropriate authentication strategy for your users.

Keep in mind as we go along that a fully functional authentication strategy will almost certainly involve a combination of the methods and protocols. A single solution will not meet the needs of an enterprise organization. Your goal as a network administrator is to create an authentication strategy that provides the optimum security for your users while allowing you to administer the network as efficiently as possible.

Single Sign-On

A key feature of Windows Server 2003 is support for single sign-on, an authentication mechanism that allows your domain users to authenticate with any computer in the domain, while only providing their logon credentials one time. This system allows network administrators to manage a single account for each user, rather than dealing with the administrative overhead of maintaining multiple user accounts for each server. It also provides greatly enhanced convenience for network users, because needing to maintain only a single password or smart card makes the network logon process much simpler.

Whether your network authentication relies on single sign-on or not, any authentication scheme is a two-step process. At the very least, the user must perform an *interactive logon* in order to access the local computer. If network access is required, *network authentication* will allow the user to access needed network services and resources. In this section, we examine both of these processes in detail.

Interactive Logon

A network user performs an interactive logon when presenting valid network credentials to the operating system of the physical computer the user is attempting to logon to—usually a desktop workstation. The logon name and password can either be a local user account or a domain account. When logging on using a local computer account, the user presents credentials that are stored in the SAM database on the local machine. Every Windows NT, 2000, and Server 2003 workstation, stand-alone server, or member server has a SAM database that is used for this purpose. Accounts stored in a SAM database can only be used for access to that specific computer.

When using a domain account, the user's logon information is authenticated against the Active Directory database. This allows the user to gain access to not only the local workstation but also to all resources he or she has been granted permission to use in the domain and any trusting domains. In this case, the user's domain account works in conjunction with the workstation's SAM database. The user authenticates to the domain, but is allowed to access the local computer because his or her domain account has the appropriate permissions and rights locally. This is typically done when the workstation joins the domain, through group membership. For example, the Domain Users global group in Active Directory is added to the local Users group on the workstation. Because all domain user accounts are members of the Domain Users group by default, they also inherit access to the local computer through this group's membership in the local Users group.

Network Authentication

Once a user has gained access to a physical workstation, it's almost inevitable that the user will require access to files, applications, or services hosted by other machines on the LAN or WAN. *Network authentication* is the mechanism that confirms the user's identity to whatever network resource the user attempts to access. Windows Server 2003 provides several mechanisms to enable this type of authentication, including Kerberos and NTLM.

Using the previous description of interactive logons, users who log on using a local computer account must provide logon credentials each time they attempt to access a network resource. This occurs because the local computer account only exists within the individual computer's SAM database rather than a centrally managed directory service like Active Directory.

On the other hand, if the user logs on using a domain account, the user's identity is proven by domain level authentication mechanisms that are automatically submitted to any network services the user is requesting to access. The mechanism used depends on the configuration of the network and the operating systems involved. Because this happens in the background, the network authentication process is transparent to users in an Active Directory environment. The network operating system handles everything behind the scenes without the need for user intervention. This feature provides the foundations for single sign-on in a Windows Server 2003 environment by allowing users to access resources in their own domains as well as other trusted domains.



TEST DAY TIP

Network authentication using a domain account can be accomplished via a username and password or with a smart card device.

Authentication Types

Windows Server 2003 offers several different authentication types to meet the needs of a diverse user base. The default authentication protocol for a homogeneous Windows 2000 or later environment is Kerberos version 5. This protocol relies on a system of tickets to verify the identity of network users, services, and devices. For Web applications and users, you can rely on the standards-based encryption offered by the SSL/TLS security protocols as well as Microsoft Digest. To provide backward compatibility for earlier versions of Microsoft operating systems, Windows Server 2003 provides support for the NTLM protocol. In this section, we examine the various authentication options available to you as a Windows administrator.

Kerberos

Within a Windows Server 2003 domain, the primary authentication protocol is Kerberos version 5. Kerberos provides thorough authentication by verifying not only the identity of

network users but also the validity of the network services themselves. This latter feature was designed to prevent users from attaching to “dummy” services created by malicious network attackers to trick users into revealing their passwords or other sensitive information. The process of verifying both the user *and* the service that the user is attempting to use is referred to as *mutual authentication*. Only network clients and servers that are running the Windows 2000, Windows Server 2003, or Windows XP Professional operating system will be able to use the Kerberos authentication protocol. When these operating systems are members of a domain, Kerberos will be enabled as their default authentication mechanism for domain-based resources. In a Windows 2000 or later Active Directory environment, pre-Windows 2000 computers that attempt to access a “Kerberized” resource will be directed to use NTLM authentication.

The Kerberos authentication mechanism relies on a key distribution center (KDC) to issue *tickets* that allow client access to network resources. Each domain controller in a Windows Server 2003 domain functions as a KDC. Network clients use DNS to locate the nearest available KDC so that they can acquire a ticket. Kerberos tickets contain cryptographic information that confirms the user’s identity to the requested service.

These tickets remain resident on the client computer system for a specific amount of time, usually 10 hours. This ticket lifetime keeps the Kerberos system from being overwhelmed, and is configurable by an administrator. If you set the threshold lower, you must ensure that your domain controllers can handle the additional load that will be placed on them. It is also important, however, not to set them too high. A ticket is good until it expires, which means that if it becomes compromised it will be valid until expiration.

Understanding the Kerberos Authentication Process

When a user enters his or her network credentials on a Kerberos-enabled system, the following steps take place. These transactions occur entirely behind the scenes. The user is only aware that he or she has entered the password or PIN number (if using a smart card) as part of a normal logon process. The following steps occur in a single domain environment:

1. Using a smart card or a username/password combination, a user authenticates to the KDC. The KDC issues a *ticket-granting ticket (TGT)* to the client system. The client retains this TGT in memory until needed.
2. When the client attempts to access a network resource, it presents its TGT to the *ticket-granting service (TGS)* on the nearest available Windows Server 2003 KDC.
3. If the user is authorized to access the service that it is requesting, the TGS issues a *service ticket* to the client.
4. The client presents the service ticket to the requested network service. Through mutual authentication, the service ticket proves the identity of the user as well as the identity of the service.

The Windows Server 2003 Kerberos authentication system can also interact with non-Microsoft Kerberos implementations such as UNIX-based Kerberos realms. In Kerberos, a

realm is similar to the concept of a domain. This “realm trust” feature allows a client in a Kerberos realm to authenticate against Active Directory to access resources, and vice versa. This interoperability allows Windows Server 2003 domain controllers to provide authentication for client systems running other types of Kerberos, including clients that are running operating systems other than Windows. It also allows Windows-based clients to access resources within a non-Windows Kerberos realm.

Secure Sockets Layer/Transport Layer Security

Any time you visit a Web site that uses an `https://` prefix instead of `http://`, you’re seeing Secure Sockets Layer (SSL) encryption in action. SSL provides encryption for other protocols such as HTTP, LDAP, and IMAP, which operate at higher layers of the protocol stack. SSL provides three major functions in encrypting TCP/IP-based traffic:

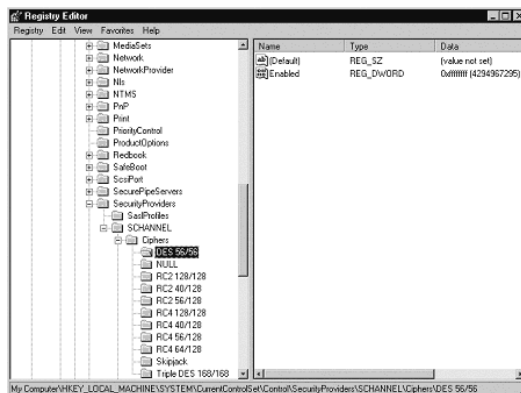
- **Server authentication** Allows a user to confirm that an Internet server is really the machine that it is claiming to be. It’s difficult to think of anyone who wouldn’t like the assurance of knowing that he or she is looking at the genuine Amazon.com site, and not a duplicate created by a hacker, before entering any credit card information.
- **Client authentication** Allows a server to confirm a client’s identity during the exchange of data. For example, this might be important for a bank that needs to transmit sensitive financial information to a server belonging to a subsidiary office. Combining server and client authentication provides a means of mutual authentication.
- **Encrypted connections** Allow all data that is sent between a client and server to be encrypted and decrypted, allowing for a high degree of confidentiality. This function also allows both parties to confirm that the data was not altered during transmission.

The Transport Layer Security (TLS) protocol is currently under development by the Internet Engineering Task Force (IETF). It will eventually replace SSL as a standard for securing Internet traffic while remaining backward compatible with earlier versions of SSL. RFC 2712 describes the way to add Kerberos functionality to the TLS suite, which will potentially allow Microsoft and other vendors to extend its use beyond LAN/WAN authentication, to use on the Internet as a whole.

SSL and TLS can use a wide range of ciphers (authentication, encryption, and/or integrity mechanisms) to allow connections with a diverse client base. You can edit the Registry in Windows Server 2003 to restrict the ciphers allowed. Within the Registry Editor on the server, browse to the following key: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Ciphers`, as shown in Figure 3.9. Each available cipher has two potential values:

- 0xffffffff (enabled)
- 0x0 (disabled)

Figure 3.9 Editing SSL/TLS Ciphers



NT LAN Manager

Versions of Windows earlier than Windows 2000 used NTLM to provide network authentication. In a Windows Server 2003 environment, NTLM is used to communicate between two computers when one or both of them is running a pre-Windows 2000 operating system. NTLM will also be used by Windows Server 2003 computers that are not members of a domain. For example, NTLM authentication would be used in the following communications:

- Windows 2000 workstations and Windows Server 2003 stand-alone servers that are participating in a workgroup instead of a domain
- Windows 2000 or Windows XP Professional computers logging on to an NT 4.0 primary domain controller (PDC) or backup domain controller (BDC)
- A Windows NT 4.0 Workstation client authenticating to an NT4.0, Windows 2000, or Windows Server 2003 domain controller
- Users in a Windows NT 4.0 domain that has a trust relationship with a Windows 2000 or Windows Server 2003 domain

NTLM encrypts user logon information by applying a mathematical function (or *hash*) to the user's password. A user's password isn't stored in the SAM or Active Directory database. Rather, the value of a hash that is generated when the user's account is first created, or the user's password is changed, is stored. If the password is less than 15 characters long, two hashes are actually stored: an NT hash and a LM hash. The LM (or LAN Manager) hash is weak and can easily be broken by password crackers. Because of this it is recommended that you configure the **Network security: Do not store LAN Manager hash value on next password change** Group Policy setting.

During logon, the domain controller sends a challenge to the client. This is a simple string of characters that the client mathematically applies to the hash value of the user's password. The result of this mathematical algorithm is a new hash that is then transmitted to the domain controller. In this way, the user's password is never actually transmitted across the network.

The domain controller also has the hash for the user's password. Moreover, it knows the challenge it sent, so it is able to perform the same calculation. It compares the hash that it mathematically calculated with the one received from the client. If they match, logon is permitted.

The NTLM hash function only exists in Windows Server 2003 for backward compatibility with earlier operating systems. Windows Server 2003 domains support both NTLM and NTLM version 2. If your network environment is exclusively running Windows 2000 or later, you might want to consider standardizing on a stronger form of authentication such as Kerberos. Using NTLM is preferable to sending authentication information using no encryption whatsoever, but NTLM has several known vulnerabilities that do not make it the best choice for network authentication if your operating system supports more advanced schemes.

Digest Authentication

Microsoft provides *digest authentication* as a means of authenticating Web applications that are running on IIS. Digest authentication uses the *Digest Access Protocol*, which is a simple challenge-response mechanism for applications that are using HTTP or Simple Authentication Security Layer (SASL) based communications. When Microsoft Digest authenticates a client, it creates a *session key* that is stored on the Web server and used to authenticate subsequent authentication requests without needing to contact a domain controller for each authentication request. Similar to NTLM, digest authentication sends user credentials across the network as an encrypted hash so that the actual password information cannot be extracted in case a malicious attacker is attempting to “sniff” the network connection. A *sniffer* is a device or software application that monitors network traffic for sensitive information, similar to a wiretap on a telephone.



NOTE

SASL is a protocol developed by Carnegie Mellon University to provide application security for client/server applications.

Before implementing digest authentication on your IIS server, you need to make sure that the following requirements have been met:

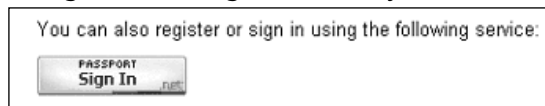
- Clients who need to access a resource or application that's secured with digest authentication need to be using Internet Explorer 5 or later.

- The server and all users attempting to log on to IIS must be members of the same domain; or domains that are connected by an appropriate trust relationship.
- The domain that the IIS server belongs to must contain a domain controller running Windows 2000 or Server 2003. The IIS server itself also needs to be running Windows 2000 or later.
- Digest authentication requires user passwords to be stored in a reversibly encrypted (cleartext) format within Active Directory. You can configure this setting from the Account tab of the user's Properties in Active Directory Users and Computers. You can also use Group Policy to enable this feature for a larger number of users. After changing this setting, your users need to change their passwords so that a reversibly encrypted hash can be created.

Passport Authentication

If you've ever logged on to the MCP Secure Site at www.microsoft.com, you've probably already seen *Passport authentication* in action. Any business that wants to provide the convenience of single sign-on to its customers can license and use Passport authentication. Passport authentication enables your company to provide a convenient means for customers to access and transact business on a given Web site. Sites that rely on Passport authentication use a centralized Passport server to authenticate users, rather than hosting and maintaining their own authentication systems. Companies can also use Passport authentication to map sign-in names to information in a sales or customer database, which can offer Passport customers a more personalized Web experience through the use of targeted ads, content, and promotional information. As Microsoft Passport has gained acceptance, the Passport sign-on logo (shown in Figure 3.10) has started to appear on more and more corporate and e-commerce Web sites.

Figure 3.10 Passport Sign-On Through www.ebay.com



From a technical perspective, Passport authentication relies on standards-based Web technologies, including SSL, HTTP redirects, and cookies. Because the technology used by Passport authentication is not proprietary, it is compatible with Microsoft Internet Explorer, Netscape Navigator, and a number of additional Web browsers, as well as operating systems such as UNIX. The single sign-on service is similar to forms-based authentication that is common throughout the Internet; it simply extends the functionality of the sign-on features to work across a distributed set of participating sites.



EXAM WARNING

Both the Internet Explorer and Netscape Navigator browsers need to be at version 4 or higher in order to access sites using Passport authentication.

Head of the Class...

Passport's Advantages for Businesses

Since Microsoft introduced the .NET Passport service in 1999, it has been responsible for authenticating more than 200 million accounts. Many prominent businesses, including McAfee, eBay, NASDAQ, and Starbucks have integrated .NET Passport into their Web authentication strategies. If you are considering integrating Passport authentication into your Web authentication plan, here are some of the advantages:

- **Single sign-in** Allows your users to sign on to the Passport site once to access information from any participating Web site. This alleviates the frustration of registering at dozens of different sites and maintaining any number of different sets of logon credentials. The Passport service allows more than 200 million Passport users quick and easy access to your site.
- **The Kids Passport service** Provides tools that help your business comply with the legal provisions of the U.S. Children's Online Privacy Protection Act (COPPA). Your company can use the Passport service to conform to the legal aspects of collecting and using children's personal information and to customize your Web site to provide age-appropriate content.
- **Maintain control of your data** Since the Passport service is simply an authentication service, your customer information and data will still be controlled in-house and is not shared with the Passport servers unless you configure your Web site to do so.

At the time of this writing, there are two fees for the use of Passport authentication: a US\$10,000 fee paid by your company on an annual basis, and a periodic testing fee of US\$1,500 per URL. The \$10,000 fee is not URL specific and covers all URLs controlled by a single company. Payment of these fees entitles your company to unlimited use of the Passport authentication service for as many URLs as you have registered for periodic testing.

Understanding Passport Authentication Security

Microsoft has created several key features within Passport authentication to ensure that the security and privacy of your customers and users can be maintained at the highest possible level. Some of the security features employed by Passport authentication are:

- The Web pages used to control the sign-in, sign-out, and registration functions are centrally hosted, rather than relying on the security mechanisms of each individual member site.
- All centrally hosted pages that are used to exchange usernames, passwords, or other credential information always use SSL encryption.
- Passport authentication-enabled sites use encrypted cookies to allow customers to access several different sites without retyping their logon information. However, a company can still opt to require users to return to the Passport sign-in screen when accessing their site for the first time.
- All cookie files related to Passport authentication use strong encryption. When you set up your site to use Passport, you receive a unique encryption key to ensure the privacy of your users' personal information.
- The central Passport servers transmit sign-in and profile information to your site in an encrypted fashion. You can then use this information to create cookies, avoiding any further client redirection to the Passport servers.
- A Web site that participates in Passport authentication will never actually receive a member's password. Authentication information is transmitted via a cookie that contains encrypted timestamps that are created when the member first signs on to Passport. The Microsoft Passport sign-out function allows users to delete any Passport-related cookies that were created on their local machines during the time that they were logged on to Microsoft Passport.
- A participating Web site only communicates directly with the central Passport server to retrieve configuration files, which are then cached locally by the individual member server. All information that is exchanged between clients and the Passport servers takes place using HTTP redirects, cookies, and encrypted queries.

Educating Users

The more highly publicized network security incidents always seem to center on a technical flaw: an overlooked patch that led to a global denial-of-service (DoS) attack, a flaw that led to the worldwide propagation of an e-mail virus, or something similar. However, many network intrusions are caused by a lack of knowledge among corporate employees. For this reason, user education is a critical component of any security plan. Make sure that your users understand the potential dangers of sharing their logon credentials with anyone

else or leaving that information in a location where others could take note of it. Your users will be far more likely to cooperate and comply with corporate security standards if they understand the reasons behind the policies and the damage that they can cause by ignoring security measures.

Security education should be both thorough and repetitive. It is not enough to simply provide security information at a new-employee orientation and never mention it again. As a network administrator, you should take steps to make sure that security awareness remains a part of your users' daily lives. You can promote this awareness through the simplest of measures, including adding a paragraph to the employee newsletter, and sending bulletins to all users when a new virus is becoming a threat. At the same time, you should avoid sending out so much information that your users become overwhelmed by it. A security bulletin that no one reads is no more useful than one that you didn't send at all. By combining user education with technical measures, such as password policies and strong network authentication, you will be well on your way to creating multiple layers of protection for your network and the data it contains.

Planning a Smart Card Authentication Strategy

Smart cards provide a portable method of providing security on a network for tasks like client authentication and securing user data. In this section, we provide an overview of smart card technology and the steps involved in using smart cards in your Windows Server 2003 network. Smart card implementations rely in part on Certificate Services, so we'll spend some time discussing the use of certificates within Windows Server 2003 as well.

Support for smart cards is a key feature within the Windows Server 2003 family. Smart cards provide tamper-resistant, safe storage for protecting your users' certificates. Certificates typically contain encryption keys that are used to encrypt and decrypt data, and can also be used in the authentication process in lieu of a standard username/password combo. Smart cards are also beneficial to use because they isolate a portion of the authentication security processes from the rest of the computer. This provides heightened security because authentication operations are performed on the smart card, which is carried by the user and not always present, rather than being carried out by processes or components that are always available within the computer or network.

Smart Cards in Action

The use of smart cards for authentication and data encryption is a new but growing trend within enterprise networks. Some cards that appear to be smart cards aren't. A good illustration of this type of card is the RSA SecurID Card from www.rsasecurity.com, shown in Figure 3.11. The RSA devices use an internal clock to generate a new synchronized number every 60 seconds. This creates a highly secure authentication method that is as portable and convenient as a common credit card or automated teller machine (ATM) card.

The number is used in the authentication process and must be verified by the server. Even if a user has the correct username and password, if this number is incorrect the user will not be able to log on. Note that this type of card does not contain a user's certificate or key. This is not the type of card-based technology Microsoft will be referring to when asking you about smart cards on the exam.

Figure 3.11 RSA SecurID Card



Not only can smart cards be used for network authentication, they can be imprinted with employee information so that they can double as identification badges. In some cases, smart card technology can also be integrated into an existing employee identification system by imprinting employee information onto a smart card. Obviously, special care needs to be taken in such implementations so that the smart card components do not become damaged through everyday use. A sample smart card/employee ID is shown in Figure 3.12.

Figure 3.12 RSA Smart Card



The advantage of this type of smart card rollout is that users do not have to remember to carry several different pieces of ID with them. The ID card that gets

Continued

them in the office door is the same one that logs them on to their computer. You'll also see smart cards that are configured as smaller "fobs," or tags, that can be carried on a keychain. Some vendors are even integrating smart card technology into handheld devices and cell phones. The smart card readers themselves can be stand-alone readers that are plugged into a serial, parallel or USB port. If a "fob" type smart card is used, it can often be directly plugged into a workstation's USB port. One such device is shown in Figure 3.13.

Figure 3.13 RSA USB Token



Using a smart card for network logons provides extremely strong authentication because it requires two factors: something the user *knows* (the PIN), and something the user *has* (the smart card itself). This system provides stronger authentication than a password alone, since a malicious user would need to have access to both the smart card and the PIN in order to impersonate a legitimate user. It's also difficult for an attacker to perform a smart card attack undetected, because the user would notice that his or her smart card was physically missing.

When to Use Smart Cards

Smart cards can provide security solutions for a number of business and technical processes within your organization. When deciding whether to add smart cards to a given system, you'll need to weigh the security benefits against the costs of deployment, both in terms of hardware and ongoing support. Smart cards can secure any of the following processes within your business:

- Using a smart card for *interactive user logons* provides enhanced security and encryption for all logon credentials.
- Requiring smart cards for *remote access logons* prevents attackers from using dial-up or VPN connections to compromise your network, even if they gain physical access to a laptop or home computer with the capability to remotely contact your network.

- *Administrator logons* are ideal candidates for smart card authentication. They have the potential to wreak far more havoc on a network installation than an account belonging to a standard network user. By requiring your administrators to use smart cards, you can greatly reduce the possibility that an attacker will be able to gain administrative access to your network.
- *Digital signing and encryption* of private user information such as e-mail and other confidential files are supported with smart cards.

Implementing Smart Cards

Using smart cards on your network involves a number of preparatory steps that we discuss in this section. First, we look at the steps involved in establishing a CA on your network, and discuss the related concepts and terminology. Next, we examine the process of establishing security permissions for users and administrators to request certificates to use with their smart cards and smart card readers. Finally, we walk through the process of setting up a smart card enrollment station to issue certificates to your end users as well as the actual procedure to issue a smart card certificate to a user on your network. We end this section with some best practices for providing technical support for the smart card users on your network.

PKI and Certificate Authorities

Smart card authentication relies on *certificates* to control which users can access the network using their smart cards. A certificate is digitally signed information that verifies the identity of a person, device, or service. Certificates can be used for a variety of functions, including Web authentication, securing e-mail, verifying the integrity of application code, and smart card authentication. The service that issues certificates is referred to as a *certificate authority (CA)*, and the person or device that receives the certificate is referred to as the *subject* of the certificate. Logon certificates typically contain quite a bit of information, including the following:

- The subject's private key
- Identifying information, such as the username or e-mail address
- The length of time that the certificate will be considered valid
- Identifier information for the CA that issued the certificate
- The digital signature of the issuing CA, which attests to the validity of the subject's key and identifying information

Many companies install their own CAs and issue certificates to their internal users to heighten the security of their network environment and save money. Using a third-party CA for a large number of certificates can be very expensive. Each certificate can cost hundreds of dollars. Using a private CA infrastructure, however, places an extra security burden on the company. If an upper-level CA becomes compromised, it can affect the validity of all certifi-

cates issued by it and subordinate CAs beneath it. Obviously, such an occurrence would be a major security breach and very expensive to recover from. It is important to ensure that there is strong security, including physical security, applied to your company's CAs.

Once a smart card certificate has expired, the user needs to obtain a new one to continue to access network resources. A CA also maintains a certificate revocation list (CRL) that can be used in case a certificate needs to be cancelled before its regular expiration date. This might result from a smart card being lost, or a certificate being compromised.

Support for smart cards is a key feature of the public key infrastructure (PKI) that's included with Windows Server 2003. You need to take several steps to prepare your Windows Server 2003 network to allow your company to use smart card devices. The first step is to install Certificate Services on at least one of your Windows Server 2003 computers. You can accomplish this from **Start | Control Panel | Add or Remove Programs | Add/Remove Windows Components | Certificate Services**. A larger discussion of configuring a Windows Server 2003 CA is beyond the scope of this chapter. For more information, view the **Certificate Services overview** article in Windows Server 2003 help.

Once you've established your server as a CA, you need to create three types of *certificate templates* to allow for smart card use on your network. Just like a document template in business application software (such as Microsoft Word), a certificate template allows multiple certificates to be created using the same basic settings. Templates are critical for this purpose because they ensure that all certificates issued will contain the same core security information and settings. Template information includes:

- Validity and renewal periods
- Whether certificates created with the template will be automatically published in Active Directory
- Allowed cryptographic service providers
- Minimum key size
- The purpose of the certificate
- What will be used for the Subject name in the certificate
- How the certificate will be issued
- Extensions associated with the template

The three security templates that you should create for use are:

- **Enrollment Agent Certificate** This template will be used to create certificates that allow a Windows Server 2003 machine to act as an enrollment station, creating certificates on behalf of smart card users who need to access the network.
- **Smart Card Logon Certificate** This template will be used to create certificates that allow your users to authenticate to Active Directory using a smart card inserted into a smart card reader.

- **Smart Card User Certificates** Like the previous template, this will be used to create certificates that allow your users to authenticate to Active Directory using a smart card inserted into a smart card reader. However, this template can also be used after login for secure e-mail.

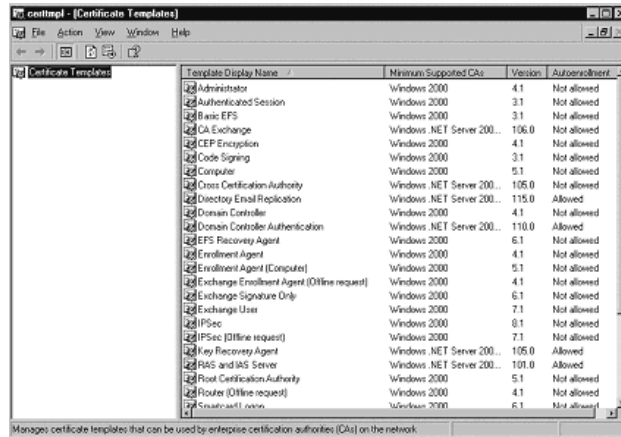
You'll be prompted to create these certificate templates automatically the first time you open the Certificate Templates snap-in. Click **Start | Run**, type **certtmpl.msc**, and click **OK**. When you're prompted to install new certificate templates, click **OK**. This step also upgrades any existing templates on your server if the machine was functioning as a CA under a previous version of Windows.

Setting Security Permissions

To implement PKI certificates, administrators and users need the appropriate permissions for the certificate templates that are installed on the CA. You can manage these permissions in the Certificate Templates snap-in, by performing the following steps:

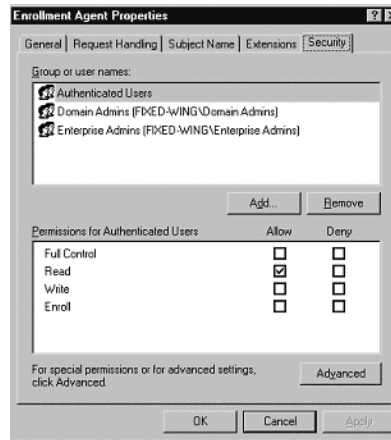
1. Open the Certificate Templates snap-in by clicking **Start | Run**, typing **certtmpl.msc**, and clicking **OK**. You'll see the screen shown in Figure 3.14.

Figure 3.14 Managing Certificate Templates



2. Right-click the certificate template whose permissions you need to change, and select **Properties**.
3. On the **Security** tab shown in Figure 3.15, add the users and groups who will need to request certificates based on this template. Under the **Allow** column, place a check mark next to the **Read** and **Enroll** permissions. Click **OK** when you've set the appropriate permissions for all necessary users and groups.

Figure 3.15 Setting Permissions for Certificate Templates



TEST DAY TIP

If you want your users to be able to retrieve and renew their certificates without any intervention on their part, you'll also need to enable autoenrollment within the properties of a certificate template. Autoenrollment is configured on the Request Handling tab. This is usually not accessible for the Smartcard Logon and Smartcard User default templates. If you need to change the default settings, you might need to right-click on each template, select Duplicate Template, and configure a copy of each template for use instead of the original. Once autoenrollment has been configured, a new permission appears on the security tab (Autoenroll) and can be granted to users and groups to allow them to autoenroll for a certificate.

Enrollment Stations

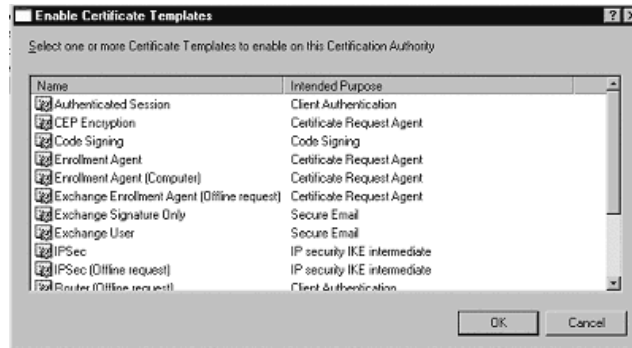
To distribute certificates and keys to your users, the CA that's included with Windows Server 2003 includes a *smart card enrollment station*. The enrollment station allows an administrator to request a smart card certificate on a user's behalf so that it can be installed onto the user's smart card. The CA signs the certificate request that's generated on behalf of the smart card user. Before your users can request certificates, you need to prepare the enrollment station to generate certificates for their use. Any machine running Windows XP Professional or Windows Server 2003 can act as an enrollment station.

Enabling Certificate Templates

To prepare your CA to issue Enrollment Agent and smart card certificates, you'll first need to enable the appropriate templates. Before you begin, make sure that the appropriate permissions have been applied to these templates. To enable the templates:

1. Open the Certification Authority snap-in by clicking **Start | Administrative Tools | Certification Authority**.
2. In the console tree, navigate to **Certification Authority | ComputerName | Certificate Templates**.
3. From the **Action** menu, click **New | Certificate Template to Issue**. You'll see the screen shown in Figure 3.16.

Figure 3.16 Enabling a Certificate Template



4. Select the **Enrollment Agent** template, and click **OK**.
5. Return to the **Action** menu, and select **New | Certificate Template to Issue**. Select one of the following options:
 - For certificates that will only be valid for user authorization, select the **Smartcard Logon** certificate template, and click **OK**.
 - For certificates that can be used both for logon and secure e-mail, click the **Smartcard User** certificate template, and then click **OK**.

Requesting an Enrollment Agent Certificate

In Exercise 3.04, we prepare a Windows Server 2003 computer to act as a smart card enrollment station. Be sure that the user account you're using has been granted the Read and Enroll permissions for the Enrollment Agent certificate template.

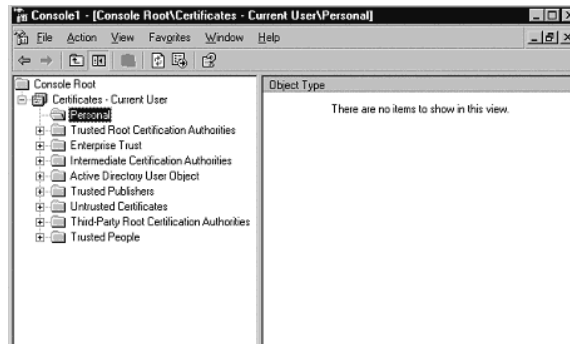
EXERCISE 3.04

CREATING A SMART CARD CERTIFICATE ENROLLMENT STATION

1. Log on to the machine as the user who will be installing the certificates.
2. Open a blank MMC console by clicking **Start | Run**, typing **mmc**, and clicking **OK**.

3. From the console window, click **File | Add/Remove Snap-in**, and then select **Add**.
4. Select the **Certificates** snap-in, click **Add**, select the option button next to **My user account**, and click **Finish**. Click **Close** and then click **OK**. You'll see the Certificates snap-in shown in Figure 3.17.

Figure 3.17 The Certificates Management Console



5. In the left pane, select **Console Root | Certificates - Current User | Personal**.
6. Click **Action | All Tasks | Request New Certificate**. Click **Next** to bypass the Welcome screen.
7. Select the **Enrollment Agent** certificate template and click **Next**.
8. Enter **Smart Card Enrollment Certificate** in the Friendly name text box and click **Next**.
9. Click **Finish** to complete the certificate request.
10. Expand **Personal** and select the **Certificates** node. In the right pane, verify that the certificate has been issued and installed.

Enrolling Users

The process of setting up your company's employees to use smart cards includes hardware, software, and administrative considerations. On the hardware side, you need to purchase and install smart card readers for all your users' workstations. Assuming that the readers are Plug-and-Play compatible, the hardware installation process should be fairly simple. Once the necessary hardware is in place, you'll use the Enrollment Station to install Smartcard Logon or User certificates in each user's smart card as well as setting initial PINs for them to use. Along

with these technical issues, you will also be required to create and document policies regarding identification requirements to receive a smart card or reset a forgotten PIN. Finally, you'll need to train your users on the new procedure to log on to a smart card-protected workstation, since the familiar **Ctrl + Alt + Del** key sequence will no longer be used.

Installing a Smart Card Reader

Most smart card readers are Plug-and-Play compatible under the Windows Server 2003 software family, so their actual installation is relatively straightforward. If you're using a reader that is not Plug-and-Play compatible or has not been tested by Microsoft, you'll need to obtain installation instructions from the card reader's manufacturer. As of this writing, the smart card readers listed in Table 3.1 are supported by Windows XP and Windows Server 2003. The corresponding device drivers will be installed on the workstation or server when the card reader has been detected by the operating system.

Table 3.1 Supported Smart Card Readers Under Windows Server 2003

Brand	Smart Card Reader	Interface	Device Driver
American Express	GCR435	USB	Grclass.sys
Bull	SmarTLP3	Serial	Bulltlp3.sys
Compaq	Serial reader	Serial	grserial.sys
Gemplus	GCR410P	Serial	Grserial.sys
Gemplus	GPR400	PCMCIA	Gpr400.sys
Gemplus	GemPC430	USB	Grclass.sys
Hewlett-Packard	ProtectTools	Serial	Scr111.sys
Litronic	220P	Serial	Lit220p.sys
Schlumberger	Reflex 20	PCMCIA	Pscr.sys
Schlumberger	Reflex 72	Serial	Scmstcs.sys
Schlumberger	Reflex Lite	Serial	Scr111.sys
SCM Microsystems	SCR111	Serial	Scr111.sys
SCM Microsystems	SCR200	Serial	Scmstcs.sys
SCM Microsystems	SCR120	PCMCIA	Pscr.sys
SCM Microsystems	SCR300	USB	Stcusb.sys
Systemneeds	External	Serial	Scr111.sys
Omnikey AG	2010	Serial	Sccmn50m.sys
Omnikey AG	2020	USB	Sccmusbm.sys
Omnikey AG	4000	PCMCIA	Cmbp0wdm.sys

To install a smart card reader on your computer, simply attach the reader to an available port, either serial or USB, or insert the reader into an available PCMCIA slot on a laptop. If

the driver for the reader is preinstalled in Windows Server 2003, the installation will take place automatically. Otherwise, the Add Hardware Wizard will prompt you for the location of the relevant software.



EXAM WARNING

If a smart card reader is attached to a serial port, it's likely that you'll need to reboot the machine before Windows will detect the device and install the appropriate driver.

Issuing Smart Card Certificates

Once you've established the appropriate security for the certificate templates and installed smart card readers on your users' workstations, you can begin the process of issuing the smart card certificates. The enrollment process must be a controlled procedure. In much the same way that employee access cards are monitored to ensure that unidentified persons do not gain physical access to your facility, smart card certificates need to be monitored to ensure that only authorized users can view network resources. In Exercise 3.05, we use the Web enrollment application to set up a smart card with a logon certificate for one of our users.

EXERCISE 3.05

SETTING UP A SMART CARD FOR USER LOGON

1. Log on to your workstation with a user account that has permissions to the appropriate certificate template in the domain where the user's account is located, and permission to enroll other users for certificates. The account used for Exercise 3.04 has these permissions.
2. Open **Internet Explorer**, and browse to **http://servername/certsrv/**, where *servername* is the name of the CA on your network.
3. Select **Request a certificate for a smart card on behalf of another user by using the smart card certificate enrollment station**.
4. A Security Warning dialog box will open asking if you'd like to install and run the Microsoft Smart Card Enrollment Control. Click **Yes**. Note that your IE security settings must be set to Low for this ActiveX control to function properly.
5. In the **Certificate Template** drop-down box select one of the following:

- **Smart Card Logon** Select this option if you want to issue a certificate that will only be valid for authenticating to the Windows domain.
 - **Smart Card User** Select this option to issue a certificate that will allow the user to use secure e-mail and log on to the Windows Server 2003 domain.
4. In the **Certification Authority** drop-down box, select the name of the CA for your domain. If there are multiple CAs in your domain, choose the one that you want to request the certificate from.
 5. In the **Cryptographic Service Provider** drop-down box, select the CSP of the smart card's manufacturer. This choice is specific to the smart card hardware you have installed. Consult the manufacturer's documentation if you are uncertain.
 6. For **Administrator Signing Certificate**, select the Enrollment Agent certificate that will sign the certificate enrollment request. This will actually display the user account that the Enrollment Agent certificate is issued to.
 7. For **User to Enroll**, click **Select User** to browse to the user account that you are associating the smart card certificate with. Insert a smart card into the smart card device attached to the system, and click **Enroll** to create a certificate for this user.
 8. You'll be prompted to set an initial PIN for the card.
 9. If another user has previously used the smart card that you're preparing, a message will appear indicating that another certificate already exists on the card. Click **Yes** to replace the existing certificate with the one you just created.
 10. On the final screen, you have the option to either view the certificate you just created or begin a new certificate request.
 11. Close your browser when you've finished so that no extraneous certificates can be created if you walk away from the enrollment station without logging off.
-

Assigning Smart Cards

Once you've preconfigured your users' smart cards, you need to establish guidelines defining how cards are assigned to users who require them. This part of your smart card deployment plan is more procedural than technical, because you need to determine accept-

able policies and service-level agreements for your smart cards and smart card readers. For example, what type of identification will you require in order for a user to obtain a smart card? Even if yours is a small organization and you recognize all of your users on sight, you should still record information from a driver's license or another piece of photo identification for auditing purposes.

Another set of issues revolves around your users' PINs. These are the equivalent of a password when using smart cards. How many unsuccessful logon attempts will you allow before locking out a smart card? Although this number will vary according to your individual business requirements, three or four PIN entry attempts are usually sufficient. Next, you need to decide whether you will allow users to reset their own PINs or if they'll need to provide personal information to, and have them reset by, the IT staff. The former option is more convenient for your user base, but that convenience will come at the expense of potential security liabilities. If user PINs need to be reset by the IT staff, decide what type of information users need to present in order to verify their identities. Document all applicable security policies, distribute them to your administration and security personnel, and make sure that your users are aware of these policies *before* they take possession of their smart cards.

Logon Procedures

To log on to a computer using a smart card, your users no longer need to enter the **Ctrl + Alt + Del** key combination. Rather, they simply insert the smart card into the smart card reader, at which point they'll be prompted to enter the PIN associated with the certificate on the card. Once the PIN is accepted, the user has access to all local and network resources to which the user's Active Directory account has been granted permissions.



TEST DAY TIP

The techniques covered here only apply to using smart card logons on computers that are attached to a domain. Third-party software is required to use smart cards on a stand-alone Windows Server 2003 computer.

Revoking Smart Cards

Along with creating policies for issuing and configuring smart cards, you should consider how your organization will handle revoking the smart card of an employee who resigns or is terminated. To be successful, this decision should be viewed as a joint effort between your company's administrative staff, such as payroll and human resources, and the IT department. Just as employees need to return ID badges and keys as part of the exit process, they should also be required to return their smart cards to the company. Whether the employee exits the company in a graceful manner or not, you should add the employee's smart card certificate(s) to your CA's CRL at the same time that you disable or delete the employee's other logon IDs and credentials. Depending on the manufacturer of the smart card, you might have an option to physically disable the smart card itself based on a serial number or other unique identifier.

Planning for Smart Card Support

Like any device or technology used to enhance network security, smart cards require you to make plans to educate your users on how to use them and provide administrative tools to support their ongoing use. First, make sure that your users understand the purpose of deploying smart cards. You'll receive a much better response if they comprehend the importance of the added security, rather than if they're simply handed a smart card and told to use it. Emphasize that the smart card is a valuable resource to protect the company and its assets, rather than simply another corporate procedure designed to annoy employees or waste their time. They should know who to call for help and technical support, and what to do if their card is lost or stolen. Maintain a printed version of this information, and distribute it to your users when they receive their smart cards. You can also publish this information on your corporate intranet, if you have one. When orienting your users to the use of smart cards, make sure you cover the following key points:

- **Protect the external smart card chip** If the chip itself becomes scratched, dented, or otherwise damaged, the smart card reader might not be able to read the data on the chip. This is similar to the magnetic strip on a credit card or an ATM card.
- **Do not bend the card** Bending the card can destroy the card's internal components. This can extend to something as simple as a user putting the smart card in a back pocket, because he or she might sit on the card and break its internal components.
- **Avoid exposing the card to extreme temperatures** Leaving a smart card on the dashboard of a car on a hot day can melt or warp the card. Extreme cold can make the card brittle and cause it to break.
- **Keep the smart card away from magnetic sources** Avoid magnetic sources such as credit card scanners at retail stores.

Along with user education, there are several settings within Active Directory Group Policy that can simplify the administration of smart cards on your network. Some of these, such as account lockout policies and restricted logon times, will impact users by default if they rely on their smart cards for domain logons. Other policy settings are specific to managing smart cards on your network. Within Group Policy, you can enable the following settings:

- **Smart card required for interactive logon** This setting prevents a user account from logging on to the network by presenting a username/password combination. When enabled, the user will only be able to authenticate by using a smart card. This provides strict security for your users; however, you should plan an alternate means of authentication in case your smart card implementation becomes unavailable.



EXAM WARNING

This policy only applies to interactive and network logons. Remote access logons are managed by separate policies on the remote access server.

- **On smart card removal** Allows you to mandate that when a user removes his or her smart card from the reader, the active session is either logged off or locked. User education is critical if you select the forced logoff option, because users need to make sure that they've saved changes to any of their documents and files before they remove their smart cards.
- **Do not allow smart card device redirection** Prevents your users from using smart cards to log on to a Terminal Services session. Set this policy if you're concerned about conserving network resources associated with your Terminal Server environment.
- **Account lockout threshold** Although this setting is not specific to smart cards, smart card PINs are more susceptible to password attacks because of their short length, so your lockout threshold settings should be adjusted accordingly.

From an administrative standpoint, there are several other important considerations in creating a support structure for smart card use. You need to identify the people within your organization who will be able to perform security-related tasks such as resetting PINs or distributing temporary cards to replace those that are lost or stolen. You also need to decide how you'll handle personnel issues such as changes in name and employment status. Finally, you'll need to carefully consider your procedures for high-level employees, traveling users, and support personnel.

Planning a Security Group Strategy

Before you can effectively start working with groups in Windows Server 2003, you need to first understand what groups are and why they are used. A group is a collection of objects (user, group, and/or computer accounts) that are managed as a single object. The objects that belong to the group are known as group members. In Windows, as with many operating systems, groups are used to simplify the administrative process of assigning permissions and rights to multiple user and computer accounts.

A set of default groups is created during the installation of Windows Server 2003 on a computer. These groups reside in the local SAM database of the stand-alone or member server, and can only be granted rights and permissions on that computer. Domain controllers also have a set of default groups. These groups reside within the Active Directory database structure and can be used throughout the domain.

You aren't limited to using the default groups. Windows Server 2003 allows you to create your own groups both at the SAM and Active Directory database levels. This book deals with Active Directory, so we will assume that you are working in a Windows Server 2003 Active Directory environment when we discuss planning group strategy.

Understanding Group Types and Scopes

In an Active Directory environment, there are two basic group characteristics: type and scope. The group type identifies the purpose of the group. There are two group types for Active Directory-based groups in Windows Server 2003:

- Distribution
- Security

Group scope refers to how the group can be used. Three group scopes can be specified for a group that resides within the Active Directory database:

- Domain local
- Global
- Universal

Security and Distribution Groups

Two types of groups can be created in Windows Server 2003:

- **Distribution groups** Distribution groups are used for distributing messages to group members. Distribution groups are used with e-mail applications, such as Microsoft Exchange. They allow a user to send e-mail to an address that is associated with the group and have it distributed to all members whose accounts are mailbox enabled. Distribution groups are not security enabled, and therefore cannot be used to assign permissions to Windows resources. The exam will not focus on distribution groups.
- **Security groups** Security groups can also be used to for the distribution of e-mail. Their main purpose, however, is to allow administrators to assign permissions and user rights to group members. Permissions can be assigned to Active Directory, file system, Registry, and printer objects. User rights include actions such as *Backup files and directories* and *Restore files and directories*, both of which are assigned to the Backup Operators group by default.

In Active Directory, security groups provide two major benefits:

- They allow you to simplify and reduce administrative requirements by assigning permissions and rights for a resource to the group rather than to each individual user that requires access. All users that are members of the group will receive the

configured permissions and rights. This is much more efficient than explicitly assigning permissions and rights to users on an individual basis. In addition, this provides you with the capability to move users in and out of groups as their job and task requirements dictate, while leaving the groups' permissions or rights unchanged.

- Security groups allow you to quickly and efficiently delegate administrative responsibilities for performing specific tasks in Active Directory. As an example, if you have a group of six help desk workers that you want to allow to reset user passwords, you can place the six users in a group and delegate the ability to the group. Again, you are able to move users in and out of the group as their job and task requirements dictate, while leaving the group's rights unchanged. This makes it very easy to simply add in other users when they require the same rights.



TEST DAY TIP

Permissions determine which users, groups, or computers can access specified resources and what they can do (read, write, execute, etc.) to that resource. By assigning these permissions to a group, instead of individual users, you can ensure that all members of the group receive the required permissions, unless they conflict with permissions that are assigned to the user explicitly or through another group.

Rights are a separate concept. Although they can relate to resources, more often rights relate to actions a user can perform involving the operating system. The right to log on locally (at the server console) or across the network are good examples. The ability to reboot a computer is also a right that must be granted to a user.

Local, Domain Local, Global, and Universal Groups

Unlike group types, which are fairly simple to understand, group scopes can be confusing to those new to working with Windows Server 2003 and Active Directory. The scope of the group identifies the extent to which the group can be applied throughout the domain or forest. Even this is not as simple as it sounds. The objects that can be members of a group, as well as the groups available, vary depending on the functional level of the domain.

Domain and Forest Functionality

Domain and forest functionality is a new feature introduced in Windows Server 2003. By having different levels of domain and forest functionality available within your Active Directory implementation, you can make different features available to your network.

If all of your network's domain controllers are Windows Server 2003 and the domain functional level is set to Windows Server 2003, then all domain features (such as the ability to rename a domain controller) become available. If your entire Active Directory forest is also set at the Windows Server 2003 functional level, then you also gain additional functionality (such as the ability to rename entire domains). In a non-upgrade environment, there are three domain functional levels available:

- **Windows 2000 mixed** This is the default domain functional level. Windows NT 4.0 BDCs, Windows 2000 domain controllers, and Windows Server 2003 domain controllers are permitted at this functional level.
- **Windows 2000 native** This is the minimum domain functional level for using universal security groups. It also enables some additional group nesting capability. This level allows for Windows 2000 and Windows Server 2003 domain controllers.
- **Windows Server 2003** This is the highest domain functional level. It provides the most features, and allows only Windows Server 2003 domain controllers.

Once you have raised the domain functional level, domain controllers running earlier operating systems cannot be used in that domain. As an example, if you raise the domain functional level to Windows Server 2003, Windows 2000 domain controllers cannot be added to the domain.

Domain Local Groups

According to Microsoft, domain local groups (DLGs) are used when assigning permissions or user rights. While we've loosely mentioned this in regard to all groups, it is this specific group scope that Microsoft wants you to use when modifying the access control list (ACL) of an object such as a file, or assigning a user right. Other groups will be added to a DLG to have their members receive the group's assigned permissions or rights.

In a Windows 2000 mixed functional level domain, domain local groups can consist of users, computers, and global groups from the domain the DLG exists in, and any trusted domain. When the functional level of the domain is raised to Windows 2000 native or Windows Server 2003, a DLG can also contain other domain local groups from its local domain, as well as universal groups. Despite the fact that this group type can contain users

and computers directly, it is important to remember that Microsoft recommends that you use it to contain other groups, which themselves contain users or computers. Specific scenarios regarding this usage are presented later in the chapter.

Global Groups

Microsoft specifies global groups (GGs) as the primary container for user and computer objects. Their models often call for grouping users according to role, function, responsibility, or department into global groups. For example, all members of the benefits team might be members of both an HR global group and a Benefits global group. Although a GG can be directly added to an ACL or assigned a user right, GGs are typically added to other groups, such as DLGs, in order to be granted access to resources.

In a Windows 2000 mixed functional level domain, a GG can contain users and computers from the same domain in which it exists. When the functional level of the domain is raised to Windows 2000 native or Windows Server 2003, a GG can also contain other GGs from its local domain.

Universal Groups

Unlike global and domain local groups, universal groups (UGs) are not stored at the domain partition level of Active Directory. They reside in the Global Catalog (GC). Because of this, adding or removing objects from a universal group triggers forest-wide replication. Microsoft recommends that other groups, and not individual user and computer accounts, be the primary members of a UG. Such members are much less likely to change.

For example, if you add a user to a UG, it triggers forest-wide replication. When you later remove that user, it again triggers forest-wide replication. However, if you add a user to a GG, which is a member of the UG, no forest-wide replication is triggered. GGs have their membership maintained at the domain level, so only domain level replication is triggered. Likewise, removing the user from the GG triggers domain level replication, not forest-wide replication.

Universal security groups do not exist in a Windows 2000 mixed functional level domain. When the functional level of the domain is raised to Windows 2000 native or Windows Server 2003, universal security groups can contain domain users, computer accounts, and global groups from any trusted domain, as well as other universal groups. Table 3.2 provides a summary of the group membership that is allowed for each domain functional level.

Table 3.2 Group Scope Behavior versus Domain Functional Level

Domain Status	Universal Group	Global Group	Domain Local Group
Windows Server 2003 or Windows 2000 native	Members can include domain user accounts, computer accounts, and global groups from any trusted domain; as well as other universal groups.	Members can include user accounts, computer accounts, and other global groups from the domain in which the global group exists.	Members can include user accounts, computer accounts, and global groups from the domain the DLG exists in or any trusted domain; universal groups; as well as other domain local groups from the domain in which the DLG exists.
Windows 2000 mixed	Universal security groups cannot be created.	Members can include user and computer accounts from the domain in which the global group exists.	Members can include user accounts, computer accounts, and global groups from the domain the DLG exists in or any trusted domain.

Changing the Group Scope

As if the concept of group scopes wasn't confusing enough, when a domain is operating at the Windows 2000 native or Windows Server 2003 functional levels, an administrator can change an existing group's scope. Universal groups can be converted to global or domain local groups, and global and domain local groups can be converted to universal groups. However, global groups cannot be converted directly to domain local groups (and vice versa).

The rules governing this are much easier than they first appear. Simply put, you cannot convert from one group type to another if the current membership of the group that is being converted is not compatible with the membership allowed for the target scope. For example, a universal security group cannot have a domain local group as a member. Therefore, if you are trying to convert a DLG into a UG, the DLG cannot have any other domain local groups as members. Table 3.3 outlines the possibilities and restrictions of changing the scope of a group.

Table 3.3 Changing the Scope of a Group

Domain Status	Universal Group	Global Group	Domain Local Group
Windows Server 2003 or Windows 2000 native	Can be changed to a global group as long as no group members are other universal groups, or user, computer, or global group accounts from any domain other than the one in which the global group will exist.	Can be converted to a domain local group with no restrictions.	Can be changed to a universal group as long as the group is not a member of any other global group.
Windows 2000 mixed	Not allowed	Not allowed	Not allowed

Configuring & Implementing...

Nesting Groups

You’ve seen how groups can have other groups as members. This concept is known as *group nesting*. Groups can be nested to help reduce management overhead. The type of nesting you can perform is determined by the domain’s functional level. If the domain functional level is set to Windows 2000 native or Windows Server 2003, the following groups have additional nesting capability:

- **Domain local groups** These groups can have other domain local groups from the same domain as well as universal groups nested in their group membership.
- **Global groups** These groups can have other global groups from the same domain nested in their membership.
- **Universal groups** These groups can have global groups from any trusted domain and other universal groups nested in their membership.

The nesting occurs in addition to the basic security group memberships that are permitted at the Windows 2000 mixed functional level.

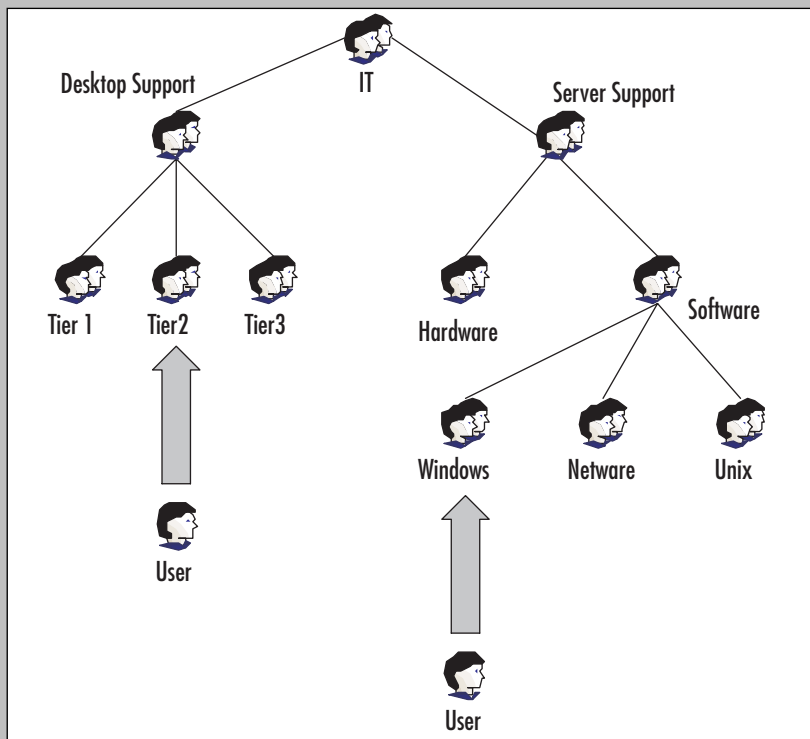
Group nesting is pictured in Figure 3.18. If a user moves from a tier 2 position in desktop support to the Windows server team, removing the user from one group and adding the user to another group automatically adjusts the permissions and rights the user is receiving from several groups. In the first example, the user is a member of the Tier2 global group, which is itself a member of the Desktop Support global group. This group is in turn nested in the IT global group. Thus, any per-

Continued

missions or rights granted to the IT, Desktop Support, and Tier2 groups will be given to the user.

When the user's account is moved, the user becomes a member of the Windows global group. The move will cause the user to lose all of the permissions and rights that were granted from the Tier2 and Desktop Support global groups. The Windows group is a member of the Software global group, which is nested in the Server Support global group. Finally, Server Support is a member of the IT global group. The user's new group membership will bring all of the permissions and rights granted to the IT, Server Support, Software, and Windows global groups.

Figure 3.18 Using Group Nesting



Security Group Best Practices

Microsoft has a number of different recommended methods for using groups in a domain environment. You should expect to be asked a number of complex questions about the appropriate use of groups. Most of their recommendations fall into one of two models:

- A single domain forest
- A multiple domain forest

Designing a Group Strategy for a Single Domain Forest


AGDLP. This simple acronym sums up everything you need to remember for the use of groups in a single domain forest environment. Each of the letters has a specific meaning:

- **A** Accounts
- **G** Global groups
- **DL** Domain local groups
- **P** Permissions

The acronym can be read as: **A**ccounts (user and computer objects) are placed into **G**lobal groups, which are placed into **D**omain **L**ocal groups, which are added to **A**CLs and granted **P**ermissions to a resource. Earlier we mentioned that Microsoft wants you to place related users (such as all HR personnel) into global groups, and use domain local groups for granting access to resources. AGDLP, their simplest model, makes these recommendations exceptionally clear and easy to remember.

Consider an environment like the one we mentioned previously. You have a new employee who is joining the benefits team within a company. The new user needs to access to both benefits-related resources and all general HR resources. Therefore, you add the user into both the Benefits and HR global groups. These global groups are themselves members of domain local groups, one of which is illustrated in Figure 3.19. The HR global group is a member of the HR_Print domain local group. This group is used to grant access to the general printers that all members of the HR department are allowed to use.

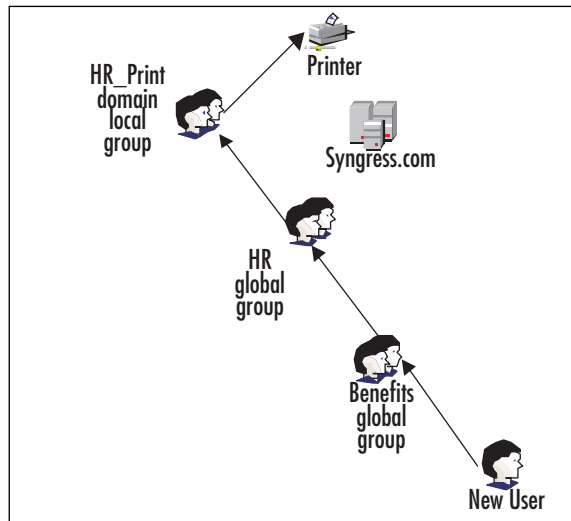
Figure 3.19 AGDLP in a Single Domain Forest



When the domain functional level is elevated to Windows 2000 native or Windows Server 2003, Microsoft specifies a new group model, AGGDLP. The meaning of the letters does not change. Therefore, this model means: **A**ccounts are placed into **G**lobal groups that can be placed into other **G**lobal groups and/or **D**omain **L**ocal groups, which are added to ACLs and granted **P**ermissions to resources. This can make a huge difference, because it allows you to potentially reduce the number of groups that you have to add a new user to.

Consider the example used previously. If you nest the Benefits global group into the HR global group, you gain a tremendous advantage. When a new user joins the benefits team, you only have to add that user's account to a single user group, Benefits. Because this group is also a member of the HR global group, the user will receive all of the permissions and rights assignments associated with both groups. Figure 3.20 shows the AGGDLP model.

Figure 3.20 AGGDLP in a Single Domain Forest



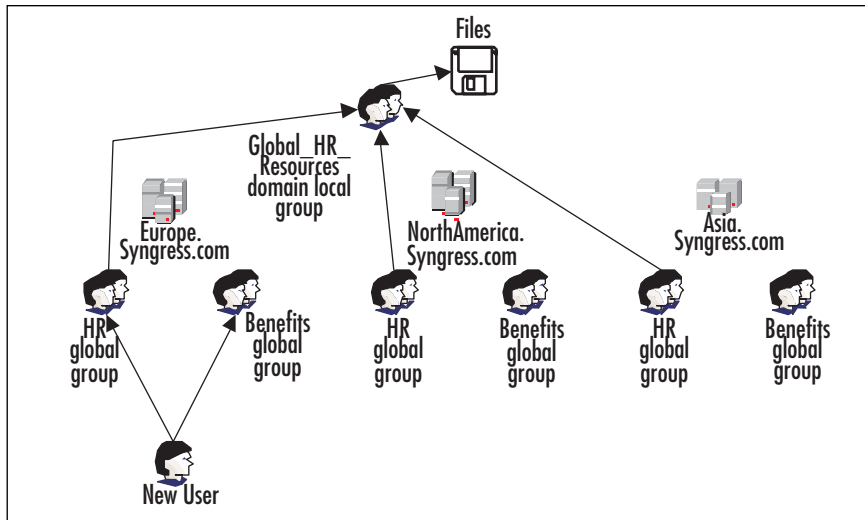
Designing a Group Strategy for a Multiple Domain Forest

These existing models can also be extended to a multiple domain forest. In a Windows 2000 mixed functional level domain, it takes quite a few resource assignments to grant permissions across domains. Extending the previous example, two additional domains will be added. Each domain is for a different region of the world, and each has an HR department. The company needs all HR employees to be able to access files that are located in the North America office. Because the domain is at the Windows 2000 mixed functional level, the AGDLP model is used.

Again, a new user joins the benefits team, this time in the Europe domain. The user is added to the Benefits and HR global groups in the Europe domain. The HR global group

in each domain has also been added to the Global_HR_Resources domain local group in the North America Domain. The Global_HR_Resources DLG has been granted the necessary permissions on the ACL for the files. Because all HR employees are (directly or indirectly) members of the HR global group in their domain, and each HR global group is a member of the Global_HR_Resources domain local group, they all have permission to access the required files. These complex relationships are shown in Figure 3.21.

Figure 3.21 AGDLP in a Multiple Domain Forest

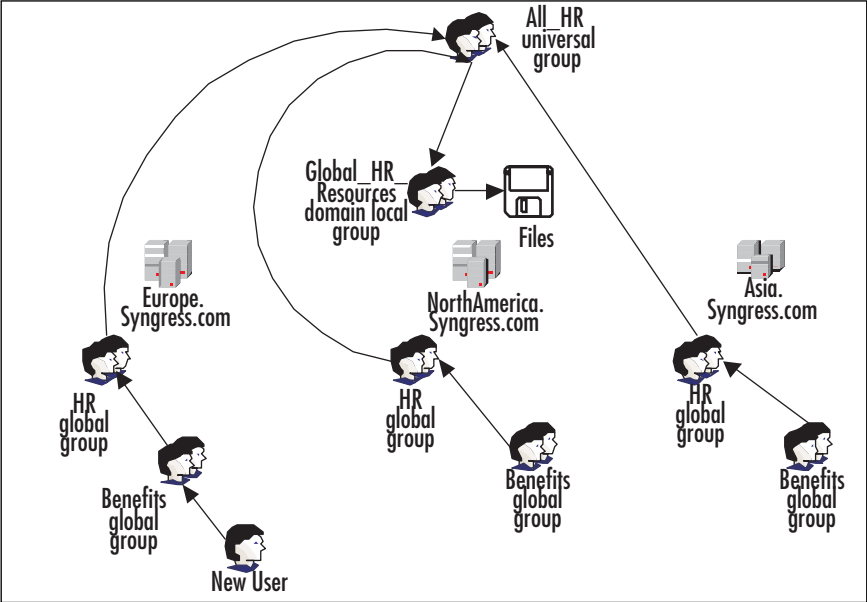


Moving all of the domains in the forest to the Windows 2000 native or Windows Server 2003 functional level greatly reduces the complexity. Just as we saw in the previous section, when a new benefits user joins the company, the only group his or her account needs to be made a member of is the Benefits global group in his or her regional domain. Again, this is because the Benefits global group is nested in the HR global group.

The real power in a multiple domain environment, however, comes in the ability to use universal security groups. You no longer have to add each HR global group into the Global_HR_Resources domain local group. Instead, you can add all of the HR global groups into a universal group called ALL_HR. You then add this group into the Global_HR_Resources DLG. These group memberships are shown in Figure 3.22.

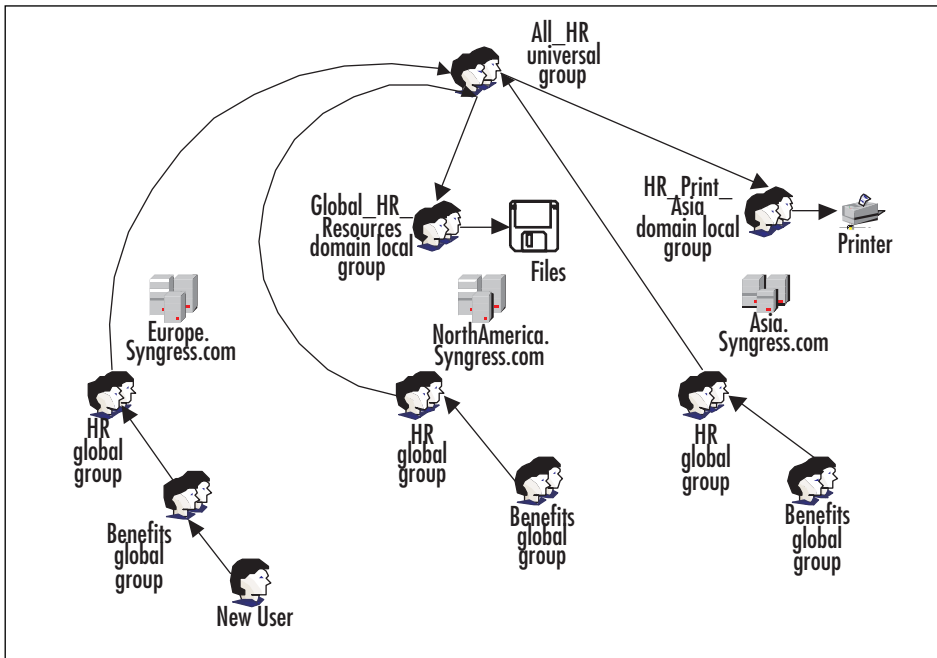
When universal groups enter the design, we are using the AGGUDLP model (sometimes abbreviated AGUDLP), where **U** represents Universal group. This model means: **A**ccounts should be placed into **G**lobal groups that can be placed into other **G**lobal groups and/or **U**niversal groups, and then into **D**omain **L**ocal groups, which are added to **A**CLs and granted **P**ermissions to resources.

Figure 3.22 AGGUDLP in a Multiple Domain Forest



While this might look like a similar amount of work when compared with Figure 3.21, the real power of this design becomes evident when you attempt to grant all HR users access to another resource, such as a printer in Asia. In this case, you simply need to create a new DLG and grant the print permission for the printer in the Asia domain to that group. In Figure 3.23, the group is called HR_Print_Asia. You then simply add the All_HR universal group to the HR_Print_Asia domain local group. Imagine what the diagram would look like if you couldn't use a universal group and how much more work would be involved. You would need to add each HR global group to the HR_Print_Asia domain local group. Now imagine that you have dozens of similar situations in your forest, and you'll no doubt appreciate the simplicity and reduced management requirements that universal groups bring with them.

Figure 3.23 Using AGGUDLP to Grant Access to an Additional Resource



Summary of Exam Objectives

This chapter addressed several key skills that are measured by Microsoft Exam 70-294. You should be well versed in the concepts presented in this chapter as well as the exercises designed to give you hands-on experience with some of the new functions and features in Windows Server 2003.

User passwords are often the weakest link in any network security implementation. To help combat this, Windows Server 2003 allows you to configure password and account lockout policies for all user accounts within a domain. You can configure passwords to expire after a specific number of days, mandate a minimum password length, and control how many unique passwords will be stored in Active Directory to control when a user can reuse an old password. Mandating password complexity forces your users to create passwords that are at least six characters long and contain three of the following types of characters: uppercase, lowercase, numeric, and non-alphanumeric. In addition, you can enforce an account lockout policy that will disable a user account after a certain number of incorrect logon attempts.

Planning a user authentication strategy involves a firm understanding of the various authentication protocols offered by Windows Server 2003. The default authentication protocol for LAN communication is Kerberos v5, although NTLM (and NTLMv2) is still supported to allow communication with pre-Windows 2000 computers on your network. Digest authentication, along with SSL/TLS, can provide secure access for users accessing resources via Web-based technology. You can implement one or more of these technologies when implementing an authentication plan for your network.

In a Windows Server 2003 environment, you can use smart cards as a stronger means of authentication for your users. Smart cards rely on Certificate Services to create enrollment certificates, as well as logon certificates to enable your users to authenticate and access network resources using their smart cards. This chapter also covered best practices for managing smart cards on your network, including preparing, issuing, maintaining, and revoking smart cards for the users.

There are two types of groups in a Windows Server 2003 domain: distribution and security. Distribution groups are used for messaging, such as e-mail. Security groups are the focus of the exam and are used to assign permissions and user rights. In addition, there are three group scopes in an AD domain: domain local, global, and universal. The capabilities of each scope vary depending on the functional level of the domain.

Domain local groups are used when assigning permissions or user rights. Microsoft wants you to use DLGs when modifying the ACL of an object such as a file, or assigning a user right. Other groups are added to a DLG to have their members receive the group's assigned permissions or rights. In a Windows 2000 mixed functional level domain, domain local groups can consist of users, computers, and global groups from the domain the DLG exists in, and any trusted domain. When the functional level of the domain is raised to Windows 2000 native or Windows Server 2003, a DLG can also contain other domain local groups from its local domain, as well as universal groups.

Microsoft specifies global groups as the primary container for user and computer objects. They call for grouping users according to role, function, responsibility, or department into global groups. In a Windows 2000 mixed functional level domain, a global group can contain users and computers from the same domain in which it exists. When the functional level of the domain is raised to Windows 2000 native or Windows Server 2003, a GG can also contain other global groups from its local domain.

Unlike global and domain local groups, universal groups (UGs) are stored in the Global Catalog (GC). Adding or removing objects from a universal group triggers forest-wide replication. To minimize this, Microsoft recommends that other groups, and not individual user and computer accounts, be the primary members of a universal group. Universal security groups do not exist in a Windows 2000 mixed functional level domain. When the functional level of the domain is raised to Windows 2000 native or Windows Server 2003, universal security groups can contain domain users, computer accounts, and global groups from any trusted domain, as well as other universal groups.

An administrator can change an existing group's scope. Universal groups can be converted to global or domain local groups, and global and domain local groups can be converted to universal groups. However, global groups cannot be converted directly to domain local groups (and vice versa). You cannot convert from one group type to another if the current membership of the group that is being converted is not compatible with the membership allowed for the target scope.

Microsoft has a number of acronyms that describe how groups should be used in different scenarios, including:

- **AGDLP** Accounts (user and computer objects) are placed into **G**lobal groups, which are placed into **D**omain **L**ocal groups, which are added to access control lists (ACLs) and granted **P**ermissions to a resource. This model is used in a single or multiple domain environment, when the Windows 2000 mixed domain functional level is in use.
- **AGGDLP** Accounts are placed into **G**lobal groups that can be placed into other **G**lobal groups and/or **D**omain **L**ocal groups, which are added to ACLs and granted **P**ermissions to resources. This model can only be used in domains that have a Windows 2000 native or Windows Server 2003 functional level.
- **AGGUDLP** (or **AGUDLP**) Accounts should be placed into **G**lobal groups that can be placed into other **G**lobal groups and/or **U**niversal groups, and then into **D**omain **L**ocal groups, which are added to ACLs and granted **P**ermissions to resources. This model can only be used in domains that have a Windows 2000 native or Windows Server 2003 functional level. In addition, it is primarily used in a multiple domain environment.

Exam Objectives Fast Track

Creating a Password Policy for Domain Users

- ☑ According to Microsoft, complex passwords consist of at least seven characters, including three of the following four character types: uppercase letters, lowercase letters, numeric digits, and non-alphanumeric characters such as & \$ * and !.
- ☑ Password policies and account lockout policies are set at the domain level in Group Policy.
- ☑ If a subset of your user base requires a different set of account policies and other security settings, you should create a separate domain to meet their requirements.
- ☑ Be sure that you understand the implications of an account lockout policy before you enable one in a production environment.

Creating User Authentication Strategies

- ☑ Within a domain, Kerberos v5 is the default communication method between two machines that are running Windows 2000 or later.
- ☑ Pre-Windows 2000 computers use NTLM (or NTLMv2) authentication in an Active Directory domain.
- ☑ To provide authentication for Web applications, you can implement either SSL/TLS or Microsoft Digest.

Planning a Smart Card Authentication Strategy

- ☑ Microsoft Windows Server 2003 relies on its public key infrastructure (PKI) and Certificate Services to facilitate smart card authentication.
- ☑ Smart card certificates are based on the following three certificate templates: Enrollment Agent, Smartcard Logon, and Smartcard User.
- ☑ Several Group Policy settings are specific to smart card implementations; most other account policy settings will also affect smart card users.

Planning a Security Group Strategy

- ☑ There are two types of groups in a Windows Server 2003 domain: distribution and security.
- ☑ Only security groups can be used to assign permissions.

- ☑ There are three group scopes in a Windows Server 2003 domain: domain local, global, and universal.
- ☑ Additional group nesting and universal security groups are only available at the Windows 2000 native and Windows Server 2003 domain functional levels.
- ☑ Existing groups can have their scopes changed in Windows 2000 native and Windows Server 2003 functional level domains.

Exam Objectives

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

Q: How can I configure a smart card user to be able to temporarily log on to the network if the user has forgotten his or her card?

A: In the **Properties** of the user's account within **Active Directory Users and Computers**, make the following changes on the **Account** tab:

1. Clear the check mark next to **Smart card is required for interactive logon**.
2. Place a check mark next to **User must change password at next logon**.

Finally, right-click the user object and select **Reset Password**. Inform the user of the new password, and that it will need to be changed at next logon.

Q: What are the advantages of implementing a “soft lockout” policy versus a “hard lockout”?

A: A hard lockout policy refers to an account that must be manually unlocked by an administrator. This setting provides the highest level of security but carries with it the risk that legitimate users will be unable to access network resources. In some circumstances, it can be used to effectively create a DoS attack against your own network. Hard lockouts place a greater burden on account administrators, because at least one must always be available for users to contact when they need their accounts unlocked. A soft lockout expires after a set amount of time and helps limit the effectiveness of password attacks against your network, while reducing the burden placed on administrators in a hard lockout environment.

- Q:** My organization is in the planning stages of a smart card rollout. What are the security considerations involved in setting up a smart card enrollment station?
- A:** Since a smart card enrollment station allows you to create certificates on behalf of any user within your Windows Server 2003 domain, you should secure these machines heavily in terms of both physical location and software patches. Imagine the damage that could be done if a malicious user were able to create a smart card logon certificate for a member of the Domain Admins group and use it to log on to your network at will.
- Q:** How can I convince my users that the company's new smart card rollout is something that is protecting them, rather than simply "yet another stupid rule to follow"?
- A:** One of the most critical components of any network security policy is securing "buy-in" from your users. A security mechanism that is not followed is not much more useful than one that doesn't exist. Try to explain the value of smart card authentication from the end-user's perspective. If you work in a sales organization, ask your sales force how they would feel if their client contacts, price quotes, and contracts fell into the hands of their main competitor. In a situation like this, providing a good answer to "What's in it for me?" can mean the difference between a successful security structure and a failed one.
- Q:** All of my workstations run Windows 95. I know that these don't support Kerberos for authentication. How can I configure the domain to use the NTLM protocol instead of the default of Kerberos protocol?
- A:** You do not need to perform any configuration to support NTLM authentication. Windows Server 2003 supports not only basic NTLM but also NTLM version 2, by default, for pre-Windows 2000 computers. In addition, NTLMv2 is more secure than NTLM, and will be automatically used if the domain controller is able to ascertain that the client supports it.
- Q:** I have a three-domain environment. All three of my domains have the same global groups. I've added the HR global group from two of the domains to an All_HR universal group. I've also added the All_HR universal group to domain local groups in these same two domains. Why can't I add the All_HR universal group to any domain local groups in my third domain?
- A:** All three domains must be at a functional level that supports universal security groups. It is possible to have a forest environment in which some domains are at the appropriate level and others are not. In this case, it sounds like two domains are at the Windows 2000 native or Windows Server 2003 functional level, but the third is at the Windows 2000 mixed functional level. Raise all domains to at least the Windows 2000 native level and try again.

Q: I'm in a single domain environment. My domain functional level is Windows Server 2003. I'm trying to convert a group from a global scope to a domain local scope. The group only contains users, but the option button is grayed out. What's wrong?

A: You cannot convert directly from a global group scope to a domain local group scope. You can only convert to and from a universal group scope. To accomplish this, you must first convert the global group to a universal group. Once this completes successfully, convert the universal group to the domain local group scope.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Creating a Password Policy for Domain Users

1. What is a potential drawback of creating a password policy on your network that requires user passwords to be 25 characters long?
 - A. Users will be more likely to write down a password that is so difficult to remember.
 - B. User passwords should be at least 30 characters long to guard against brute-force password attacks.
 - C. There are no drawbacks; this solution creates network passwords that will be impossible for an unauthorized user to penetrate.
 - D. Windows Server 2003 will not allow a password of more than eight characters.
2. You have recently started a new position as a network administrator for a Windows Server 2003 network. Shortly before the previous administrator left the company, the *syskey* utility was used on one of your domain controllers to create a password that needs to be entered when the machine is booted. You reboot the domain controller, only to discover that the password the previous administrator documented is incorrect. You are unable to contact your predecessor to obtain the correct one. How can you return this DC to service as quickly as possible?
 - A. Reformat the system drive on the server and reinstall Windows Server 2003.
 - B. Boot the server into Directory Services Restore Mode and restore the DC from a point before the previous administrator ran the *syskey* utility.

- C. Boot the server into Safe Mode and run *syskey* again to change the password.
 - D. Use *ntdsutil* to seize the PDC Emulator role and transfer it to another DC.
3. According to Microsoft, which of the following would be considered weak passwords for a user account named jronick? (Choose all that apply.)
- A. S#n\$IUsN7
 - B. soprano
 - C. ronickrj
 - D. Oo!dIx2
 - E. new
4. You have implemented a password policy that requires your users to change their passwords every 30 days and retains their last three passwords in memory. While sitting in the lunch room, you hear someone advise his coworker that all she needs to do to get around that rule is to change her password four times so that she can go back to using the password that she is used to. What is the best way to modify your domain password policy to avoid this potential security liability?
- A. Increase the maximum password age from 30 days to 60 days.
 - B. Enforce password complexity requirements for your domain users' passwords.
 - C. Increase the minimum password age to seven days.
 - D. Increase the minimum password length of your users' passwords.
5. You are a new network administrator for a Windows Server 2003 domain. In making user support calls, you have noticed that many users are relying on simplistic passwords such as their children's or pets' names. Passwords on the network are set to never expire, so some users have been using these weak passwords for years. You change the default Group Policy to require strong passwords. Several weeks later, you notice that the network users are still able to log on using their weak passwords. What is the most likely reason why the weak passwords are still in effect?
- A. You must force the users to change their passwords before the strong password settings will take effect.
 - B. The Group Policy settings have not replicated throughout the network yet.
 - C. Password policies need to be set at the organizational unit (OU) level, not the domain level.
 - D. The users reverted back to their passwords the next time they were prompted to change them.

Creating User Authentication Strategies

6. You have created an e-commerce Web application that allows your customers to purchase your company's products via the Internet. Management is concerned that customers will not feel comfortable providing their credit card information over the Internet. What is the most important step to secure this application so that your customers will feel confident that they are transmitting their information securely and to the correct Web site?
 - A. Use IP restrictions so that only your customers' specific IP addresses can connect to the e-commerce application.
 - B. Issue each of your customers a smart card that they can use to authenticate to your e-commerce Web site.
 - C. Place your company's Web server behind a firewall to prevent unauthorized access to customer information.
 - D. Install a Secure Sockets Layer (SSL) certificate on your Web server.
7. Your network environment consists of Windows 2000 Professional, Windows XP Professional, and Windows NT 4.0 Workstation computers. You have just upgraded all domain controllers to Windows Server 2003. The domain and forest functional levels are both set to Windows Server 2003. The company does not use any Web applications or services. Which of the following authentication protocols will be used on the network? (Choose all that apply.)
 - A. Digest
 - B. NTLM
 - C. Kerberos
 - D. SSL
8. You've decided to implement Web-based authentication. You have a wide range of domains, domain controllers, and domain functional levels in your enterprise Windows Server 2003 forest. Because you are a homogenous Windows environment, you decide to implement digest authentication. Which of the following requirements must you keep in mind when planning to implement digest authentication? (Choose all that apply.)
 - A. Digest authentication requires IE 5 or later on the clients.
 - B. There must be at least one Windows Server 2003 DC in the IIS server's domain.
 - C. User passwords must be stored with reverse encryption.
 - D. There must be at least one Windows 2000 or later DC in the IIS server's domain.

Planning a Smart Card Authentication Strategy

9. Your network configuration includes a Terminal Server designed to allow users at remote branches to access network applications. The Terminal Server often becomes overloaded with client requests, and you have received several complaints regarding response times during peak hours. You have recently issued smart cards for the users located at your corporate headquarters and would like to prevent those users from using their smart cards to access the Terminal Server. How can you accomplish this goal in the most efficient manner possible?
 - A. Enable auditing of logon/logoff events on your network to determine which smart card users are accessing the Terminal Server, and then speak to their supervisors individually.
 - B. Create a separate OU for your Terminal Server. Create a global group containing all smart card users, and restrict the logon hours of this group for the Terminal Server's OU.
 - C. Enable the "Do not allow smart card device redirection" setting within Group Policy.
 - D. Create a global group containing all smart card users, and deny this group the "Log on locally" right to the computers on your network.
10. You have attached a smart card reader to your Windows XP Professional workstation's serial port. The reader is not detected when you plug it in and is not recognized when you scan for new hardware within Device Manager. The smart card reader is listed on the Microsoft Web site as a supported device, and you have verified that all cables are connected properly. Why is your workstation refusing to recognize the smart card reader?
 - A. The manufacturer-specific installation routine is not compatible with Windows Server 2003.
 - B. The workstation needs to be rebooted before it will recognize the card reader.
 - C. Smart card readers are only supported on machines running Windows Server 2003.
 - D. You are not logged on as a member of the Domain Admins group.
11. You have recently deployed smart cards to your users for network authentication. You configured the Smartcard Logon certificates to expire every six months. One of your smart card users has left the company without returning her smart card. You have disabled this user's logon account, but management is concerned that she will still be able to use the smart card to access network resources. How can you be sure that the information stored on the former employee's smart card cannot be used to continue to access network resources?

- A. Monitor the security logs to ensure that the former employee is not attempting to access network resources.
- B. Use the smart card enrollment station to delete the user's Smartcard Logon certificate.
- C. Deny the Autoenroll permission to the user's account on the Smartcard Logon Certificate template.
- D. Add the user's certificate to the CRL on your company's CA, and publish the CRL.

Planning a Security Group Strategy

12. One of your coworkers is trying to grasp the concept of distribution and security group types. He asks you what the two primary benefits are for the security group type. What do you tell him? (Choose two.)
- A. You tell him that they can have permissions and user rights assigned to them.
 - B. You tell him that they can function for messaging just like a distribution group type.
 - C. You tell him that they allow for quick and efficient delegation of administrative responsibility in Active Directory.
 - D. You tell him that they can only be used for messaging and granting permissions to Active Directory, file system, Registry, and printer objects.
13. Your boss has been looking over marketing material from Microsoft. She asks you how you plan on using universal groups. You administer a single domain environment that is about to be upgraded to Windows Server 2003. What do you tell her?
- A. You tell her that because you will be using a Windows Server 2003 functional level domain, you will be using only universal groups.
 - B. You tell her that because you will be using a Windows 2000 native functional level domain, you will be using only universal groups.
 - C. You tell her that you will use universal groups to replace global groups, but will still be using domain local groups for resource access.
 - D. You tell her that you will not be using universal groups.
14. Last night you finished configuring a complex set of groups for your new Windows Server 2003 Active Directory environment. You spent this morning adding users to their appropriate groups. Now that the Active Directory environment is configured, you are trying to add the groups into ACLs in the file system. For some reason, they aren't showing up in the list of groups to select from. You can see all the default groups that the operating system and Active Directory installed. Why can't you see the groups you created?

- A. You don't have permission.
 - B. You didn't activate the groups in Active Directory.
 - C. You created distribution groups.
 - D. You created security groups.
15. Your company has a single domain environment that will be upgraded to Windows Server 2003. One of the company's existing Windows NT 4.0 BDCs must remain in place because a custom application requires it. This application will not be migrated until sometime next year. The company has many departments, each of which has sub-departments and teams. The company would like to take advantage of Windows Server 2003's new group nesting capabilities. Which of the following group models is appropriate for this company?
- A. AGDLP
 - B. AGGDLP
 - C. AGGUDLP
 - D. AGUDLP

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|-------------------|-----------------|
| 1. A | 9. C |
| 2. B | 10. B |
| 3. B, C, E | 11. D |
| 4. C | 12. A, C |
| 5. A | 13. D |
| 6. D | 14. C |
| 7. B, C | 15. A |
| 8. A, C, D | |

MCSA/MCSE 70-294

Working with Forests and Domains

Exam Objectives in this Chapter:

- 1.3.5 Set an Active Directory forest and domain functional level based on requirements.
- 1.3 Implement an Active Directory directory service forest and domain structure.
 - 2.1 Manage an Active Directory forest and domain structure.
 - 1.3.1 Create the forest root domain.
 - 1.3.2 Create a child domain.
 - 1.3.3 Create and configure Application Data Partitions.

- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

A Microsoft Active Directory network has both a physical and a logical structure. Forests and domains define the logical structure of the network, with domains organized into domain trees in which subdomains (called child domains) can be created under parent domains in a branching structure. Domains are logical units that hold users, groups, computers, and organizational units (OUs) (which in turn can contain users, groups, computers, and other OUs). Forests are collections of domain trees that have trust relationships with one another, but each domain tree has its own separate namespace.

In this chapter, you will learn all about the functions of forests and domains in the Windows Server 2003 Active Directory infrastructure, and we will walk you through the steps of creating a forest and domain structure for a network. You'll learn to install domain controllers (DCs), create the forest root domain and a child domain, find out how to name and rename domains, and how to set the functional level of a forest and domain.

The Domain Name System (DNS) is an integral part of a Windows Server 2003 network, as it is used for providing name resolution within the network. We will discuss the role of DNS in the Active Directory environment, and you'll learn about the relationship of the DNS and Active Directory namespaces, how DNS zones are integrated into Active Directory, and how to configure DNS servers for use with Active Directory.

Understanding Forest and Domain Functionality

A Windows Server 2003 *domain* is a group of networked computers that share a common Active Directory database, and a common namespace. You can think of a domain as a limited boundary of network security and administrative control. A *namespace* is a hierarchical collection of service and object names, typically stored within DNS and Active Directory.

There are some similarities between the Active Directory namespace and the DNS namespace, both of which are required by Windows Server 2003. For example, the name of an Active Directory tree is derived from the DNS name of the tree root, which means that both namespaces share the same root. When you rename the root domain, you must automatically rename all child domains in the tree to match; hence, all levels of both namespace hierarchies. The Active Directory and DNS namespaces, by Microsoft definition, must have the same name. Exceptions do exist, however, such as during a domain rename procedure.

How Can I Tell the Difference between the Active Directory Namespace and the DNS Namespace?

Among the differences in the two namespaces is the ability of DNS to split a domain name into two separate zones. In split-DNS, one zone typically provides name resolution for resources outside the firewall, while the other zone provides name resolution for the inside. Inside users can locate and use external resources. An Active Directory domain cannot be split in the same way and continue to fully interoperate.

Another difference is where the data is stored. Even given identical names, and even with Active Directory integrated DNS, the two namespaces occupy different partitions within the directory. This gives them different logical addresses, although replication of the two is accomplished in the same way. With non-Active Directory-integrated DNS, the namespaces do not reside in the same directory and do not need to reside on the same servers. Non-integrated DNS must also provide its own replication topology. In either case, the data is always discretely separated. DNS records and Active Directory objects work together, but never truly intermingle.

One of the most distinct differences is the real-time nature of dynamic DNS. When a server is shut down, dynamic DNS removes the resource records associated with that server from its database. Unless you created static records, as you might for an e-mail or web server, DNS retains no knowledge of the machine. Active Directory, by contrast, requires the stability of constant knowledge for all hosts. If a server were to be removed and re-added to Active Directory, the host would receive a new Security Identifier (SID) and be treated as a new and unique system. In Active Directory, hosts within the same domain are often subdivided into sites and OUs, while DNS hosts are only differentiated by record types.

These distinctions help clarify the forest and domain structure, the namespaces they define, and the interoperability between them.

Active Directory is composed of a number of components, each associated with a different concept, or layer of functionality. You should understand each of these layers before making any changes to the network. The Active Directory itself is a distributed database, which means it can be spread across multiple computers within the forest. Among the major logical components are:

- Forests
- Trees
- Domains
- The domain namespace

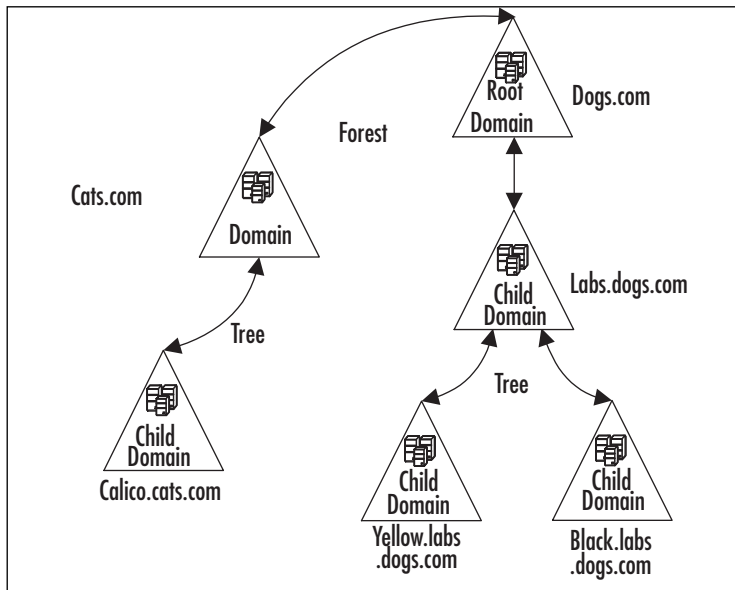
Aspects of the physical structure include the following:

- Sites
- Servers
- Roles
- Links

Administrative boundaries, network and directory performance, security, resource management, and basic functionality are all dependent on the proper interaction of these elements.

Figure 4.1 shows the logical view of a Windows Server 2003 Active Directory. Note that the differentiation between forests and trees is most obvious in the namespace. By its nature, a *tree* is one or more domains with a contiguous namespace. Each tree consists of one or more domains, while each *forest* consists of one or more trees. Because a forest can be composed of discrete multiple trees, a forest's namespace can be discontinuous. By *discontinuous*, we mean that the namespaces anchor to different forest-root DNS domains, such as cats.com and dogs.com. Both are top-level domains and are considered two trees in a forest when combined into a single directory as shown in Figure 4.1.

Figure 4.1 The Forest Structure



The Role of the Forest

An Active Directory always begins with a *forest root domain*, which is automatically the first domain you install. This root domain becomes the foundation for additional directory components. As the cornerstone of your enterprise-computing environment, you should protect

it well. Fault tolerance and good backups are not optional—they are essential. If an administrative error or hardware failure results in the unrecoverable loss of this root structure, the entire forest becomes inoperable. Certain forest objects and services are only present at the root (for example, the Enterprise Administrators and Schema Administrators groups, and the Schema Master and Domain Naming Master roles). These cannot be easily recreated, depending on the type of failure.

New Forestwide Features

Many of the new features offered by Windows Server 2003 are only available in a forest where you have raised the forest functional level to Windows Server 2003. For more information on functional levels and a breakdown of when these new features become available, see the section *Forest and Domain Functional Levels* later in the chapter

Defunct Schema Objects

In Windows 2000 Active Directory, you could deactivate a schema class or attribute. Now, once your forest has been raised to the Windows Server 2003 functional level, you cannot only deactivate them, you can even rename and redefine them. This feature protects against the possibility of one application irreversibly claiming another application's schema. It allows for the redefinition of classes and attributes without changing their unique identities. These items are called *reused*. If the class or attribute is left deactivated, it is called *defunct*.

Where this becomes important is where, for example, you make an error in the definition of an attribute. In Windows 2000, the best you can do is deactivate the attribute with the incorrect syntax and create a new one with a different name. If you have an application that requires a certain attribute name, there's little you can do but operate with the incorrect definition, get by without it altogether, or find a different application. Restoring the schema from a state backup is possible, but risky. Now, with the new functionality of Windows Server 2003, you can deactivate the incorrect attribute and safely create a new one that uses the same object identifier (OID) and Lightweight Directory Access Protocol (LDAP) display name as the old one, but with the correct syntax.

Another case is when an object identifier collision occurs. This is where a needed OID conflicts with an existing one, a situation usually created by mistyping a number. By deactivating the first OID, the second can be created. There are several situations in which classes and attributes cannot be deactivated, and it is an operation that should always be performed with great care and planning.

Domain Rename

This is a complex and sweeping modification to the namespace of a domain. DNS names, and NetBIOS names of any child, parent, or forest-root domain can now be changed. As far as Windows Server 2003 Active Directory is concerned, the identity of a domain rests in its domain Globally Unique Identifier (GUID), and its domain SID. Creating new DNS or NetBIOS names will leave those attributes unchanged. The domain rename function is not

able to promote a domain to the forest root role. Even if you rename the forest root domain, its role will remain unchanged.

The renaming process will temporarily interrupt the functionality of the domain and its interaction with the forest, until the DCs are rebooted. Client workstations will not function properly until they are each rebooted *twice*. Due to the complexity of the operation, the risks of such a sweeping change, and the unavoidable domain and workstation service interruptions, domain renaming should not be considered a routine operation.

Forest Restructuring

Existing domains can now be moved to other locations within the namespace. During this restructuring, you will manually break and reestablish the appropriate trust relationships among the domains. A requirement for namespace changes, or a need to decrease administrative overhead, typically drives forest restructuring. This reduction in overhead is accomplished by reducing replication traffic, reducing the amount of user and group administration required, and simplifying the administration of Group Policy. The smallest possible number of domains will provide the most efficient design. Minimizing the number of domains reduces administrative costs and increases the efficiency of your organization. Reasons to restructure include:

- Decommissioning a domain that is no longer needed
- Changing the internal namespace
- Upgrading your network infrastructure to increase your bandwidth and replication capacity, which enables you to combine domains

Before you begin restructuring Windows Server 2003 domains within your forest, make sure that the forest is operating at the Windows Server 2003 functional level.

Universal Group Caching

Before Windows Server 2003, some sites had to make a decision to deploy a Global Catalog (GC) at each remote site regardless of the number of users at that location, because each DC contacts a GC server during a Windows 2000 native mode logon. The problem was that a GC generated a lot of replication traffic and required a lot of disk space, memory, and WAN bandwidth. The solution in Windows Server 2003 is Universal Group caching.

Universal Group caching is a new feature of the Windows Server 2003 DC, which caches a user's complete Universal Group membership. The cache is populated at first logon, and subsequent logons use the cache, which is refreshed periodically.

Some of the benefits of Universal Group caching include faster logon times. Authenticating DCs no longer have to consult a GC to get Universal Group membership information. In addition, you can save the cost of upgrading a server to handle the extra load for hosting the GC. Finally, network bandwidth is minimized because a DC no longer has to handle replication for all of the objects located in the forest.

Application Partitions

Another DC enhancement allows for the creation of application-specific Active Directory partitions, also known as *naming contexts*. Active Directory stores the information in a hierarchy that can be populated with any type of object except for security principles such as users, groups, and computers. This dynamic body of data can be configured with a replication strategy involving DCs across the entire forest, not just a single domain. With application partitions, you can define as many or as few replicas as you want. Site topologies and replication schedules are observed, and the application objects are not replicated to the GC. Conveniently, application partitions can leverage DNS for location and naming. The Windows Server 2003 Web Edition cannot host application partitions because they do not support the DC role.

New & Noteworthy...

Active Directory Application Partitions Can Exist on a Non-DC

Another new type of application partition is the Active Directory in Application Mode (ADAM) stand-alone product that allows Windows Server 2003 web edition and other member servers and workstations to participate in a form of application partitions without being DCs. It is maintained and replicated independent of the central Active Directory, although it interfaces with directory-enabled Kerberos and NTLM for authentication services. One advantage with this configuration is that schema changes made to support Web-based applications do not have to clutter up the core operating system's (OS's) schema. It gives you local control and naming flexibility in addition to the autonomous schema, and can be run on Windows XP or Windows Server 2003. ADAM is sometimes referred to as *Active Directory "Light."*

ADAM runs as a non-OS service. This means that multiple instances can run concurrently on a single server, with each instance being independently configurable. It is an extended capability that allows you to deploy Active Directory as a lightweight directory service for the rapid and flexible implementation of directory-enabled applications.

ADAM can be particularly helpful in the following areas:

- **Application-specific directories**, where you can store "private" directory data relevant only to the application.
- **Application developer activities**, where ADAM uses the same programming model and administration as Active Directory. This enables the developer to work with a local instance on the developer workstation and then later move the application to Active Directory.
- **Extranet Access Management (EAM) solutions**, such as hosting user objects that are not Active Directory security principals. This allows you to use LDAP to authenticate non-Windows or external users.

Continued

- **Migration scenarios**, where an organization has an established X.500 directory that must be maintained to serve legacy applications.

Install from Backups

The *Install from backups* feature provides the capability to install a DC using backup media rather than populating the Active Directory through a lengthy replication period. This is especially useful for domains that cross-site boundaries using limited WAN connectivity. To do this, back up your directory store using **Windows Backup**, restore the files at the remote site's candidate DC, and run **dcpromo** using the **source replication from files** option. This also works for GC servers.

Active Directory Quotas

The new Active Directory quotas (not to be confused with disk quotas) are defined as the number of objects that can be owned by a given user in a given directory partition. Fortunately, Domain Admins and Enterprise Administrators are exempt from the quota, and they do not apply at all to the schema partition. Replicated operations do not count toward the quota; only the original operations do. Quota administration is performed through a set of command-line tools, including *dsadd*, *dsmode*, *dsget*, and *dsquery*. No graphical interface exists for quota administration.

Linked Value Replication

Linked value replication provides an answer to Windows 2000's limit of 5000 direct group members. Instead of treating a large group as a single replication unit, linked value replication allows a single member to be added or removed from the group during replication, thereby reducing network traffic. Without it, for example, any changes to a 10,000-member distribution group will trigger a complete replication. With a group that large, this would be likely to occur many times in a typical day.

Improved Knowledge Consistency Checker

The Windows 2000 Knowledge Consistency Checker (KCC) would not operate properly within a forest containing more than 200 sites due to the complexity of the inter-site replication topology generator algorithms. The service had to be turned off in that case, and the replication topology had to be managed manually. The Windows Server 2003 KCC can automatically manage replication among up to 5000 sites due to new, more efficient algorithms. In addition, it uses greatly improved topology generation event logging to assist in troubleshooting.

Reduced NTDS.DIT Size

The Windows Server 2003 directory takes advantage of a new feature called *Single Instance Store* (SIS). This limits the duplication of redundant information. The new directory store is about 60 percent smaller than the one in Windows 2000.

Forest Trusts

In Windows NT 4.0, there were few options for the interoperability of business units; for example, either Calico.cats.com trusted Labs.dogs.com or they didn't. There were no other real options. In addition, if trust existed at all, it tended to be complete. When Windows 2000 introduced the Active Directory, many more options became available so that partnerships and integrated project teams could form on the network just as they did in real life. The problem with that approach was that there always had to be a dominant partner at the root—the playing field could never be completely even.

The idyllic utopia of a single forest cannot handle certain situations. The root owner employs Administrators, Domain Admins, and Enterprise Admins, any of which can gain access to any resource in the forest with nothing more than a little persistence. Domains make good administrative boundaries, and domains and sites make good replication boundaries, but only a forest can provide a viable security boundary.


Understanding the politics of business, Microsoft stepped in with a solution called *multiple-forest trusts* in Windows Server 2003, which, when used, result in a configuration called *federated forests*. Without the forest trust, Kerberos authentication between forests would not work. Remember that having two forests means two Active Directory databases and two completely distinct sets of directory objects, such as user accounts. Accessing resources across the federated forest boundary requires a more complex trust path than the one between domains within a single forest. See Figure 4.2 for an example of a multiple-forest trust path.



NOTE

Note that “federated forest” is not a term you'll find in the Windows Server 2003 Help files. However, this terminology has been used in TechNet articles on Windows Server 2003. For more information on the concept and implementation, see *Planning and Implementing Federated Forests in Windows Server 2003*: www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windowsserver2003/maintain/security/fedffin2.asp.

Figure 4.2 The Forest Trust Path



Head of the Class...

How Can I Share Resources between Two Active Directories?

Here's how sharing resources between two Active Directories works. Say that Tabby, a user in the Windows 2000 Calico.cats.com domain, tries to access the public folder on a file server called Onyx in the Windows Server 2003 Black.labs.dogs.com domain as shown in Figure 4.2.

First, Tabby logs on to her workstation using Kerberos authentication and tries to access a public folder on Onyx. Her workstation naïvely contacts one of the Calico.cats.com DCs, which hosts the Kerberos KDC, requesting a service ticket for the server principle name (SPN) of Onyx.black.labs.dogs.com. Naturally, the DC's database doesn't contain that information, so it queries the Cats.com GC to see if any of the other domains in its forest contain such a machine. As it turns out, the GC isn't so global, and Tabby gets an error.

This is because the Windows 2000 GC is limited to its own forest. Tabby wisely purrs and convinces her manager to upgrade their forest to Windows Server 2003. Being very catty, the Enterprise Administrators in Cats.com quickly take care of the

Continued

prerequisites for the establishment of a forest trust with Dogs.com as soon as the upgrade is complete.

This time, instead of generating an error, Tabby's newly upgraded GC checks its database for forest trusts. When it finds one, it looks at the forest trust *trusted domain object* to see if its listed name suffixes correspond with the target SPN. Sure enough, a match is found, and it generates a *routing hint* back to the Calico.cat.com DC, which in turn hints to Tabby's workstation that it needs to go climb a different tree.

Undaunted, the workstation asks a forest-root DC at Cats.com for a referral to one of the DCs at the forest root of Dogs.com, based on the routing hint just received. This generates more electronic red tape. Now the Calico workstation has to make a request to the Dogs.com KDC for a service ticket to Onyx. Not the brightest bulb in the pack, the KDC has to ask its own GC server in the Dogs.com domain to see if it knows this file server. Being as global as needed this time, a match is found and the SPN goes back to the Dogs.com forest-root KDC, which sends it off to Tabby's workstation back in the Calico child domain.

Success! Well, almost. Starting all over again, this time with a resolved SPN, the workstation negotiates with the KDC in Calico.cats.com for Tabby to access Onyx, and receives the appropriate server service ticket. Finally, sending the service ticket directly to Onyx through a *trust path* of one *forest* trust, two *tree-root* domain trusts, and one *parent and child* domain trust, the file server examines Tabby's credentials and sends her an access token.

Windows Explorer opens and displays the filenames in the \\Onyx\public folder. Tabby, unaware of the complex chain of events set off by her request, accesses the files.

Routing Hints for Forest Trusts

Routing hints are a new feature of GCs. The problem with creating trusts between forests is that all traditional authentication channels stop at the forest boundary. DCs and traditional GCs are sometimes not enough. When these fail to produce an SPN describing the location of the service being requested, routing hints from the Windows Server 2003 GC help guide the workstation toward the correct forest within the *Federated Forest* boundary. The GC server does this by checking the forest trust's *trusted domain object* (TDO) for trusted name suffixes that match the one found in the destination SPN. The routing hint always goes back to the originating device so that it can resume its search for the SPN location in the other forest. This new functionality has some limitations. If the TDO contains outdated or incorrect information, the *hint* might be incorrect since the GC does not actually check for the existence of the other forests.

Cross-Forest Authentication

Although some types of data access are supported, Windows Server 2003 does not support NetBIOS name resolution or Kerberos delegation across forests. NTLM authentication for down-level clients continues to be fully supported, however. A Universal Group in one forest might contain global groups from one or more additional forests across any available forest trusts.

Federated Forest, or cross-forest, authentication takes two forms. In the default *forest-wide authentication*, an “allow-all deny-some” approach is used. In other words, external users have the same level of access to local resources as the local users do. The other form of access control takes the security conscious approach of “deny-all allow-some.” This optional method is called *selective authentication*, and requires more administrative overhead by granting explicit control over the outside use of local resources. You must set a control access right called *allowed to authenticate* on an object for the users and groups that need access from another forest. If selective authentication is enabled, an *Other Organization* SID is associated with the user. This SID is then used to differentiate the external user from local users and determines if an attempt can be made to authenticate with the destination service.

For reliable authentication using Kerberos, system time must be accurate across every workstation and server. Servers are best synchronized with the same time source, while workstations are synchronizing time with the servers. In an upgraded Active Directory domain, this is usually not a problem. The Windows Server 2003 W32Time service provides time synchronization for all Windows XP and Windows 2003 OSs. Kerberos version 5 is particularly time sensitive and might falsely interpret logon requests as intrusion attempts if the time is off. In that case, user access will be denied. Earlier versions of Windows might need some assistance with the *net time* command in a logon script to stay current. In a federated forest, individual enterprises can choose to attune with different time sources. If these sources diverge, although each forest is chronologically homogenous, they might not agree with each other, resulting in a failure of all cross-forest authentications.

The Role of the Domain

The *domain* is the starting point of Active Directory. It is the most basic component that can functionally host the directory. Simply put, Active Directory uses the domain as a container of computers, users, groups, and other object containers. Objects within the domain share a common directory database partition, replication boundaries and characteristics, security policies, and security relationships with other domains.

Typically, administrative rights granted in one domain are only valid within that domain. This also applies to Group Policy Objects (GPOs), but not necessarily to trust relationships, which you will learn more about later in the book. Security policies such as the password policy, account lockout policy, and the Kerberos ticket policy are defined on a per-domain basis. The domain is also the primary boundary defining your DNS and

NetBIOS namespaces. The DNS infrastructure is a requirement for an Active Directory domain, and should be defined before you create the domain.

There are several good reasons for a multiple domain model, although the best overall practice consists of an empty root domain with a single user domain. Do not install additional domains unless you have a specific reason for them. Some of the more common reasons include:

- Groups of users with different security policy requirements, such as strong authentication and strict access controls.
- Groups of users requiring additional autonomy, or administrative separation for security reasons.
- A requirement for decentralized administration due to political, budgetary, time zone, or policy pressures.
- A requirement for unique namespaces.
- Controlling excessive directory replication traffic by breaking the domain into smaller, more manageable pieces. This often occurs in an extremely large domain, or due to a combination of geographical separation and unreliable WAN links.
- Maintaining a pre-existing NT domain structure.

The primary Active Directory partitions, also called *naming contexts*, are replicated among all DCs within a domain. These three partitions are the schema partition, the configuration partition, and the domain partition.

- The **schema partition** contains the *classSchema* and the *attributeSchema* objects that make up the directory schema. These classes and attributes define all possible types of objects and object properties within the forest. Every DC in the entire forest has a replica of the same schema partition.
- The **configuration partition**, replicated identically on all DCs throughout the forest, contains Active Directory's replication topology and other configuration data.
- The **domain partition** contains the local domain objects, such as computers, users, and groups, which all share the same security policies and security relationships with other domains. If multiple DCs exist within a domain, they contain a replica of the same domain partition. If multiple domains exist within a forest, each domain contains a unique domain partition.

Since each domain contains unique principles and resources, there must be some way for other domains to locate them. Active Directory contains objects that adhere to a naming convention called the DN, or *distinguished name*. The DN contains enough detail to locate a replica of the partition that holds the object in question. Unfortunately, most users and applications do not know the DN, or what partition might contain it. To fulfill that

role, Active Directory uses the GC, which can locate DNs based on one or more specific attributes of the needed object.

The GC contains a portion of every naming context in the directory, including the schema and configuration partitions. In order to be able to find everything, the GC must contain a replica of *every object in the Active Directory*. Fortunately, it only maintains a small number of attributes for each object. These attributes are those most commonly used to search for objects, such as a user's first, last, and logon names. The GC extends an umbrella of awareness throughout the discontinuous namespace of the enterprise.

Although the GC can be modified and optimized, it typically requires infrequent attention. The Active Directory replication system automatically builds and maintains the GC, generates its replication topology, and determines which attributes to include in its index.



EXAM WARNING

Remember this distinction between the GC and the Schema Master: The GC contains a limited set of attributes of all objects in the Active Directory. The *Schema Master* contains formal definitions of every object class that can exist in the forest and every object attribute that can exist within an object.

In other words, the GC contains *every object*, while the schema contains *every definition* of every type of object.

New Domainwide Features

Active Directory technology debuted with Windows 2000. Now, with Windows Server 2003, it has been refined and enhanced. Active Directory is now easier to deploy, more efficient at replication, has improved administration, and poses a better end-user experience. Some features are enabled right away, while others require a complete migration of DCs to the new release before they become available. There are countless new features, the most significant of which we discuss next.

Domain Controller Rename

Not to be confused with domain renaming, *domain controller rename* is the ability to rename a DC without following the Windows 2000 procedure of demoting, renaming, and promoting again. In a large domain, this saves considerable time, especially over a slow WAN link, since the process of re-promoting the DC requires a replication of the Active Directory.

Renaming a DC in Windows Server 2003 is much easier than it was in 2000, but that does not mean it has become a simple procedure. If you have multiple DCs, before you rename one of them you must make sure of a few things first. If any Operational Master roles reside on the DC, you need to transfer them to another DC. If the DC is a GC server, you have to move that role as well. Remember that the first DC you install in the forest is

the root DC. This DC is responsible for the GC and for *all* Flexible Single Master Operations (FSMO) roles unless you have spread them out manually. You need to transfer all of these functions to another DC before you rename the server.

Universal Groups and Group Conversions

Universal Groups are able to contain members from any domain in any forest, and they replicate to the GC. They are particularly useful for administrative groups. One of the best uses for groups with universal scope is to consolidate groups above the domain level. To do this, add domain user accounts to groups with global scope and nest these Global Groups within Universal Groups. Using this strategy, changes to the Global Groups do not directly affect the membership of groups with universal scope. Taking it one step further, a Universal Group in one forest can contain Global Groups from one or more *additional* forests across any available forest trusts.

Here is an example. Refer to Figure 4.2. You have two domains in different forests with NetBIOS names of CATS and DOGS. Each domain contains a Global Group called Birdwatchers. To take advantage of this new capability, you add both of the Global Groups, CATS\Birdwatchers and DOGS\Birdwatchers, to a Universal Group you create called ALLBirdwatchers. The second step is to create an identical Universal Group in the other forest as well. The ALLBirdwatchers group can now be used to authenticate users anywhere in both enterprises. Any changes in the membership of the individual Birdwatchers groups will not cause replication of the ALLBirdwatchers group.

You should strive to manage your Universal Groups in such a way as to minimize the frequency of changes, since every change causes the entire membership of the group to be replicated to every GC in the forest. A newly created group, by default is configured as a Security Group with global scope regardless of the current domain functional level. Refer to Table 4.1 for a summary of Universal Group capabilities that are available at the various domain functional levels.

Groups can also be changed from one scope to another, within certain limitations. Changing a group scope is *not* allowed in domains with a functional level of Windows 2000 mixed or Windows Server 2003 interim. The following scope conversions *are* allowed in domains with a functional level of Windows 2000 native or Windows Server 2003:

- Global to Universal, *if* the group you want to change is not a member of another Global Group.
- Domain Local to Universal, *if* the group you want to change does not have another Domain Local Group as a member.
- Universal to Global, *if* the group you want to change does not have another Universal Group as a member.
- Universal to Domain Local, with no restrictions.

Table 4.1 Summary of Universal Group Capabilities by Domain Functional Level

Functional Level	Universal Group Members	Universal Group Nesting
Windows 2000 mixed	None	None
Windows 2000 native	User and computer accounts, Global Groups, and Universal Groups from any domain	Universal Groups can be added to other groups and assigned permissions in any domain
Windows Server 2003 interim	None	None
Windows Server 2003	User and computer accounts, Global Groups, and Universal Groups from any domain	Universal Groups can be added to other groups and assigned permissions in any domain

Security Group Nesting

Security Groups are used to grant access to resources. Using nesting, you can add a group to a group. This reduces replication traffic by nesting groups to consolidate member accounts. A *Security Group* can also be used as an e-mail distribution list, but a *Distribution Group* cannot be used in a discretionary access control list (DACL), which means it cannot be used to grant access to resources. Sending e-mail to a *Security Group* sends the message to all members of the group.

In the Windows 2000 mixed domain functional level, *Security Groups* are restricted to the following members:

- Global Groups can only have user accounts as members.
- Domain Local Groups can have other Global Groups and user accounts as members.
- Universal Groups cannot be created.



EXAM WARNING

It is very important to know the different restrictions on group memberships at different domain functional levels.

Distribution Group Nesting

Distribution Groups are collections of users, computers, contacts, and other groups. They are typically used only for e-mail applications. *Security Groups*, on the other hand, are used to grant access to resources *and* as e-mail distribution lists. Using nesting, you can add a

group to a group. *Group nesting* consolidates member accounts and reduces replication traffic. Windows NT did not support Distribution Groups within the OS, but they are supported in all versions of Active Directory. Distribution Groups cannot be listed in DACLs in any version of Windows, which means they cannot be used to define permissions on resources and objects, although they *can* be used in DACLs at the application layer. Microsoft Exchange is a common example. If you do not need a group for security purposes, create a Distribution Group instead.

Number of Domain Objects Supported

In Windows 2000, group membership was stored in Active Directory as a single multivalued attribute. When the membership list changed, the entire group had to be replicated to all DCs. So that the store could be updated in a single transaction during the replication process, group memberships were limited to 5000 members. In Windows Server 2003, *Linked Value Replication* removes this limitation and minimizes network traffic by setting the granularity of group replication to a single principle value, such as a user or group.

Distribution Groups

Distribution Groups, unlike Security Groups, are not primarily used for access control, although they can be used in an ACL at the application layer. Distribution groups are designed to be used with e-mail applications only. You can convert a Distribution Group to a Security Group (or vice versa), if the functional level is Windows 2000 native or higher. You have to be a domain or enterprise admin, or a member of the Account Operators Group (or have the appropriate authority delegated) to convert a group. Changing the group type is as simple as right-clicking the group in **Active Directory Users and Computers**, clicking **Properties**, and clicking the desired group type on the **General** tab.

Domain Trees

A domain tree can be thought of as a DNS namespace composed of one or more domains. If you plan to create a forest with discontinuous namespaces, you must create more than one tree. Referring back to Figure 4.1, you see two trees in that forest, Cats.com and Dogs.com. Each has a *contiguous namespace* because each domain in the hierarchy is directly related to the domains above and below it in each tree. The forest has a *discontiguous namespace* because it contains two unrelated top-level domains.

Forest and Domain Functional Levels

Functional levels are a mechanism that Microsoft uses to remove obsolete backward compatibility within the Active Directory. It is a feature that helps improve performance and security. In Windows 2000, each domain had two functional levels (which were called “modes”), native mode and mixed mode, while the forest only had one functional level. In Windows Server 2003, there are two more levels to consider in both domains and forests. To enable all Windows Server 2003 forest and domainwide features, all DCs must be run-

ning Windows Server 2003 and the functional levels must be set to *Windows Server 2003*. Table 4.2 summarizes the levels, DCs supported in each level, and each level's primary purpose.

Table 4.2 Domain and Forest Functional Levels

Type	Functional Level	Supported DCs	Purpose
Domain Default	Windows 2000 mixed	NT, 2000, 2003	Supports mixed environments during upgrade; low security, high compatibility
Domain	Windows 2000 native	2000, 2003	Supports upgrade from 2000 to 2003
Domain	Windows Server 2003 interim	NT, 2003	Supports upgrade from NT to 2003; low security, no new features
Domain	Windows Server 2003	2003	Ideal level, best security, least compatibility, all new Active Directory features are enabled
Forest Default	Windows 2000	NT, 2000, 2003	Supports mixed environments during upgrade; low security, high compatibility
Forest	Windows Server 2003 interim	NT, 2003	Supports upgrade from NT to 2003; low security, some new features
Forest	Windows Server 2003	2003	Ideal level, best security, least compatibility, all new Active Directory features are enabled

Domain Functionality

When considering raising the domain functionality level, remember that the new features will directly affect only the domain being raised. The four available functional levels are:

- Windows 2000 mixed
- Windows 2000 native
- Windows Server 2003 interim

- Windows Server 2003

Once the domain functional level has been raised, no prior version DCs can be added to the domain. In the case of the Windows Server 2003 domain functional level, no Windows 2000 servers can be promoted to DC status after the functionality has been raised. Table 4.2 summarizes the levels, DCs supported in each level, and the level's primary purpose. See Table 4.3 for a summary of the capabilities of the current Windows 2000 and new Windows Server 2003 domain functional levels.

Table 4.3 Domain Functional Level Features

Domain Feature	Windows 2000 Mixed	Windows 2000 Native	Windows Server 2003 Interim	Windows Server 2003 Native
Local and Global Groups	Enabled	Enabled	Enabled	Enabled
Distribution Groups	Enabled	Enabled	Enabled	Enabled
GC support	Enabled	Enabled	Enabled	Enabled
Number of domain objects supported	40,000	1,000,000	40,000	1,000,000
Kerberos KDC key version numbers	Disabled	Disabled	Disabled	Enabled
Security Group nesting	Disabled	Enabled	Disabled	Enabled
Distribution Group nesting	Enabled	Enabled	Enabled	Enabled
Universal Groups	Disabled	Enabled	Disabled	Enabled
SIDHistory	Disabled	Enabled	Disabled	Enabled
Converting groups between Security Groups and Distribution Groups	Disabled	Enabled	Disabled	Enabled
DC rename	Disabled	Disabled	Disabled	Enabled
Logon timestamp attribute updated and replicated	Disabled	Disabled	Disabled	Enabled
User password support on the <i>InetOrgPerson</i> objectClass	Disabled	Disabled	Disabled	Enabled

Continued

Table 4.3 Domain Functional Level Features

Domain Feature	Windows 2000 Mixed	Windows 2000 Native	Windows Server 2003 Interim	Windows Server 2003 Native
Constrained delegation	Disabled	Disabled	Disabled	Enabled
Users and Computers container redirection	Disabled	Disabled	Disabled	Enabled

Windows 2000 Mixed Domain Functional Level

The Windows 2000 mixed domain functional level is primarily designed to support mixed environments during the course of an upgrade. Typically, this applies to a transition from Windows NT to Windows 2000, although it is also the default mode for a newly created Windows Server 2003 domain. It is characterized by lowered security features and defaults, and the highest compatibility level possible for Active Directory.

- All Windows DCs are supported.
- Active Directory domain features *not* supported in this mode:
 - Group nesting
 - Universal Groups
 - SIDHistory
 - Converting groups between Security Groups and Distribution Groups
 - Domain controller rename
 - Logon timestamp attribute updated and replicated
 - User password support on the *InetOrgPerson objectClass*
 - Constrained delegation
 - Users and Computers container redirection
- Can be raised to Windows 2000 native mode or directly to the Windows Server 2003 domain level.
- Can never be lowered since no lower domain functional level exists.

In the *Windows 2000 mixed* functional level, which is the default level, Windows 2000 and greater DCs can exist, as well as Windows NT backup domain controllers (BDCs). Newly created Windows Server 2003 domains always start at this level. Windows NT primary domain controllers (PDCs) do not exist in any version of Active Directory.

Windows 2000 Native Domain Functional Level

The Windows 2000 native domain functional level is primarily intended to support an upgrade from Windows 2000 to Server 2003. Typically, this applies to existing Active Directory implementations since mixed and interim modes support the upgrade from Windows NT. It is characterized by better security features and defaults, and an average compatibility level.

- Microsoft Windows NT 4.0 DCs are not supported.
- Active Directory domain features *not* supported in this mode:
 - Domain controller rename
 - Logon timestamp attribute updated and replicated
 - User password support on the *InetOrgPerson objectClass*
 - Constrained delegation
 - Users and Computers container redirection
- Can be raised to the Windows Server 2003 domain level.
- Can never be lowered back to the Windows 2000 mixed mode.

In *Windows 2000 native* functional level, DCs have all been upgraded to Windows 2000 or Windows Server 2003. Native mode enables Universal Security Groups, nested groups, group conversion between distribution and security types, and SIDHistory.

Windows Server 2003 Interim Domain Functional Level

The Windows Server 2003 interim domain functional level is the preferred method of supporting Windows NT environments during the course of an upgrade. This level *only* applies to a transition from Windows NT to Windows Server 2003 because it does not allow for the presence of Windows 2000 DCs. It is characterized by lowered security features and defaults, similar to the Windows 2000 mixed domain functional level, and a high compatibility level for Windows NT.

- Microsoft Windows 2000 DCs are not supported.
- New Active Directory domain features *not* supported in this level:
 - Group nesting
 - Universal Groups
 - SIDHistory
 - Converting groups between Security Groups and Distribution Groups
 - Domain controller rename
 - Logon timestamp attribute updated and replicated

- User password support on the *InetOrgPerson objectClass*
- Constrained delegation
- Users and Computers container redirection
- Can only be raised to the Windows Server 2003 domain level.
- Can never be lowered since the Windows 2003 interim domain level only exists during an upgrade from Windows NT 4.0 to Windows Server 2003, bypassing Windows 2000.
- Reasons to use the Windows 2003 interim domain functional level:
 - Upgrading a Windows NT 4.0 domain *directly* to Windows 2003.
 - Windows NT 4.0 BDCs will not upgrade immediately.
 - Your Windows NT 4.0 domain contains groups with more than 5000 members, not including the Domain Users group.
 - You have no plans to implement Windows 2000 DCs at any time.
 - Since the Windows 2003 interim domain level greatly improves group replication efficiency, it is better to upgrade directly from Windows NT 4.0 to Windows Server 2003 instead of to Windows 2000, and then to 2003.

In the *Windows Server 2003* interim domain functional level, no domainwide features are activated, although many forest level features are activated at this level (see the section *Windows Server 2003 Interim Forest Functional Level* later in the chapter). This mode is only used during the upgrade of Windows NT 4.0 DCs to Windows Server 2003 DCs. If a Windows 2000 Active Directory domain already exists, then the Windows Server 2003 interim domain level cannot be achieved.

Remember that any domain joined to an existing forest inherits its domain functional level from the child, top-level, or root-level domain that it connects to during the joining process. The domain level of Windows 2000 is only the default when you create a new forest root.



EXAM WARNING

Remember the difference between domain and forest functional levels of the same name. For example, the Windows Server 2003 interim *domain* functional level *can never* be reversed. The Windows Server 2003 interim *forest* functional level *can* be reversed temporarily for the purpose of joining a Windows NT 4.0 domain as a new domain in an existing forest *during an upgrade* of the NT 4.0 domain to the Windows Server 2003 level. To revert your Windows Server 2003 forest back to the interim level for an upgrade, you must manually configure the forest level with LDAP tools such as Ldp.exe or Adsiedit.msc, and then back again. As you can see from this example, domain functional levels and forest functional levels are not the same.

Windows Server 2003 Domain Functional Level

The Windows Server 2003 domain functional level is the ideal level. This level does not allow for the presence of Windows NT or Windows 2000 DCs. It starts out with the best security defaults and capabilities, and the least compatibility with earlier versions of windows. All new 2003 Active Directory domain features are enabled at this level, providing the most efficient and productive environment.

- DCs *not* supported at this level:
 - Windows NT 4.0 DCs
 - Windows 2000 DCs
- All new Active Directory domain features are supported at this level.
- Cannot be raised to any other level, since no higher level exists at this time.
- Can never be lowered to the Windows 2000 mixed mode, the Windows 2000 native mode, or the Windows Server 2003 interim level.

In the Windows Server 2003 domain functional level, only Windows Server 2003 DCs can exist.

Forest Functionality

The Windows Server 2003 *forest functional levels* are named similarly to the domain levels. Windows 2000 originally had only one level, and that level was carried over into Windows 2003. The two other available functional levels are Windows Server 2003 interim and Windows Server 2003, sometimes referred to as Windows Server 2003 native mode. Table 4.2 summarizes the levels, DCs supported in each level, and the level's primary purpose.

As with domain functional levels, each forest functional level carries over the features from lower levels, and activates new features as well. These new features apply across every domain in your forest. After you raise the forest functional level, earlier OSs cannot be promoted to DCs. For example, Windows NT 4.0 and Windows 2000 DCs cannot be part of the forest at any level, except through external or forest trusts, once the forest level has been raised to Windows Server 2003 native. See Table 4.4 for a summary of the capabilities of the new Windows Server 2003 forest functional levels.

Table 4.4 New Forest Functional Level Features

Forest Feature	Windows 2000	Windows Server 2003 Interim	Windows Server 2003 Native
Support for more than 5000 members per group	Not available	Enabled	Enabled
Universal Group caching	Enabled	Enabled	Enabled
Application partitions	Enabled	Enabled	Enabled

Continued

Table 4.4 New Forest Functional Level Features

Forest Feature	Windows 2000	Windows Server 2003 Interim	Windows Server 2003 Native
Install from backups	Enabled	Enabled	Enabled
Quotas	Enabled	Enabled	Enabled
Rapid GC demotion	Enabled	Enabled	Enabled
SIS for system access control lists (SACL) in the Jet Database Engine	Enabled	Enabled	Enabled
Improve topology generation event logging	Enabled	Enabled	Enabled
Windows Server 2003 DC assumes the Intersite Topology Generator (ISTG) role	Enabled	Enabled	Enabled
Efficient group member replication using linked value replication	Disabled	Enabled	Enabled
Improved KCC inter-site replication topology generator algorithms	Disabled	Enabled	Enabled
ISTG aliveness no longer replicated	Disabled	Enabled	Enabled
Attributes added to the GC, such as: <i>ms-DS-Entry-Time-To-Die</i> , <i>Message Queuing-Secured-Source</i> , <i>Message Queuing-Multicast-Address</i> , <i>Print-Memory</i> , <i>Print-Rate</i> , and <i>Print-Rate-Unit</i>	Disabled	Enabled	Enabled
Defunct schema objects	Disabled	Disabled	Enabled
Cross-forest trust	Disabled	Disabled	Enabled
Domain rename	Disabled	Disabled	Enabled
Dynamic auxiliary classes	Disabled	Disabled	Enabled
<i>InetOrgPerson objectClass</i> change	Disabled	Disabled	Enabled
Application groups	Disabled	Disabled	Enabled
15-second intrasite replication frequency for Windows Server 2003 DCs upgraded from Windows 2000	Disabled	Disabled	Enabled

Continued

Table 4.4 New Forest Functional Level Features

Forest Feature	Windows 2000	Windows Server 2003 Interim	Windows Server 2003 Native
Reduced NTDS.DIT size	Disabled	Disabled	Enabled
Unlimited site management	Disabled	Disabled	Enabled

Windows 2000 Forest Functional Level (default)

The Windows 2000 forest functional level is primarily designed to support mixed environments during the course of an upgrade. Typically, this applies to a transition from Windows 2000 to Windows Server 2003. It is also the default mode for a newly created Windows Server 2003 domain. It is characterized by relatively lower security features and reduced efficiency, but maintains the highest compatibility level possible for Active Directory. The Windows 2003 interim forest functional level handles upgrades from Windows NT to Windows Server 2003.

- All Windows DCs are supported.
- Active Directory forest features *not* supported in this mode:
 - Efficient group member replication using linked value replication
 - Improved KCC inter-site replication topology generator algorithms
 - ISTG aliveness no longer replicated
 - Attributes added to the GC, such as *ms-DS-Entry-Time-To-Die*, *Message Queuing-Secured-Source*, *Message Queuing-Multicast-Address*, *Print-Memory*, *Print-Rate*, and *Print-Rate-Unit*
 - Defunct schema objects
 - Cross-forest trust
 - Domain rename
 - Dynamic auxiliary classes
 - *InetOrgPerson objectClass* change
 - Application groups
 - 15-second intra-site replication frequency for Windows Server 2003 DCs upgraded from Windows 2000
 - Reduced NTDS.DIT size
 - Unlimited site management
- Can only be raised to the Windows 2003 native forest level.
- Can never be lowered back to the Windows 2000 level.

In the *Windows 2000* functional level, which is the default level, Windows 2000 and greater DCs can exist, as well as Windows NT BDCs. Newly created Windows Server 2003 forests always start at this level. Windows NT PDCs do not exist in any version of Active Directory. Features available in the Windows 2000 forest functional level of Windows Server 2003 carry over the old features and add many new ones.

Windows Server 2003 Interim Forest Functional Level

The Windows Server 2003 interim forest functional level is the preferred method of supporting Windows NT environments during the course of an upgrade. This level *only* applies to a transition from Windows NT to Windows Server 2003 because it does not allow for the presence of Windows 2000 DCs anywhere in the forest. It is characterized by lowered security features and defaults, but provides many efficiency improvements over the Windows 2000 forest functional level.

- Microsoft Windows 2000 DCs are not supported.
- New Active Directory forest features *not* supported in this level:
 - Defunct schema objects
 - Cross-forest trust
 - Domain rename
 - Dynamic auxiliary classes
 - *InetOrgPerson objectClass* change
 - Application groups
 - 15-second intrasite replication frequency for Windows Server 2003 DCs upgraded from Windows 2000
 - Reduced NTDS.DIT size
 - Unlimited site management
- Can only be raised to the Windows Server 2003 forest functional level.
- Can never be lowered to the Windows 2000 level, since the Windows 2003 interim domain level only exists during an upgrade *directly* from Windows NT 4.0 to Windows Server 2003.
- Reasons to use the Windows 2003 interim forest functional level:
 - Upgrading a root Windows NT 4.0 domain *directly* to Windows 2003.
 - Windows NT 4.0 BDCs will not upgrade immediately.
 - Your Windows NT 4.0 domain contains groups with more than 5000 members, not including the Domain Users group.
 - You have no plans to implement Windows 2000 DCs at any time.

- Since the Windows 2003 interim domain level greatly improves group replication efficiency, it is better to use the Windows 2003 interim forest functional level instead of upgrading to Windows 2000, and then to Windows 2003.
- You are upgrading a Windows NT 4.0 PDC as the first DC of a new root domain in an existing Windows Server 2003 forest.
- Simultaneously upgrading and joining a Windows NT 4.0 domain as a child domain in an existing Windows Server 2003 forest.

In the *Windows Server 2003* interim forest functional level, unlike the Windows Server 2003 interim domain functional level, many new features are activated while still allowing Windows NT 4.0 BDC replication. This mode is only used during the upgrade of a Windows NT 4.0 domain to a Windows Server 2003 forest. If a Windows 2000 Active Directory forest already exists, then the Windows Server 2003 interim forest level cannot be achieved.

To revert your Windows Server 2003 forest back to the interim level for an upgrade, you must manually configure the forest level with LDAP tools such as Ldp.exe or Adsiedit.msc. Remember that any domain joined to an existing forest inherits its domain functional level from the child, top-level, or root-level domain that it connects to during the joining process. The default forest level of Windows 2000 only applies when you create a new forest.

Windows Server 2003 Forest Functional Level

The Windows Server 2003 forest functional level is the ideal level. This level does not allow for the presence of Windows NT or Windows 2000 DCs anywhere in the forest. It starts out with the best security defaults and capabilities, and the least compatibility with earlier versions of Windows. All new 2003 Active Directory forest features are enabled at this level, providing the most efficient and productive environment.

- DCs *not* supported at this level:
 - Windows NT 4.0 DCs
 - Windows 2000 DCs
- All new Active Directory forest features are supported at this level.
- Cannot be raised to any other level, since no higher level exists in Windows Server 2003.
- Can never be lowered back to the Windows 2000 level, but can be temporarily lowered to the Windows Server 2003 interim level for the purpose of joining a Windows NT 4.0 domain as a new domain in an existing forest *during an upgrade* of the NT 4.0 domain to the Windows Server 2003 level.

In the Windows Server 2003 forest functional level, only Windows Server 2003 DCs can exist.

EXAM
70-294
OBJECTIVE
1.3.5

Raising the Functional Level of a Domain and Forest

Before increasing a functional level, you should prepare for it by performing the following tasks. First, inventory your entire forest for earlier versions of DCs. The Active Directory Domains and Trusts MMC snap-in can generate a detailed report should you need it. You can also perform a custom LDAP query from the Active Directory Users and Computers MMC snap-in that will discover Windows NT DC objects within the forest. Use the following search string:

```
(&(objectCategory=computer)(operatingSystem Version=4*)
(userAccountControl:1.2.840.113556.1.4.803:=8192))
```

There should be no spaces in the query, and type it in all on one line. The search string is shown on two lines for readability.

Second, you need to *physically* locate all down-level DCs for the new functional level in the domain or forest as needed, and either upgrade or remove them.

Third, verify that end-to-end replication is working in the forest using the Windows Server 2003 versions of Repadmin.exe and Replmon.exe.

Finally, verify the compatibility of your applications and services with the version of Windows that your DCs will be running, and specifically their compatibility with the target functional level. Use a lab environment to test for compatibility issues, and contact the appropriate vendors for compatibility information.

Domain Functional Level

Before raising the functional level of a domain, all DCs must be upgraded to the minimum OS level as shown in Table 4.2. Remember that when you raise the domain functional level to Windows 2000 native or Windows Server 2003, it can never be changed back to Windows 2000 mixed mode. Exercise 4.01 takes you systematically through the process of verifying the current domain functional level. Exercise 4.02 takes you through the process of raising the domain functional level. To raise the level, you must be an enterprise administrator, a domain administrator in the domain you want to raise, or have the appropriate authority.

EXERCISE 4.01

VVERIFYING THE DOMAIN FUNCTIONAL LEVEL

1. Log on as a Domain Admin of the domain you are checking.
2. Click on **Start | Control Panel | Performance and Maintenance | Administrative Tools | Active Directory Users and Computers**, or use the Microsoft Management Console (MMC) preconfigured with the Active Directory Users and Computers snap-in.

3. Locate the domain in the console tree that you are going to raise in functional level. Right-click the domain and select **Raise Domain Functional Level**.
 4. In the Raise Domain Functional Level dialog box, the current domain functional level appears under **Current domain functional level**.
-



NOTE

This check can also be performed using the Active Directory Domains and Trusts MMC snap-in.

EXERCISE 4.02

RAISING THE DOMAIN FUNCTIONAL LEVEL

1. Log on locally as a Domain Admin to the PDC or the PDC Emulator FSMO of the domain you are raising.
 2. Click on **Start | Administrative Tools | Active Directory Domains and Trusts**, or use the MMC preconfigured with the Active Directory Domain and Trusts snap-in.
 3. Locate the domain in the console tree that you are going to raise in functional level. Right-click the domain and select **Raise Domain Functional Level**.
 4. A dialog box will appear entitled **Select an available domain functional level**. There are only two possible choices, although both might not be available.
 - Select **Windows 2000 native**, and then click the **Raise** button to raise the domain functional level to Windows 2000 native.
 - Select **Windows Server 2003**, and then click the **Raise** button to raise the domain functional level to Windows Server 2003.
-

**NOTE**

You can also use the Active Directory Users and Computers MMC snap-in to perform this exercise.

Forest Functional Level

Before raising the functional level of a forest, all DCs in the forest must be upgraded to the minimum OS level as shown in Table 4.2. In practice, since the only forest functional level that will be available to you is Windows Server 2003, all DCs in the forest must be running Windows Server 2003. Locate all down-level DCs and either upgrade them or remove them from the domain. You do not have to upgrade the domain functional level before the forest functional level. The reason for this is that all domains in the forest will automatically raise to the level of Windows Server 2003 to match the forest level after Active Directory replicates the changes. The forest Schema Master performs this operation. Exercise 4.03 takes you through the process of verifying the current forest functional level. Exercise 4.04 shows you how to raise the forest functional level. To raise the forest level, you must be an enterprise administrator, a domain administrator at the forest root, or have the appropriate authority.

EXERCISE 4.03

VERIFYING THE FOREST FUNCTIONAL LEVEL

1. Log on as an Enterprise Administrator in the forest you are checking.
 2. Click on **Start | Administrative Tools | Active Directory Domains and Trusts**, or use the MMC preconfigured with the Active Directory Domains and Trusts snap-in.
 3. In the console tree, right-click the **Active Directory Domains and Trusts** folder and select **Raise Forest Functional Level**.
 4. In the Raise Forest Functional Level dialog box, the current forest functional level appears under **Current forest functional level**.
-

EXERCISE 4.04

RAISING THE FOREST FUNCTIONAL LEVEL

1. Log on locally as an Enterprise Administrator on the PDC Emulator FSMO of the forest root domain you are raising.
2. Click on **Start | All Programs | Administrative Tools | Active Directory Domains and Trusts**, or use the MMC preconfigured with the Active Directory Domains and Trusts snap-in.
3. In the console tree, right-click the **Active Directory Domains and Trusts** folder and select **Raise Forest Functional Level**.
4. Where it asks you to Select an available forest functional level, click Windows Server 2003, and then click the Raise button.

Optimizing Your Strategy for Raising Functional Levels

There are two basic strategies for traveling the path from the Windows 2000 native level and Windows 2000 mixed-mode levels to the goal of Windows Server 2003 functional levels across your forest.

- The Windows 2000 native mode path.
 - Raise the level of all domains to the Windows 2000 native functional level.
 - Raise the forest level to Windows Server 2003.

Benefits of this method include:

- You do not have to perform the domain level-raising procedure on every domain before raising the forest level.
- It automatically does the work of tracking down all down-level domains and DCs for you. The process fails if these exist, but then you have a ready list of preparation work to do. This is helpful if your forest is not well documented. See the sidebar *If Raising the Forest Functional Level Fails* for more information.
- The Windows Server 2003 level path.
 - Raise the level of all domains to the Windows 2000 native functional level.
 - Raise the level of all domains to the Windows Server 2003 functional level.
 - Raise the forest level to Windows Server 2003.

The benefits of this method are:

- All of the new Windows Server 2003 domain-level features are turned on before you make the commitment to raising the level of the forest.
- You can perform integration and interoperability testing on a smaller scale without committing the forest to the functional upgrade.

There are three basic approaches for the use of interim modes when upgrading Windows NT to Windows Server 2003. Interim level should be avoided if you will ever have a need to implement Windows 2000 DCs. Here are the three strategies:

- When upgrading the Windows NT PDC into a *new* Windows Server 2003 forest, select the interim level from the `dcpromo` utility.
- When upgrading the Windows NT PDC into an *existing* Windows Server 2003 forest, manually set the interim level with `Ldp.exe` or `Adsiedit.msc`, and join the forest during the upgrade. The upgraded domain inherits the interim setting from the forest.
- Upgrade or remove all Windows NT BDCs, and then upgrade the Windows NT PDC. Since no Windows NT DCs remain in the domain, the Windows Server 2003 interim functionality level is not needed.

If Raising the Forest Functional Level Fails

Typically, when the procedure for raising the forest functional level fails, it is because of the existence of DCs somewhere in the forest that must still be upgraded from Windows NT 4.0 or Windows 2000. To find out, click **Save As** right after the failure from the **Raise Forest Functional Level** dialog box. This will save a helpful log file to aid you in locating the DCs and domains that do not meet the requirements.

If you have verified that all DCs have been upgraded to Windows Server 2003 but continue to get a failure, consider the following two possibilities:

- You have upgraded the DC in question, but Active Directory replication has not yet completed.
- The DC causing the error is no longer in service, but its corresponding computer object has not been removed from the domain.

Go to “Q” article number 216498 in the Microsoft Knowledge Base for information on how to remove data in Active Directory after an unsuccessful DC demotion. This is a likely cause of the error.



Creating the Forest and Domain Structure

The process of creating the forest and domain structure is centered on the use of the Active Directory Installation Wizard. This utility installs and configures DCs, which in turn provide the Active Directory directory service to networked computers and users. The first step is to install Windows Server 2003 as a member server or a stand-alone server. At this point, you should be familiar with that process, so it will not be covered here.

Next comes the decision process leading to the installation of a DC. Essentially, there are two reasons to install a DC: to create a new domain, or to add an additional DC to an existing domain. Depending on your current forest structure, you will end up with one of four results:

- A new forest
- A new domain tree in an existing forest
- A new child domain in an existing domain
- A new DC in an existing domain

Deciding When to Create a New DC

Since a domain cannot exist without a DC, you must create at least one for each domain. The process of creating the first DC also creates the domain itself. The domain can be either a new child domain or the root of a new tree. The difference is in the namespace. See the section *Domain Trees earlier* in the chapter. Here are the four main reasons to create a new DC:

- Creating the first domain in your network
- Creating a new domain in your forest
- Improving a domain's reliability
- Improving network performance between sites

If you want to create a domain with a name that is not related to any other namespace in your forest, you will create a new tree. If you want to create a domain that will function as an additional subunit within an existing domain, you will create a child domain.

To improve a domain's reliability, you should always create at least a second DC in each domain. That way, if the first one fails, you will still be able to use the second. If your existing DCs are overloaded, simply adding another DC to your domain will help spread the load. If any of your domains are divided by WAN links, then it is a good practice to place a DC in each site. Besides lowering WAN bandwidth utilization and improving logon response times, you also provide a level of fault tolerance. If the WAN link fails, users on both sides of the link will continue to be able to log on if your domain is at the Windows Server 2003 functional level.

Legacy DNS Compatibility

You might find it useful at times to interoperate with a legacy DNS server, such as that found on older UNIX servers and Windows NT, especially during a migration. One way to accomplish this is to configure the legacy server to support Active Directory. Every Windows Server 2003 DC creates a file in the %Systemroot\System32\Config folder called Netlogon.dns. It contains the DNS records necessary to register the resources of the DC.

If you are using a DNS server that does not support dynamic updates, but *does* supports the SRV resource record, you should import the Netlogon.dns records from all of your Active Directory DCs into the primary zone file corresponding to the namespace your Active Directory is using. Windows DNS clients configured with the legacy DNS server's IP address will then use the static SRV records to locate a DC to begin the logon process. The flaw in this setup is that any time one of your DCs is offline, you will get some number of logon failures until you manually remove the corresponding SRV records.

In this type of environment, DNS issues and misconfigurations are common. For example, if you try to execute any of the Active Directory MMC snap-ins, such as Active Directory Users and Computers, and you get an error saying, "No authority can be contacted for authentication," then DNS might be incorrectly configured. First, make sure that your DNS client is using the intended DNS server. Verify that the zones on your DNS server are configured correctly, that your DNS server has authority for the zone that contains your Active Directory namespace, and that the correct SRV records have been imported into the applicable primary zone file.

Active Directory requires DNS. The Active Directory Installation Wizard will look for an authoritative DNS server that accepts dynamic updates for the domain. If a DNS server that can accept dynamic updates is not available, the Active Directory Installation Wizard will optionally create one for you that is preconfigured for the name of your domain. When you restart the new DC, it will register itself with DNS.

Installing Domain Controllers

You should know what type of domain you want to install before you begin, and the namespace it will use. Read the procedure for the type of domain you want to install and know what your responses will be. For example, if you want the *shared system volume* (SYSVOL) on its own disk volume, you will need to prepare it ahead of time. Before you run the Active Directory Installation Wizard, make sure that the authoritative DNS zone allows dynamic updates and that your DNS server supports SRV records. As always, no matter how small the domain, it should always have two DCs for fault tolerance and availability.

What Else Do I Need to Know Before I Install a DC?

Before running the Active Directory Installation Wizard, you should know the answers to the questions it will ask. Some of the questions require preparation ahead of time. Here are a few of the important points to consider before typing `dcpromo`.


1. **Where do I put my database and log folders?** The default setting of `%SystemRoot\ntds` for both the Active Directory database and logs is good for small networks, but consider putting these two items on their own disk volume for better performance. The Active Directory is a database and is susceptible to the same performance issues of any other database. The best practice for the `Ntds.dit` database store is to use a RAID-5 arrangement of the appropriate size for your enterprise. As an additional performance enhancement for large enterprises, use two mirrored physical disks for the log folders. Install and configure these disk structures before running `dcpromo`. The Active Directory sizer tool, found at www.microsoft.com/windows2000/techinfo/reskit/tools/new/adsizer-o.asp, allows you to estimate the hardware that you need to deploy Active Directory in your organization. `Adsizer.exe` performs its estimates based on your enterprise usage profile, and your domain and site topologies.
2. **Where do I put the SYSVOL?** The contents of the SYSVOL are replicated to all DCs in the domain. It contains a copy of all the domain's public files. The default setting of `%SystemRoot\sysvol` is acceptable for small networks, but consider putting it on its own disk volume for better performance.
3. Select **Permissions compatible with pre-Windows 2000 server operating systems** only after considering the impact of weaker permissions. This option adds the *Everyone* group to the *Pre-Windows 2000 Compatible Access* group. Don't be hasty; if you aren't sure you need it, select **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**. Even with a pure Windows XP and Server 2003 network, you might find that some of your applications are not working. Manually add the group as described previously. If the problem is fixed, then you have applications that are not compatible with Windows Server 2003 domains. Once those applications are upgraded, remove the backward compatibility setting to enhance security. This will prevent anonymous users from accessing your user and group information.
4. Select a strong password for the **Directory Services Restore Mode Administrator Password**. Use a mixture of numbers, uppercase letters, lowercase letters, and special characters. The best practice is to put the password in a sealed envelope under lock and key for emergencies.

EXAM
70-294
OBJECTIVE
1.3.1

Creating a Forest Root Domain

The initial DC that you install, as shown in Figure 4.3, will provide your users with the Active Directory. Consider making this an empty root domain where your Enterprise Administrators have accounts, but no regular users. With the procedure in Exercise 4.05, you will simultaneously create your first domain, called the root domain, and your first forest.

Figure 4.3 Creating a New Domain in a New Forest



EXERCISE 4.05

CREATING A NEW DOMAIN IN A NEW FOREST

1. Log on as a local Administrator.
2. Click **Start | Run**.
3. Type **dcpromo**.
4. Click **OK** to start the Active Directory Installation Wizard.
5. In the **Welcome to the Active Directory Installation Wizard** window, click **Next** as shown in Figure 4.4.

Figure 4.4 The Welcome Dialog Box for dcpromo



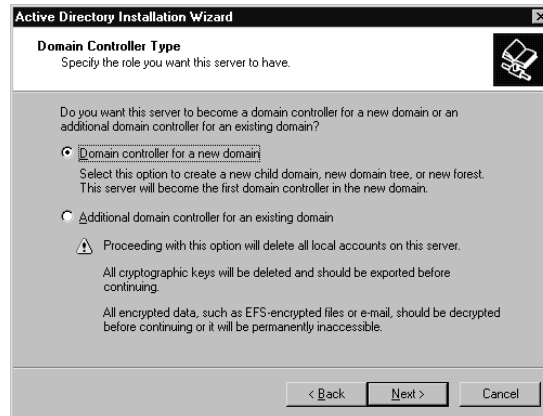
6. In the **Operating System Compatibility** window, click **Next** as shown in Figure 4.5.

Figure 4.5 The Operating System Compatibility Dialog Box for dcpromo



7. In the **Domain Controller Type** window, click **Domain controller for a new domain | Next** as shown in Figure 4.6.

Figure 4.6 The Domain Controller Type Dialog Box Used for a New DC in a New Domain



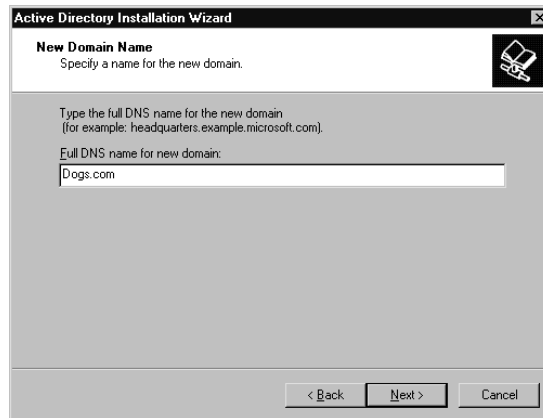
8. In the **Create New Domain** window, click **Domain in a new forest | Next** as shown in Figure 4.7.

Figure 4.7 The Create New Domain Dialog Box Used for Creating a New Domain in a New Forest

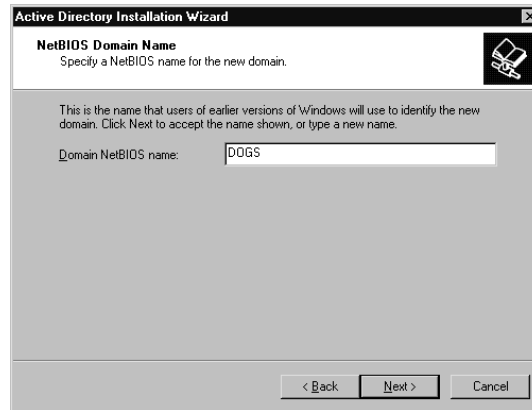


9. In the **New Domain Name** window, type the full DNS domain name for the new domain, and click **Next** as shown in Figure 4.8.

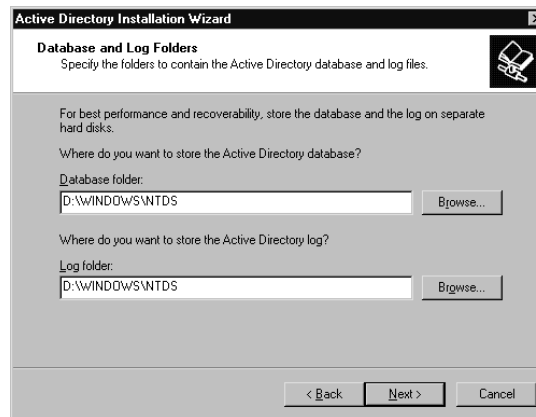
Figure 4.8 The New Domain Name Dialog Box Used for Creating a New Domain in a New Forest



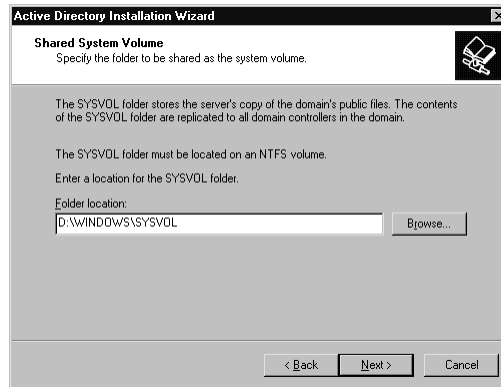
10. In the **NetBIOS Domain Name** window, verify the NetBIOS name and click **Next** as shown in Figure 4.9. The default name is generally the best one to use.

Figure 4.9 The NetBIOS Domain Name Dialog Box for dcpromo

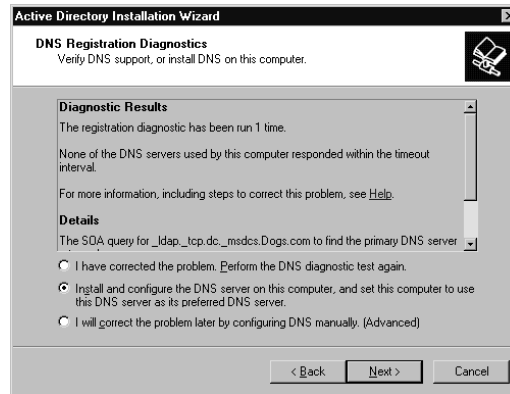
11. In the **Database and Log Folders** window, type or browse to the location where you want the database and log folders. Click **Next** as shown in Figure 4.10.

Figure 4.10 The Database and Log Folders Dialog Box for dcpromo

12. In the **Shared System Volume** window, type or browse to the location where you want the SYSVOL folder. Click **Next** as shown in Figure 4.11.

Figure 4.11 The Shared System Volume Dialog Box for dcpromo

13. In the **DNS Registration Diagnostics** window, verify an existing DNS server to be authoritative for this new forest, or click **Install** and configure the DNS server on this computer, and set this computer to use this DNS server as its preferred DNS server. Click **Next** as shown in Figure 4.12.

Figure 4.12 The DNS Registration Diagnostics Window with No Current DNS Server Available

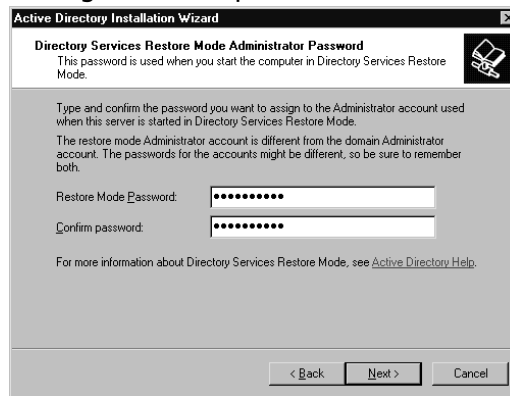
14. In the Permissions window, you have two options: **Permissions compatible with pre-Windows 2000 server operating systems** and **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**. Select one, and then click **Next** as shown in Figure 4.13.

Figure 4.13 The Permissions Dialog Box for dcpromo



15. In the **Directory Services Restore Mode Administrator Password** window, input and confirm the password for the **Directory Services Restore Mode**. Click **Next** as shown in Figure 4.14.

Figure 4.14 The Directory Services Restore Mode Administrator Password Dialog Box for dcpromo



16. Read the **Summary** window. Click **Next** as shown in Figure 4.15. The installation will continue for several minutes.
17. Restart your new DC.
18. Verify that the installation was successful. Open a command prompt and enter the **Net Share** command. It should report the existence of the **Netlogon** and **SYSVOL** shares. To verify that the DNS service locator records for the new DC were successfully created, follow these steps:

1. Click **Start | Administrative Tools | DNS** to start the DNS administrator console.
2. Expand the server name.
3. Expand **Forward Lookup Zones**.
4. Expand the domain.
5. Verify that the `_msdcs`, `_sites`, `_tcp`, and `_udp` folders are present and contain records for your new DC. These service location records are crucial to the operation of the DC. See Table 4.7 for a more detailed description of the required records, and Figure 4.16 for a view of the DNS administrator tool used to view them.

Figure 4.15 The Summary Window Describing the Forest Root

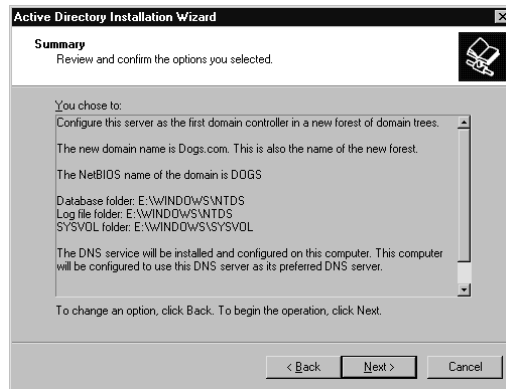
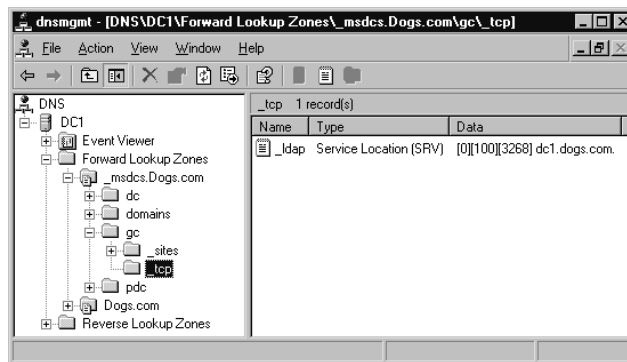


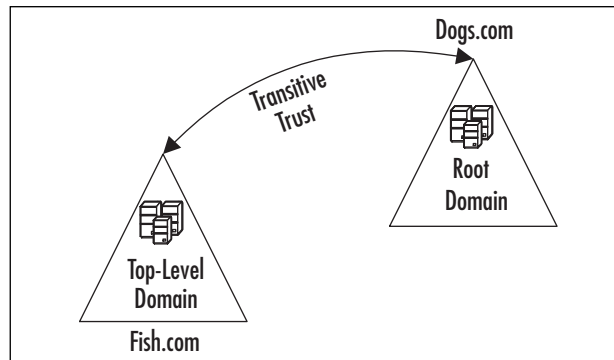
Figure 4.16 The DNS Administrator Tool Used to Verify a Successful Forest-Root Installation



Creating a New Domain Tree in an Existing Forest

This will often be the second domain that you install, as shown in Figure 4.17. This type of arrangement accommodates a forest comprised of two different company divisions, or two companies within a larger corporation. Domains are used as boundaries for security and administration. With the procedure in Exercise 4.06, you will simultaneously create your first nonroot top-level domain, and the second tree in your forest. Note that a new bidirectional, transitive trust is automatically created with the forest root.

Figure 4.17 Creating a New Domain Tree in an Existing Forest

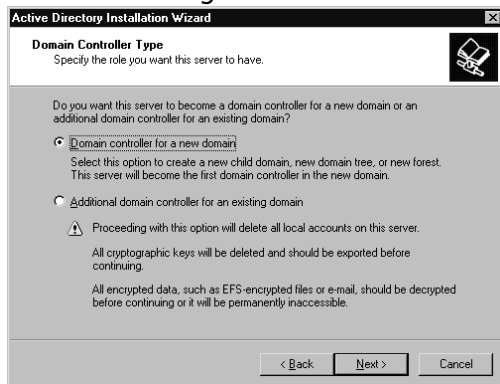


EXERCISE 4.06

CREATING A NEW DOMAIN TREE IN AN EXISTING FOREST

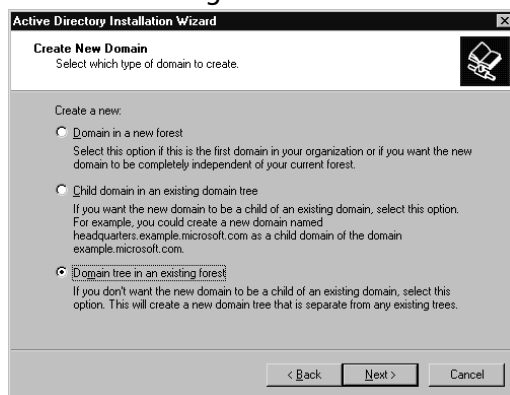
1. Log in as a local Administrator.
2. Click **Start | Run**.
3. Type **dcpromo**.
4. Click **OK** to start the Active Directory Installation Wizard.
5. In the **Welcome to the Active Directory Installation Wizard** window, click **Next**.
6. In the **Operating System Compatibility** window, click **Next**.
7. In the **Domain Controller Type** window, click **Domain controller for a new domain | Next**, as shown in Figure 4.18.

Figure 4.18 The Domain ControllerType Dialog Box Used for a New Domain Tree in an Existing Forest



8. In the **Create New Domain** window, click **Domain in an existing forest** | **Next**, as shown in Figure 4.19.

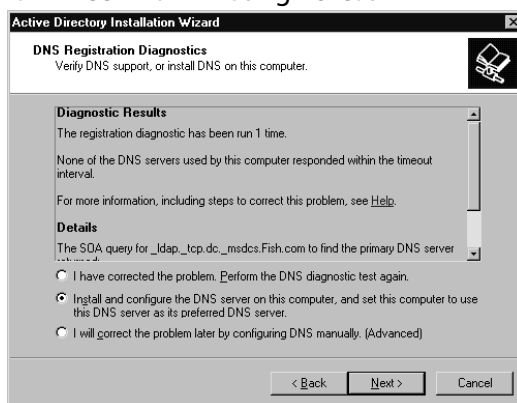
Figure 4.19 The Create New Domain Dialog Box Used for a New Domain Tree in an Existing Forest



9. In the **Network Credentials** window, type in the username, password, and domain name of an Enterprise Administrator or Domain Admin in the forest-root domain. Click **Next**.
10. In the **New Domain Tree** window, type the full DNS domain name for the new domain, and click **Next**.
11. In the **NetBIOS Domain Name** window, verify the NetBIOS name and click **Next**. The default name is generally the best one to use.
12. In the **Database and Log Folders** window, type or browse to the location where you want the database and log folders. Click **Next**.

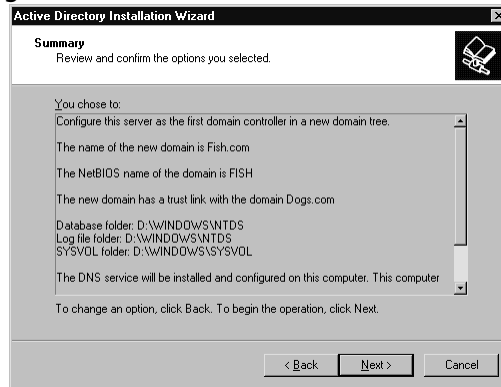
13. In the **Shared System Volume** window, type or browse to the location where you want the SYSVOL folder. Click **Next**.
14. In the **DNS Registration Diagnostics** window, configure an existing DNS server to be authoritative for this tree, or click **Install and configure the DNS server on this computer, and set this computer to use this DNS server as its preferred DNS server**. Click **Next**, as shown in Figure 4.20.

Figure 4.20 The DNS Registration Diagnostics Dialog Box Used for a New Domain Tree in an Existing Forest



15. In the Permissions window you have two options: **Permissions compatible with pre-Windows 2000 server operating systems** and **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**. Select one, and then click **Next**.
16. In the **Directory Services Restore Mode Administrator Password** window, input and confirm the password for the **Directory Services Restore Mode**. Click **Next**.
17. Read the **Summary** window. Click **Next**, as shown in Figure 4.21. The installation will continue for several minutes.
18. Restart your new DC.
19. Verify that the installation was successful. Open a command prompt and enter the **Net Share** command. It should report the existence of the Netlogon and SYSVOL shares. To verify that the DNS service locator records for the new DC were successfully created, follow these steps:

Figure 4.21 The Summary Dialog Box Used for a New Domain Tree in an Existing Forest



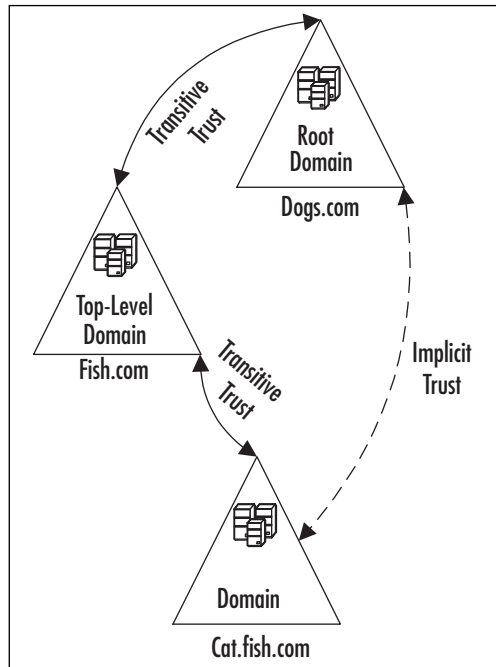
1. Click **Start | Administrative Tools | DNS** to start the DNS administrator console.
2. Expand the server name.
3. Expand Forward Lookup Zones.
4. Expand the domain.
5. Verify that the `_msdcs`, `_sites`, `_tcp`, and `_udp` folders are present and contain records for your new DC. These service location records are crucial to the operation of the DC. See Table 4.7 for a more detailed description of the required records, and Figure 4.16 for a view of the DNS administrator tool used to view them.

EXAM
70-294

OBJECTIVE
1.3.2

Creating a New Child Domain in an Existing Domain

This will often be the third domain that you install, as shown in Figure 4.22. This type of arrangement accommodates a tree comprised of two different company groups, sometimes in physically separate locations. Since domains are used as boundaries for security and administration, there are many reasons for segregating a subgroup. If a group requires higher or lower levels of security, or if a different group of administrators requires complete control, then a child domain is a good idea. With the procedure in Exercise 4.07, you will create your first child domain within any existing tree of your forest. Note that a new bidirectional, transitive trust is automatically created with the parent domain, simultaneously creating an implicit trust between the child domain and the forest-root. See the section *Establishing Trust Relationships* later in this chapter for more information on the transitivity and direction of trusts.

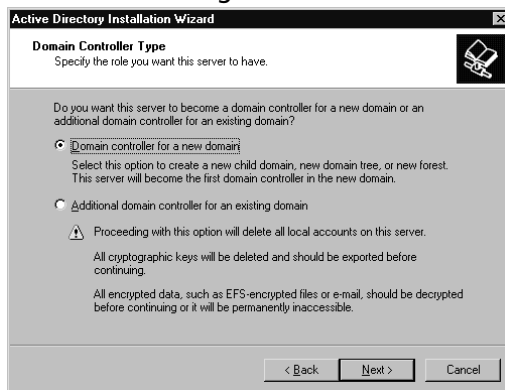
Figure 4.22 Creating a New Child Domain in an Existing Domain

EXERCISE 4.07

CREATING A NEW CHILD DOMAIN IN AN EXISTING DOMAIN

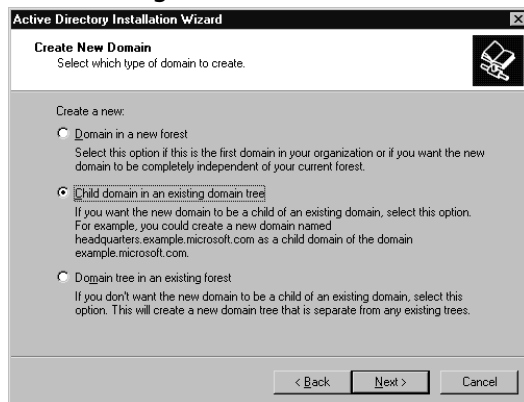
1. Log in as a local Administrator.
2. Click **Start | Run**.
3. Type **dcpromo**.
4. Click **OK** to start the Active Directory Installation Wizard.
5. In the **Welcome to the Active Directory Installation Wizard** window, click **Next**.
6. In the **Operating System Compatibility** window, click **Next**.
7. In the **Domain Controller Type** window, click **Domain controller for a new domain | Next**, as shown in Figure 4.23.

Figure 4.23 The Domain ControllerType Dialog Box Used for a New Child Domain in an Existing Domain



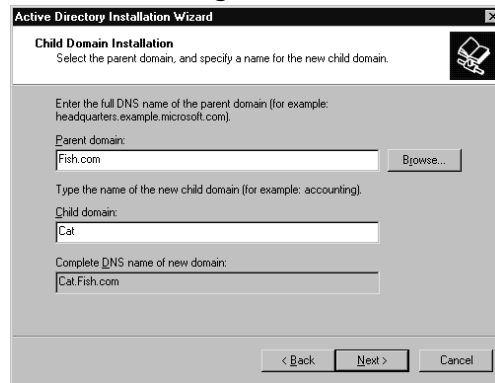
8. In the **Create New Domain** window, click **Child Domain in an existing domain tree** | **Next**, as shown in Figure 4.24.

Figure 4.24 The Create New Domain Dialog Box Used for a New Child Domain in an Existing Domain



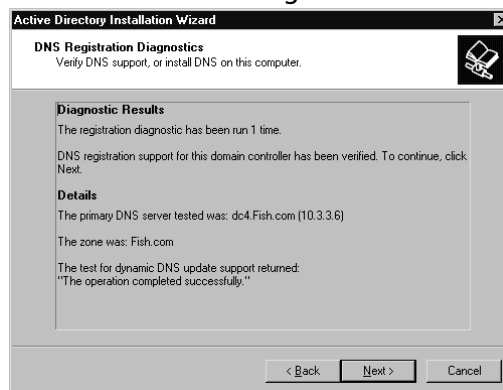
9. In the **Network Credentials** window, type in the username, password, and domain name of an Enterprise Administrator in the forest-root domain, or a Domain Admin in the parent domain, and click **Next**.
10. In the **Child Domain Installation** window, verify the parent domain and type the name of the new child domain. Click **Next**, as shown in Figure 4.25.

Figure 4.25 The Child Domain Installation Dialog Box Used for a New Child Domain in an Existing Domain



11. In the **NetBIOS Domain Name** window, verify the NetBIOS name and click **Next**. The default name is generally the best one to use.
12. In the **Database and Log Folders** window, type or browse to the location where you want the database and log folders. Click **Next**.
13. In the **Shared System Volume** window, type or browse to the location where you want the SYSVOL folder. Click **Next**.
14. In the **DNS Registration Diagnostics** window, check to see if the DNS settings are correct for the parent domain. Click **Next**, as shown in Figure 4.26.

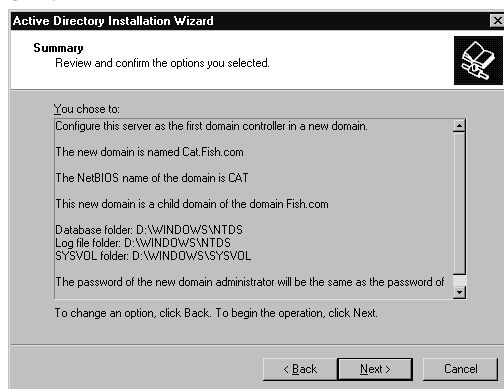
Figure 4.26 The DNS Registration Diagnostics Dialog Box Used for a New Child Domain in an Existing Domain



15. In the Permissions window you have two options: **Permissions compatible with pre-Windows 2000 server operating systems** and **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**. Select one, and then click **Next**.

16. In the **Directory Services Restore Mode Administrator Password** window, input and confirm the password for the Directory Services Restore Mode. Click **Next**.
17. Read the **Summary** window. Click **Next**, as shown in Figure 4.27. The installation will continue for several minutes.

Figure 4.27 The Summary Dialog Box for a New Child Domain in an Existing Domain



18. Restart your new DC.
19. Verify that the installation was successful. Open a command prompt and enter the **Net Share** command. It should report the existence of the *Netlogon* and *SYSVOL* shares. To verify that the DNS service locator records for the new DC were successfully created, follow these steps:
 1. Click **Start | Administrative Tools | DNS** to start the DNS administrator console.
 2. Expand the server name.
 3. Expand **Forward Lookup Zones**.
 4. Expand the domain.
 5. Verify that the *_msdcs*, *_sites*, *_tcp*, and *_udp* folders are present and contain records for your new DC. These service location records are crucial to the operation of the DC. See Table 4.7 for more a more detailed description of the required records, and Figure 4.16 for a view of the DNS administrator tool used to view them.

What If dcpromo Fails?

There are a few things to check for after a failure of the Active Directory Installation Wizard. First, check the contents of the %SystemRoot%\Debug\Dcpromo.log file.


- If the log file reports that *The system cannot find the file specified*, then check for the presence of the %SystemRoot%\System32\Ntds.dit file. This is a default directory services file on a member server. The way to fix this is to expand **Ntds.di_** from any server CD. Note that this file should be in the System32 folder on a member server. Once you run a successful dcpromo, the active Ntds.dit file will be in the folder you specified during the promotion.
- If you receive an *Access is denied* error, check for incorrect permissions on the default Ntds.dit file, as well as on your new and existing NTDS folders.
- If SRV records fail to show up in the appropriate DNS zone, check first to see if the new DC's Primary DNS Server TCP/IP property is set to the correct DNS server. If the DC *is* a DNS server, then this value should point to itself.

Creating a New DC in an Existing Domain

This is the only situation where you will run the Active Directory Installation Wizard without creating a new domain. See Exercise 4.08 for the necessary steps. Usually, you will need to perform this procedure when your domain has grown to the point that it needs additional DCs to spread the workload.

Using the Conventional Across-the-Network Method

In Windows 2000 Server, this is the standard method. Copying the Active Directory database along with other state data over the network during DC promotion is now an option in Windows Server 2003, as illustrated in Figure 4.28. The other option is to run dcpromo with a backup file. Note that several wizard choices are missing from this procedure in comparison to the earlier ones because you are not creating a new domain, and that different windows appear based on your use of the */ADV* option. The */ADV* switch with dcpromo is required for promoting from a backup file, but is optional in the across-the-network method.

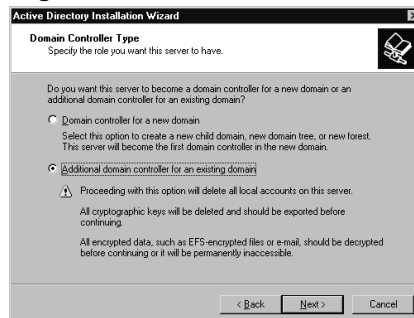
Figure 4.28 Using the Conventional Across-the-Network Method

EXERCISE 4.08

CREATING A NEW DOMAIN CONTROLLER IN AN EXISTING DOMAIN USING THE CONVENTIONAL ACROSS-THE-NETWORK METHOD

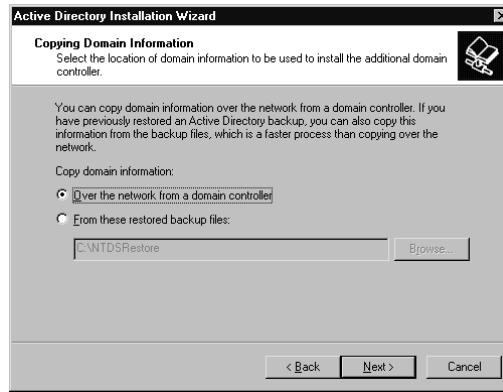
1. Log in as a local Administrator.
2. Click **Start | Run**.
3. Type `dcpromo /adv`.
4. Click **OK** to start the Active Directory Installation Wizard.
5. In the **Welcome to the Active Directory Installation Wizard** window, click **Next**.
6. In the **Operating System Compatibility** window, click **Next**.
7. In the **Domain Controller Type** window, click **Additional domain controller for an existing domain | Next**, as shown in Figure 4.29.

Figure 4.29 The Domain Controller Type Dialog Box Used for a New DC in an Existing Domain



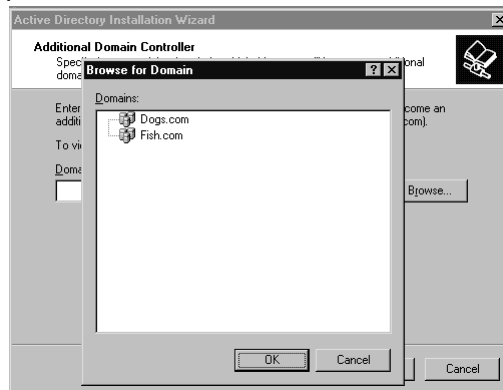
8. In the **Copying Domain Information** window, click **Over the network** | **Next**, as shown in Figure 4.30.

Figure 4.30 The Copying Domain Information Dialog Box Used for the Across-the-Network Method



9. In the **Network Credentials** window, type in the username, password, and domain name of an Enterprise Administrator in the forest-root domain, or a Domain Admin in the parent domain, and click **Next**.
10. In the **Additional Domain Controller** window, type in or browse to the top-level domain name where you are adding the new DC, and click **Next**, as shown in Figure 4.31.

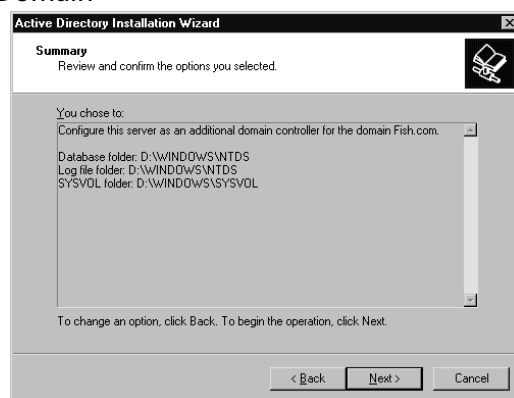
Figure 4.31 The Additional Domain Controller Dialog Box Browse Option



11. In the **Database and Log Folders** window, type or browse to the location where you want the database and log folders. Click **Next**.

12. In the **Shared System Volume** window, type or browse to the location where you want the SYSVOL folder. Click **Next**.
13. In the **Directory Services Restore Mode Administrator Password** window, input and confirm the password for the **Directory Services Restore Mode**. Click **Next**.
14. Read the **Summary** window. Click **Next**, as shown in Figure 4.32. The installation will continue for several minutes.

Figure 4.32 The Summary Dialog Box Used for an Additional DC in an Existing Domain



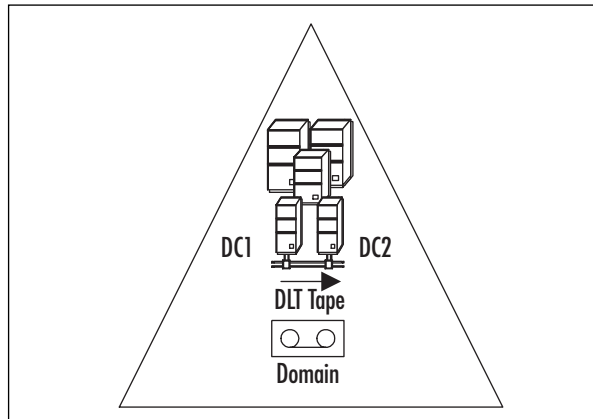
15. Restart your new DC.
16. Verify that the installation was successful. Open a command prompt and enter the **Net Share** command. It should report the existence of the Netlogon and SYSVOL shares. To verify that the DNS service locator records for the new DC were successfully created, follow these steps:
 1. Click **Start | Administrative Tools | DNS** to start the DNS administrator console.
 2. Expand the server name.
 3. Expand **Forward Lookup Zones**.
 4. Expand the domain.
 5. Verify that the `_msdcs`, `_sites`, `_tcp`, and `_udp` folders are present and contain records for your new DC. These service location records are crucial to the operation of the DC. See Table 4.7 for a more detailed description of the required records, and Figure 4.16 for a view of the DNS administrator tool used to view them.

Using the New System State Backup Method

Windows 2000 only offered two choices when deploying DCs and GC servers for remote sites, and neither choice was ideal for many companies. The first choice was to build the server at the home office where it could replicate over the LAN, and ship it to the remote location. This worked, as long as you got the new server online within the 60-day tombstone lifetime. If you didn't, the DC or GC could reanimate previously deleted Active Directory objects, including user accounts.

The second choice was to promote the server at the remote location and hope that replication would finish before users needed to use your WAN bandwidth for something else—like logging on to the domain. Replication can take days over a slow link, depending on the size of your directory and the available bandwidth. The tombstone problem still exists, but it is easier to ship new backup media from the main office than it is to ship an entire server.

Figure 4.33 Using the New System State Backup Method



This final DC installation procedure covers the new method of installing the Active Directory database on your new DC from backups, as illustrated in Figure 4.33. You should use a healthy Windows Server 2003 DC as the source of the system state, and DNS should be working before you begin. Exercise 4.09 is an advanced procedure, and assumes certain skills such as installing Windows Server 2003 as a member server, the use of Windows Backup, and general Windows administrative abilities. You should also test this procedure in a lab environment before trying it on an operational network. In addition, Exercise 4.09 will show you how to use an *answer file* to automate the promotion process, making this the optimal procedure for unattended installations. Figure 4.34 shows a sample answer file. The `/ADV` switch with `dcpromo` is only necessary for promoting from a backup file.

EXERCISE 4.09

CREATING A NEW DOMAIN CONTROLLER IN AN EXISTING DOMAIN USING THE NEW SYSTEM STATE BACKUP METHOD

Steps 1 through 3 walk you through taking the snapshot.

1. Log in as a local Administrator on the healthy DC.
2. Create a directory called C:\Backup. If the folder already exists, remove any files that it contains.
3. Using **Windows Backup**, save the *system state*. It is a good practice to name the file after your source DC, giving it a .bkf extension.

You now must transport the file. Use the backup media of your choice, ensuring your ability to perform the restore at the other end. Remember that the backup file can be many GBs in size. If you choose to use the network to transport the file, you can perform the restore and the copy at the same time using the following steps. There are various ways to accomplish this. If you choose to use a third-party backup program to transport the file on physical media such as DLT tape, CD, or DVD, you will still need to use **Windows Backup** at the other end to extract the data from the backup file. Adjust the procedure to your preferences.

4. Log on as a local Administrator on the member server that you want to promote, and create a shared folder called C:\Restore. It might be on your LAN or across a WAN at this point, so you might need a helping hand at the other end.
5. Back at the DC, map a drive to the shared folder created previously if you choose to copy the file over the network.
6. You have two options, depending on your choice of transport. If you are copying the file across the network, use the **Restore Wizard** within **Windows Backup** from the existing DC to restore the domaincontrollername.bkf file to the shared folder in the member server. If you have created *physical* media for transport, use the **Restore Wizard** directly on the member sever using the local physical media.

NOTE



Within **Windows Backup**, select the **System State** as the file to restore. Most importantly, select **Advanced Options** and specify the mapped drive as the destination if you are restoring from the source machine. If you don't, it will assume the file's original location and you will revert your DC to the state it was in at the time of the backup.

7. Create a file on the member server containing the following settings. For the exercise, we call this file DCUnattend.txt. Examine the options in Figure 4.34. They allow for unattended Active Directory installations in other configurations such as directly across the network from an established DC. Remember to rename the member server *before* promoting it, or you will be faced with the opportunity to perform a *domain controller rename* procedure, which is another new feature of Windows Server 2003.

Figure 4.34 Sample DCUnattend.txt File

```
[Unattended]
Unattendmode=fullunattended
[DCINSTALL]
UserName=(domain or Enterprise admin account)
Password=(password)
UserDomain=(domain of the user account)
DatabasePath=c:\Windows\ntds
LogPath=c:\Windows\ntds
SYSVOLPath=c:\Windows\sysvol
SafeModeAdminPassword
CriticalReplicationOnly
SiteName=
ReplicaOrNewDomain=Replica
ReplicaDomainDNSName=(domain name, not including any server name)
ReplicationSourceDC=
ReplicateFromMedia=yes
ReplicationSourcePath=c:\NTDSrestore
RebootOnSuccess=yes
```



NOTE

Coordinate with your computer security department if necessary, since the username and password of an administrator account will be hard coded into the answer file. You might have local regulations prohibiting this. Microsoft has provided a risk mitigation method of automatically erasing the password from the file as soon as it is used. If you want to run the answer file again, you must edit the file and re-enter the password. One way around this is to delegate permission for this operation to a regular domain user account and use that account within the answer file, although other security issues have to be considered in that case.

8. Open a command prompt and type the following command: **Dcpromo /adv /answer:C:\DCUnattend.txt**. After it is complete, the system will reboot. If dcpromo stops and asks for information, then some information was missing from the answer file.
 9. Verify that the installation was successful. Open a command prompt and enter the **Net Share** command. It should report the existence of the Netlogon and SYSVOL shares. To verify that the DNS service locator records for the new DC were successfully created, follow these steps:
 1. Click **Start | Administrative Tools | DNS** to start the DNS administrator console.
 2. Expand the server name.
 3. Expand Forward Lookup Zones.
 4. Expand the domain.
 5. Verify that the **_msdcs**, **_sites**, **_tcp**, and **_udp** folders are present and contain records for your new DC. These service location records are crucial to the operation of the DC. See Table 4.7 for a more detailed description.
-

Assigning and Transferring Master Roles

In Windows NT 4.0, the domain had only one authoritative source for domain-related information, the PDC. With the implementation of Active Directory came the multi-master model, where objects and their properties can be modified on any DC and become authoritative through replication conflict resolution measures. This scalability effort came with a price in complexity, and Active Directory FSMO roles were designed. The problem with the multi-master architecture is that some domain and enterprisewide operations are not well suited for it. The best design placed those functions on a single DC within the domain or forest.

The advantage of a single-master model is that conflicts cannot be introduced while the Operations Master is offline. The alternative involves resolving conflicts later, with possibly negative results. The disadvantage is that all Operations Masters must be available at all times to support all dependant activities within the domain or forest. The Active Directory supports five operational master roles: the Schema Master, Domain Master, RID Master, PDC Emulator, and the Infrastructure Master. Two of these operate at the forest level only, the Schema Master and the Domain Naming Master. Conversely, the RID Master, PDC Emulator, and Infrastructure Master operate at the domain level. To examine this role relationship between master roles and the required authorization for administering them in the forest and domains, refer to Table 4.5. Additionally, since it is often important for you to be

aware of mixing certain operational master roles with the GC function, use Exercise 4.21 to view the GC status of any DC. Most importantly, ensure that your Infrastructure Master never coexists with a GC.

- **Schema Master:** To update the schema of a forest, you must have access to the Schema Master DC, which controls all schema updates and modifications. There can be only one Schema Master in the forest.
- **Domain Naming Master:** The Domain Naming Master DC controls the addition or removal of domains in the forest as well as adding and removing any cross-references to domains in external LDAP directories. There can be only one Domain Naming Master in the forest.
- **Infrastructure Master:** The Infrastructure Master is responsible for updating references from objects in the local domain to objects in other domains. There can be only one Infrastructure Master DC in each domain.
- **Relative ID (RID) Master:** The RID Master processes RID pool requests from all DCs in the local domain. These *relative* identifiers are the unique part of the SID, which is a security identifier used to uniquely identify objects and group memberships. There can be only one RID Master DC in each domain.
- **PDC:** The PDC Emulator is a DC that advertises itself as the PDC to workstations, member servers, and BDCs running Windows NT. It is also the Domain Master Browser, and handles Active Directory password collisions, or discrepancies. There can be only one PDC Emulator in each domain.

Table 4.5 Valid Authorization Levels for Viewing, Transferring, and Seizing Operations Master Roles

Role	Task	Domain Administrator on the Local Domain	Domain Administrator on the Forest-Root Domain	Enterprise Administrator
Schema Master	Viewing, transferring, or seizing		X (Plus <i>Schema Admins</i> membership)	X
Domain Naming Master	Viewing, transferring, or seizing		X	X
Infrastructure Master	Viewing, transferring, or seizing	X		X
RID Master	Viewing, transferring, or seizing	X		X
PDC Emulator	Viewing, transferring, or seizing	X		X

The forest level, therefore, has five roles—one of each. Each domain added after the forest root domain has three additional masters. With that information, we can determine the number of operations master servers required in a given forest with the following formula:

$$((\text{Number of domains} * 3) + 2)$$

Given the formula, we can determine that the forest depicted in Figure 4.22, with three domains, needs a maximum of 11 server platforms to support the 11 FSMO roles ($3*3=9$, and $9+2=11$), unless you assign multiple roles to a single DC. Often, small domains, empty root domains, or best practices will make combining several of these roles onto a single DC desirable. In the example shown in Figure 4.22, the following roles exist:

- One Schema Master in Dogs.com.
- One Domain Naming Master in Dogs.com
- Three PDC Emulators (one each in Dogs.com, Fish.com, and Cat.fish.com)
- Three RID Masters (one each in Dogs.com, Fish.com, and Cat.fish.com)
- Three Infrastructure Masters (one each in Dogs.com, Fish.com, and Cat.fish.com)

The first DC that you install in the forest root will automatically host all five roles. The first DC that you install in any additional domains will automatically host the three roles of PDC Emulator, RID Master, and Infrastructure Master.



EXAM WARNING

Remember that in a Windows 2000 mixed-mode domain or forest, or in a Windows Server 2003 interim-level domain or forest, Windows NT 4.0 DCs *cannot* host FSMO roles.

You can use the `ntdsutil.exe` command-line utility to transfer FSMO roles, or you can use an MMC snap-in tool. Depending on which role you want to transfer, you can use one of the following three MMC snap-in tools:

- Active Directory Schema snap-in (Schema Master role)
- Active Directory Domains and Trusts snap-in (Domain Naming Master role)
- Active Directory Users and Computers snap-in (RID Master, Infrastructure Master, and PDC Emulator roles)

To seize a role, you must use the `ntdsutil` utility. If a computer cannot be contacted due to a hardware malfunction or long-term network failure, the role must be seized.



NOTE

The Microsoft recommendation is that you perform any seizing procedure *from* the DC that is taking the FSMO roles. All five roles should be in the forest for full functionality. Seize a role only when the other DC that hosted it is not returning to the domain; otherwise, fix the broken DC and transfer the roles instead of seizing them. If the role is functional, `ntdsutil.exe` defaults back to a transfer, aborting the seizure attempt.

Locating, Transferring, and Seizing the Schema Master Role

The DC that hosts the Schema Master role controls each update or modification to the schema. You must have access to the Schema Master to update the schema of a forest.



NOTE

You must be a member of the Schema Admins group to perform this operation. The Enterprise Administrator account is automatically a member of this group.

Refer to Exercise 4.10 for instructions on how to identify the DC that is performing the Schema Master operation role for your forest using the command line or the GUI. Refer to Exercise 4.11 for instructions on how to transfer the Schema Master operations role for your forest to a different DC, and Exercise 4.08 for steps to seize the role to another DC in case of a failure.

Temporary loss of the Schema Master is not noticeable to domain users. Enterprise and domain administrators will not notice the loss either, unless they are trying to install an application that modifies the schema during installation or trying to modify the schema themselves. You should seize the schema FSMO role to the standby operations master only if your old Schema master will be down permanently.

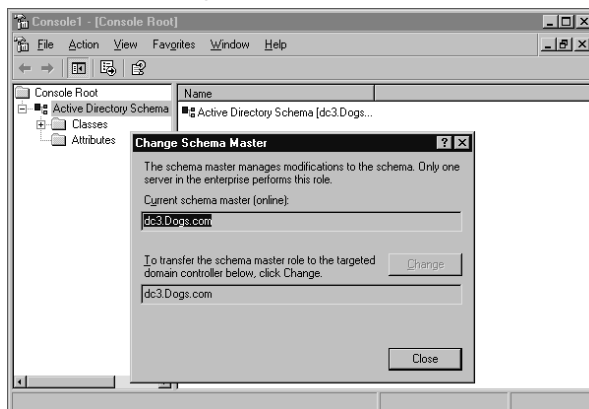
EXERCISE 4.10

LOCATING THE SCHEMA OPERATIONS MASTER

1. Log on as an Enterprise Administrator in the forest you are checking.
2. Click **Start | Run**.
3. Type `regsvr32 schmmgmt.dll` in the **Open** box, and click **OK**. This registers the *Schmmgmt.dll*.
4. Click **OK** in the dialog box showing that the operation succeeded.
5. Click **Start | Run**, type `mmc`, and then click **OK**.

6. On the menu bar, click **File | Add/Remove Snap-in**, click **Add**, double-click **Active Directory Schema**, click **Close**, and then click **OK**.
7. Expand and then right-click **Active Directory Schema** in the top left pane, and then select **Operations Masters** to view the server holding the Schema Master role as shown in Figure 4.35.

Figure 4.35 Locating the Schema Operations Master



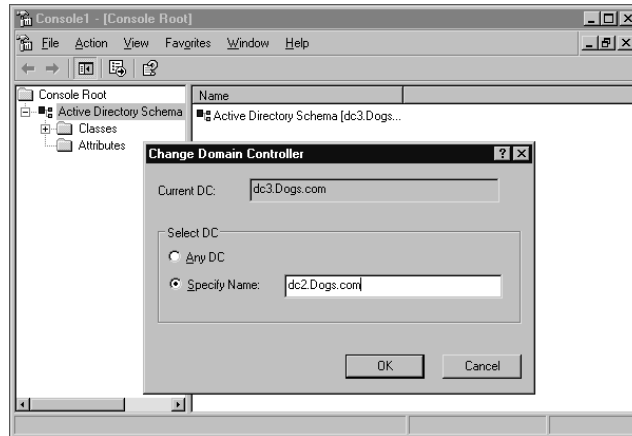
EXERCISE 4.11

TRANSFERRING THE SCHEMA OPERATIONS MASTER ROLE

1. Log on as an Enterprise Administrator in the forest where you want to transfer the Schema Master role.
2. Click **Start | Run**.
3. Type `regsvr32 schmmgmt.dll` in the **Open** box, and then click **OK**. This registers the *Schmmgmt.dll*.
4. Click **OK** in the dialog box showing that the operation succeeded.
5. Click **Start | Run**, type `mmc`, and then click **OK**.
6. On the menu bar, click **File | Add/Remove Snap-in**, click **Add**, double-click **Active Directory Schema**, click **Close**, and then click **OK**.
7. Right-click **Active Directory Schema** in the top left pane, and then click **Change Domain Controller**.

8. Click **Specify Name** as shown in Figure 4.36, type the name of the DC that will be the new role holder, and then click **OK**.

Figure 4.36 Transferring the Schema Operations Master Role



9. Right-click **Active Directory Schema** again, and then click **Operations Master**.
10. Click **Change**.
11. Click **OK** to confirm that you want to transfer the role, and then click **Close**.

Locating, Transferring, and Seizing the Domain Naming Master Role

The Domain Naming Master DC controls the addition or removal of domains in the forest, AND adding and removing any cross-references to domains in external LDAP directories. There can be only one Domain Naming Master in the forest.

Refer to Exercise 4.12 for instructions on how to identify the DC that is performing the Domain Naming Master operation role for your forest. Refer to Exercise 4.13 for instructions on how to transfer the Domain Naming Master operations role for your forest to a different DC, and Exercise 4.16 for steps to seize the role to another DC in case of a failure.



EXAM WARNING

Remember that in a Windows 2000 or Windows 2000 functional-level forest, the Domain Naming Master must also be a GC server. After the upgrade to the Windows Server 2003 forest functional level, that restriction is lifted and you can then separate the two functions.

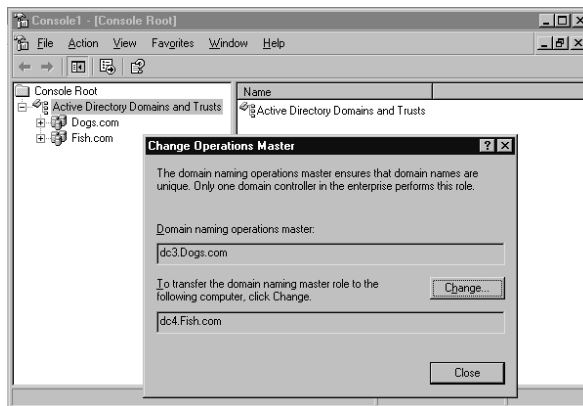
Windows Server 2003 Domain Naming Masters are no longer required to host the GC.

EXERCISE 4.12

LOCATING THE DOMAIN NAMING OPERATIONS MASTER

1. Log on as an Enterprise Administrator in the forest you are checking.
2. Click **Start | Run**, type: **mmc**, and then click **OK**.
3. On the menu bar, click **File | Add/Remove Snap-in**, click **Add**, double-click **Active Directory Domains and Trusts**, click **Close**, and then click **OK**.
4. Right-click **Active Directory Domains and Trusts** in the top left pane, and then click **Operations Masters** to view the server holding the domain naming master role as shown in Figure 4.37.

Figure 4.37 Locating the Domain Naming Operations Master



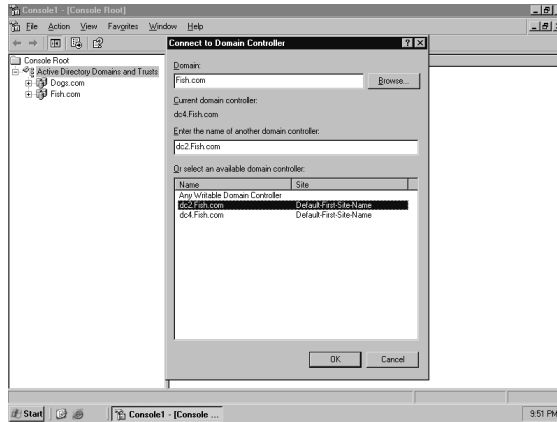
EXERCISE 4.13

TRANSFERRING THE DOMAIN NAMING MASTER ROLE

1. Click **Start | Administrative Tools | Active Directory Domains and Trusts**.

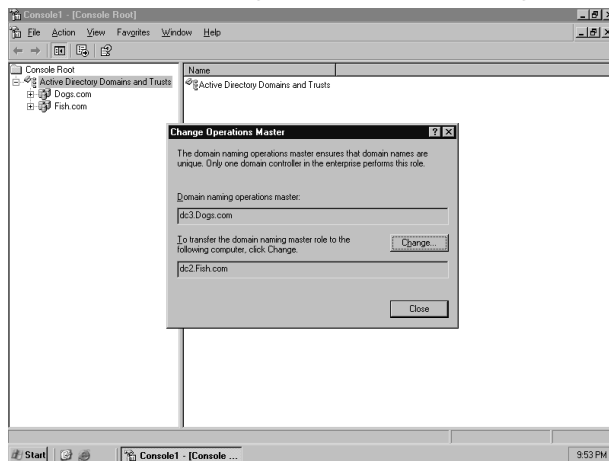
- Right-click **Active Directory Domains and Trusts**, and click **Connect to Domain Controller**, unless you are already on the DC to which you are transferring to the role. In the **Enter the name of another domain controller** window, type the name of the DC that will be the new role holder, and then click **OK**. Optionally, in the **Or, select an available domain controller** list, click the DC that will be the new role holder, and click **OK**. See Figure 4.38.

Figure 4.38 Connecting to a New DC



- In the console tree, right-click **Active Directory Domains and Trusts**, and then select **Operations Master** as shown in Figure 4.39.

Figure 4.39 Transferring the Domain Naming Master Role



- Click **Change**.

5. Click **OK** for confirmation, and click **Close**.
-

Locating, Transferring, and Seizing the Infrastructure, RID, and PDC Operations Master Roles

The Infrastructure Master is responsible for updating references from objects in the local domain to objects in other domains. There can be only one Infrastructure Master DC in each domain. The RID Master processes Relative ID (RID) pool requests from all DCs in the local domain. There can be only one RID Master DC in each domain. The PDC Emulator is a DC that advertises itself as the PDC to workstations, member servers, and BDCs running Windows NT. It is also the Domain Master Browser, and handles Active Directory password collisions, or discrepancies. There can be only one PDC Emulator in each domain.

Refer to Exercise 4.14 for instructions on how to identify the DCs that are performing the FSMO roles for your forest using the Active Directory Users and Computers GUI interface. Refer to Exercise 4.15 for instructions on how to transfer the Infrastructure, RID, and PDC Master operations roles for your forest to different DCs, and Exercise 4.16 for instructions on how to seize the roles to other DCs in case of a failure.

EXERCISE 4.14

LOCATING THE INFRASTRUCTURE, RID, AND PDC OPERATIONS MASTERS

1. Log on as an Enterprise Administrator in the forest you are checking.
 2. Click **Start | Run**, type **dsa.msc**, and click **OK**. This is an alternate method for opening the **Active Directory Users and Computers** administrative tool.
 3. Right-click the selected Domain Object in the top left pane, and then click **Operations Masters**.
 4. Click the **Infrastructure** tab to view the server holding the Infrastructure Master role.
 5. Click the **RID** tab to view the server holding the RID Master role.
 6. Click the **PDC** tab to view the server holding the PDC Master role.
-

EXERCISE 4.15

TRANSFERRING THE INFRASTRUCTURE, RID, AND PDC MASTER ROLES

1. Click **Start | Administrative Tools | Active Directory Users and Computers**.
2. Right-click **Active Directory Users and Computers**, and click **Connect to Domain Controller** *unless you are already on the DC you are transferring to*. In the **Enter the name of another domain controller** window, type the name of the DC that will be the new role holder, and then click **OK**; or in the **Or, select an available domain controller** list, click the DC that will be the new role holder, and click **OK**.
3. In the console tree, right-click **Active Directory Users and Computers**, and click **All Tasks | Operations Master**.
4. Take the appropriate action below for the role you want to transfer.
5. Click the **Infrastructure** tab, and click **Change**.
6. Click the **RID** tab, and click **Change** as shown in the example in Figure 4.40.
7. Click the **PDC** tab, and click **Change**.

Figure 4.40 Transferring the RID Master Role



8. Click **OK** for confirmation, and click **Close**.

EXERCISE 4.16

SEIZING THE FSMO MASTER ROLES

1. Log on to any working DC.
2. Click **Start | Run**, type **ntdsutil** in the Open box, and then click **OK**.
3. Type **roles**, and press **Enter**.
4. In **ntdsutil**, type **?** at any prompt to see a list of available commands, and press **Enter**.
5. Type **connections**, and press **Enter**.
6. Type **connect to server servername**, where *servername* is the name of the server that will receive the role, and press **Enter**.
7. At the server connections: prompt, type **q**, and press **Enter**.
8. Type the appropriate seizing command as shown next. See the example in Figure 4.41. If the FSMO role is available, *ntdsutil.exe* will perform a transfer instead. Respond to the Role Seizure Confirmation Dialog box as shown in Figure 4.42.

```

seize Schema master
seize domain naming master
seize Infrastructure master
seize RID master
seize PDC

```

Figure 4.41 Seizing the PDC Master Role

```

D:\WINDOWS\system32\ntdsutil.exe: roles
fsmo maintenance: connections
server connections: connect to server DC4
Binding to DC4 ...
Connected to DC4 using credentials of locally logged on user.
server connections: q
fsmo maintenance: seize PDC
Attempting safe transfer of PDC FSMO before seizure.
FSMO transferred successfully - seizure not required.
Server "DC4" knows about 5 roles
Schema - CN=NTDS Settings,CN=DC3,CN=Servers,CN=Default-First-Site-
Name,CN=Sites,
CN=Configuration,DC=Dogs,DC=com

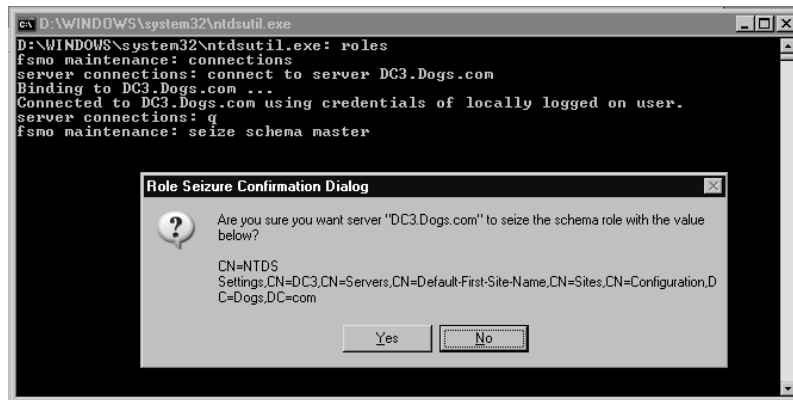
```

```

Domain - CN=NTDS Settings,CN=DC3,CN=Servers,CN=Default-First-Site-
        Name,CN=Sites,
        CN=Configuration,DC=Dogs,DC=com
PDC - CN=NTDS Settings,CN=DC4,CN=Servers,CN=Default-First-Site-
        Name,CN=Sites,CN=
        Configuration,DC=Dogs,DC=com
RID - CN=NTDS Settings,CN=DC4,CN=Servers,CN=Default-First-Site-
        Name,CN=Sites,CN=
        Configuration,DC=Dogs,DC=com
Infrastructure - CN=NTDS Settings,CN=DC4,CN=Servers,CN=Default-First-Site-
        Name,C
        N=Sites,CN=Configuration,DC=Dogs,DC=com
fsmo maintenance:q

```

Figure 4.42 Seizing the Schema Operations Master Role



9. After you seize the role, type **q**, and then press **Enter** repeatedly until you quit the Ntdsutil tool.

Placing the FSMO Roles

It is a good idea to place the RID and PDC Emulator roles on the same DC. Down-level clients and applications target the PDC, making it a large consumer of RIDs. Good communication between these two roles is important. If performance demands it, place the RID and PDC Emulator roles on separate DCs, but make sure they stay in the same site and that they are direct replication partners with each other.

As previously stated, you should place the Infrastructure Master on a non-GC server to maintain proper replication. Additionally, ensure that the Infrastructure Master has a direct connection object to a GC server somewhere in the forest, preferably in the same site. There are two exceptions to this rule:

- **Single domain forest** If your forest contains only one Active Directory domain, then there can be no phantoms. The Infrastructure Master has no functionality in a single domain forest. In that case, you can place the Infrastructure Master on any DC.
- **Multidomain forest where every DC holds the GC** Again, there can be no phantoms if every DC in the domain hosts a GC. There is no work for the Infrastructure Master to perform. In that case, you can place the Infrastructure Master on any DC.

Considering the forest level, the Schema Master and Domain Naming Master roles are rarely used and should be tightly controlled. For that reason, you can place them on the same DC. Another Microsoft-recommended practice is to place the Domain Naming Master FSMO on a GC server. Taking all of these practices together, a Microsoft-recommended best-practice empty root domain design would consist of two DCs with the following FSMO/GC placement:

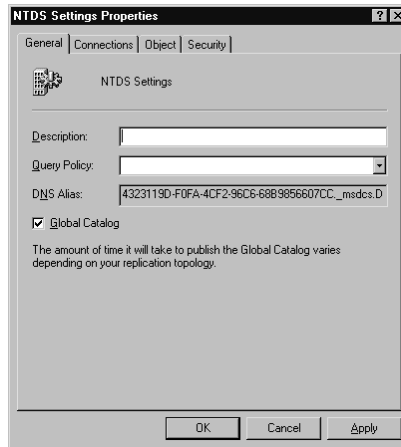
- **DC 1:**
 - Schema Master
 - Domain Naming Master
 - GC
- **DC 2:**
 - RID Master
 - PDC Emulator
 - Infrastructure Master

This preferred design remains valid until performance degradation forces you to separate the roles. Consider upgrading the hardware instead, or adding additional GCs, since the recommended configuration is the most efficient. For extremely large forests, install additional DCs and separate roles as needed. For these reasons and more, you need to be able to locate and assess your GC placement in relation to your FSMO roles. Here is how you find GCs:

1. Log on to any working DC.
2. Click **Start | Programs | Administrative Tools | Active Directory Sites and Services**.
3. Double-click **Sites** in the left console pane, and browse to the appropriate site, or click Default-first-site-name if no other sites are available.

4. Expand the **Servers** folder, and click the name of the DC that you want to check.
5. In the DC's folder, double-click **NTDS Settings**.
6. Click **Action | Properties**.
7. On the **General** tab, locate the **Global Catalog** check box to see if it is selected as shown in Figure 4.43.

Figure 4.43 Locating the Global Catalog Function



EXAM
70-294

OBJECTIVE
1.3.3

Using Application Directory Partitions

The Active Directory contains several partitions for the storage of object data. These *directory partitions*, also called *naming contexts*, are contiguous Active Directory subtrees that are replicated across DCs. As a minimum, each DC contains a replica of three partitions: the schema partition, the configuration partition, and the domain partition in addition to any application directory partitions that you might choose to create. An instance of an application directory partition on another DC is called a *replica*.

The default security descriptor for objects in the application directory partition is defined by an attribute called the *security descriptor reference domain*. By default, this attribute is the parent domain of the application directory partition. If the partition is a child of another application directory partition, the default security descriptor reference domain is the security descriptor reference domain of its parent. If it has no parent, the forest root domain becomes the default security descriptor reference domain. This attribute can be modified as shown in Exercise 4.17.



EXAM WARNING

Some important points to remember about application directory partitions:

- *Ntdsutil.exe* is the tool that you use to administer application directory partitions.
 - Two common uses of application directory partitions are TAPI ([Telephony]) and DNS.
 - An application directory partition does not have a NetBIOS name.
 - GPOs cannot reference application directory partitions.
-

Exercise 4.17 covers the basic procedures related to the use of application directory partitions.

EXERCISE 4.17

ADMINISTERING APPLICATION DIRECTORY PARTITIONS

1. Log on as an Enterprise Administrator.
2. Click **Start | Run**, type **ntdsutil**, and click **OK**.
3. At the **ntdsutil** command prompt, type **domain management**.
4. At the **domain management** command prompt, type **connection**.
5. At the **connection** command prompt, type **connect to server servername**, where *servername* represents the DNS name of the DC where you want to create the application directory partition.
6. At the **connection** command prompt, type **quit**.
7. At the **domain management** command prompt, consult the following list of commands for the function you want to perform:
 - Create an application directory partition: use the command **create nc application_directory_partition domain_controller**
 - Delete an application directory partition: use the command **delete nc application_directory_partition**
 - Add an application directory partition replica: use the command **add nc replica application_directory_partition domain_controller**
 - Remove an application directory partition replica: use the command **remove nc replica application_directory_partition**
 - Display application directory partition information: use the command **list**

- Add an application directory partition replica: use the command **set nc reference domain application_directory_partition domain_controller**

In this context, *application_directory_partition* is the DN of the application directory partition that you want to operate on, and *domain_controller* is the DNS name of the DC where you want to perform the operation. If you are operating on the DC that you connected to in step 5, use "NULL" as the *domain_controller* parameter.

8. Enter **q** until `ntdsutil` exits.



NOTE

If you remove the last replica of an application directory partition, you might permanently lose the data in that partition. Before you delete the last copy of a particular partition, make sure it is no longer needed.

Establishing Trust Relationships

As the name implies, trusts are all about sharing information. For security purposes, you should carefully consider your reasons before creating a new trust relationship, as well as knowing which type of trust to implement. In Active Directory, a *shortcut* trust doesn't add more trust; rather, it can make the trusts you already have more efficient. *External* trusts are a concept left over from Windows NT, but are still necessary for sharing resources with a Windows NT domain or any other Windows domain outside your forest. A *realm* trust allows cross-platform interoperability with non-Windows Kerberos V5 (version 5) realms, such as those commonly used with UNIX systems. Finally, you should consider the new Windows Server 2003 *forest* trust as discussed earlier in the chapter. As you can see, trusts are varied in properties and purposes. The most important concepts to understand about trusts before you create them are *direction* and *transitivity*. Always be aware of the extent of any internal access that you grant to external users.


Direction and Transitivity

As the name implies, trusts are predetermined avenues of access to forest resources. It is like giving someone a key to your house and hoping that he or she won't misuse your trust. DCs do the authenticating, but not all DCs necessarily trust each other. That's where you come in, setting the relationships between domains that govern the flow of information.

Two primary attributes of trusts are direction and transitivity. The *direction of trust* flows from the trusting domain to the trusted domain as shown by the arrow in Figure 4.44.

Cats.com *trusts* Dogs.com. The *direction of access* is always in the opposite direction; Dogs.com *accesses* resources in Cats.com. This is a one-way trust. Likewise, Dogs.com trusts Fish.com, but does not trust Cats.com. Two one-way trusts can combine to simulate a single two-way trust.

Figure 4.44 The Nontransitive Trust



The second attribute of the trust is *transitivity*, or a measure of how far the trust extends. A nontransitive trust has limits. The trusted domain, *and only the trusted domain*, can access resources through the trust to the trusting domain. As shown in Figure 4.44, if the Dogs.com domain has trusts to other domains such as Fish.com, those other domains are barred from access to Cats.com unless they have a nontransitive trust of their own. The absence of the third leg of the trust breaks the circle of access. This is the behavior of all trusts in Windows NT.

Conversely, transitive trusts, like the ones shown in Figure 4.45, are the skeleton keys of access. Anyone on the trusted side of the trust relationship can enter, including anyone trusted by the trusted domain. When a user or process requests access to a resource in another domain, a series of hand-offs occurs within the authentication process down the *trust path* as shown in Figure 4.45. When Cats.com trusts Dogs.com, they must trust all Dogs.com child domains equally at the level of the trust. There are two types of trusts in Figure 4.45, *parent and child* and *tree-root*. All trusts shown are bidirectional and transitive, as they are by default in Windows Server 2003. Calico.cats.com has a trust relationship with Yellow.labs.dogs.com because of the trust path that extends through all three intervening domains. If Calico.cats.com has no reason to trust Yellow.labs.dogs.com, then the cats must apply permissions to limit or block the access.



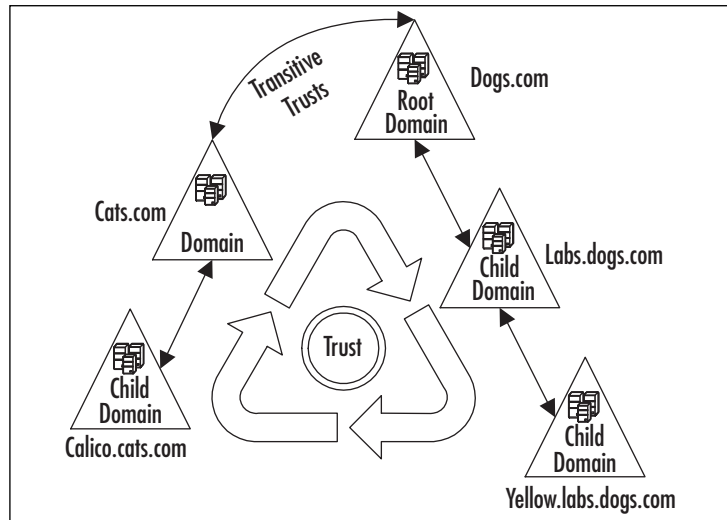
TEST DAY TIP

Remember that default Windows Server 2003 trust relationships are friendly. The default and most common trusts in Active Directory, which are *parent and child* and *tree-root* trusts, are both *bidirectional* and *transitive*, meaning that the trust

path extends throughout the entire forest. You can remember this type of transitive trust with the old saying, “Any friend of yours is a friend of mine.”

Other types of Windows Server 2003 trusts exist, such as *forest*, *shortcut*, *external*, and *realm*, each of which can be bidirectional or unidirectional and have different transitivity properties. One of the first things you should do when you sit down at the testing station is to write down the trusts and their properties on your scratch paper. Do this before starting the test so as not to waste valuable time.

Figure 4.45 The Transitive Trust



Types of Trusts

A *trust* is a logical authentication path between two domains. A *trust path* is the number of trusts that must be traversed between the source and destination of a resource request. Two trusts, *tree-root* and *parent and child*, are created by default when running the Active Directory Installation Wizard. The other four trusts—*shortcut*, *external*, *realm*, and *forest*—can be created as needed with the New Trust Wizard or the *Netdom.exe* command-line tool.

When creating those four trusts, you have the option of creating two one-way relationships, simulating bidirectional capabilities. As with any use of passwords, it is a security best practice to use long, random, and complex passwords in the establishment of trusts. The best option is to use the New Trust Wizard to create both sides simultaneously, in which case the wizard generates a strong password for you. Naturally, you must have the appropriate administrative credentials in both domains for this to work.

**NOTE**

See Chapter 5, “Working with Trusts and Organizational Units,” for more detailed information on the different types of trusts.

Restructuring the Forest and Renaming Domains

In Windows Server 2003, you can rename domains in an Active Directory forest after the forest structure is in place. This was not true in the Windows 2000 Server family. You build your Active Directory forest structure one domain at a time, and the resulting relationships are the result of the order in which you create them and the DNS names you assign.

Renaming domains allows you to change the forest structure. For example, you can raise a child domain to be a new tree-root domain, or lower a top-level domain to child status in another tree. In each case, you rename an existing domain to create a different forest structure. In cases where restructuring is not your goal, you can rename domains without affecting the trust relationships between domains. For example, you do not create a different domain-tree structure if you rename a root domain, although the names of all child domains below it are also changed.

This is a complex and sweeping modification to the namespace of one or more domains. During the domain rename procedure, you can change DNS and NetBIOS names, but the true identity of a domain lies in its domain GUID and its domain SID. Creating new DNS or NetBIOS names will leave those attributes unchanged.

**WARNING**

The renaming process will temporarily interrupt the functionality of the domain and its interaction with the forest, until the DCs are rebooted. Client workstations will not function properly until they are each rebooted twice. Due to the complexity of the operation, the risks of such a sweeping change, and the unavoidable domain and workstation service interruptions, do not consider domain renaming as a routine operation.

Domain Rename Limitations

Windows NT 4.0 had no *supported* method for domain renaming, other than a complete rebuild of the new domain. The best option that Windows 2000 offers is only the first half of the solution—you still have to create a new domain from scratch. Microsoft released a support tool called the Active Directory Object Manager that can migrate users, computers, and groups into the new empty domain structure. Windows Server 2003 now supports the full solution. However, even with the new restructuring capability, certain types of structural changes are not supported, and many forests cannot be renamed due to limitations of the pro-

cedure. These limitations include the presence of Exchange 2000. Other problematic issues arise, such as the failure of enterprise certificates with certain types of embedded pointers, and network saturation due to the replication of sweeping changes to the directory.

Domain Rename Limitations in a Windows 2000 Forest

Windows Server 2003's forest restructuring capabilities provide solutions to some of the problems that the Windows 2000 Server family did not address. In a Windows 2000 forest, renaming domains is not directly possible after the forest structure is in place. The only way to accomplish it is to move or recreate the domain contents. These constraints make domain name changes or forest restructuring prohibitive in Windows 2000.

- You cannot change the DNS name or the NetBIOS name of a domain. You can, however, achieve similar results by moving its contents into a new domain using the Active Directory Object Manager (MoveTree) in the Windows 2000 Support Tools.
- Using the Active Directory Object Manager method, you cannot move a domain within a forest in a single operation.
- Using the Active Directory Object Manager method, you cannot split a domain into two domains in a single operation.
- Using the Active Directory Object Manager method, you cannot merge two domains into a single domain in a single operation.

Domain Rename Limitations in a Windows Server 2003 Forest

Windows Server 2003 Standard, Enterprise, and Datacenter Editions provide tools that you can use to safely rename domains. Since domain renaming is at the core of forest restructuring, you can leverage this capability with very powerful results. When considering restructuring an existing Windows Server 2003 forest, be sure to consider the limitations of domain renaming. Adding, removing, merging, and splitting domains are operations outside the scope of the domain rename process.

- You cannot change which domain is the forest root domain, although you can still give it a new DNS or NetBIOS name.
- You cannot remove or add domains to the forest. The number of domains before and after the restructuring must remain the same (you can, of course, add new domains after the name change).
- You cannot move a domain name from one domain to another in a single operation.

The resulting forest, no matter how sweeping the DNS and NetBIOS changes are, must result in a well-formed forest. A *well-formed forest* has the following characteristics:

- All domains within the forest must form one or more DNS trees.
- The forest-root domain must be the root of one of these trees.
- A domain directory partition must not have an application directory partition as a parent.

Domain Rename Dependencies

Other conditions that must be eliminated, or prerequisites that must be met before you can attempt the domain rename procedure, include the following:

- Domain rename is not supported in a domain where Exchange 2000 is installed.
- All of your DCs must be running Windows Server 2003, and the Active Directory *forest* functional level must be raised to Windows Server 2003.
- The domain rename procedure requires Enterprise Administrator privileges.
- The control station for the domain rename operation must not be a DC. You must use a member server to perform the operation.
- All domain-based DFSroot servers must be running Windows 2000 with Service Pack 3 or higher.

The authentication levels needed in each step of the domain rename procedure are identified in Table 4.6.



WARNING

Rendom.exe versions prior to version 1.2 (such as the one on the Windows Server 2003 distribution CD) failed to detect Exchange 2000 and incorrectly permitted domain rename operations. The updated tool is available at www.microsoft.com/windowsserver2003/downloads/domainrename.mspx.

Table 4.6 Required Authorization Levels for Each Step of the Domain

Rename Procedure

STEP	Task	Tool(s)	Local Administrator on the Control Machine	Domain Administrator on the Local Domain	Domain Administrator on a Different Domain	Enterprise Admin- istrator
1	Back up all DCs	Backup, or third-party application		X (or Backup Operator)		
2	Set up the control station	Copy, and Windows Server 2003 Support tools setup	X			
3	Generate the current forest description	Copy, and <code>rendom /list</code>	X			X
4	Specify the new forest description	<code>rendom /showforest</code> , and Notepad or other plain-text editor	X			
5	Generate domain rename instructions	<code>rendom /upload</code>	X			X
6	Push domain rename instructions to All DCs and verify DNS readiness	<code>dsquery</code> , <code>repadmin</code>				X
7	Verify readiness of DCs	<code>rendom /prepare</code>	X			X
8	Execute domain rename instructions	<code>rendom /execute</code>	X			X
9	Unfreeze the forest configuration	<code>rendom /end</code>				X
10	Re-establish external trusts	Active Directory Domains and Trusts, or <code>netdom</code>		X	X	
11	Fix Distributed File System (DFS) topology	DFS (MMC snap-in), or <code>dfsutil</code>		X	X	
12	Fix GPOs and links	<code>gpfixup</code> , and <code>repadmin</code>				X

Domain Rename Conditions and Effects

The domain rename procedure is complex, requires a great deal of care in planning and execution, and should always be tested in a lab environment before performing it on an operational forest. The time required to go through a complete domain rename operation varies; the number of domains, DCs, and member computers is directly proportional to the level of effort required.



NOTE

There is a good reason for caution. Read this entire procedure before attempting any part of it, including the pre- and post-procedure steps. You might find limitations that preclude the procedure altogether on your network. Consult Microsoft documentation, read Technet articles, and search for patches, hotfixes, and service packs that can affect domain renaming and forest restructuring. Every attempt is made in this chapter to address all pertinent topics and concerns, but issues and conflicts continue to be exposed over time. Search Microsoft.com for new “Q” articles detailing conditions that might have an affect on this procedure. Most importantly, consider hiring a consultant who has recently and successfully performed a domain renaming operation.

Before undertaking a domain rename operation, you must fully understand the following conditions and effects. They are inherent in the process and must be dealt with or accommodated.

- Each DC requires individual attention. Some changes are not replicated throughout the Active Directory. This does not mean that every DC requires a physical visit. Headless management can greatly reduce the level of effort required, depending on the size and structure of the domain and the number of sites it contains.
- The entire forest will be out of service for a short period. Close coordination is required with remote sites, especially those in other time zones. During this time, DCs will perform directory database updates and reboot. As with other portions of the procedure, the time involved is proportional to the number of DCs affected.
- Any DC that is unreachable or fails to complete the rename process must be eliminated from the forest for you to declare the procedure complete.
- Each client workstation requires individual attention. After all DCs have updated and rebooted, each client running Windows 2000 or Windows XP must be rebooted two times to fully adapt to the renamed domain. Windows NT workstations must disjoin from the old domain name and rejoin the new domain name, a manual process that requires a reboot of its own.

- The DNS host names of your DCs are not changed automatically by the domain rename process. To make them reflect the new domain name, you must perform the *domain controller rename procedure* on each DC. Having the host name of a DC decoupled from its domain name does not affect forest service, but the discrepancy will be confusing until you change the names.
- The DNS suffix of client workstations and member servers will automatically update through the domain renaming process, but not all computers will match the DNS name of the domain immediately. As with most portions of this process, the period of time required is proportional to the number of hosts in the domain.

Domain Rename Preliminary Steps

Prerequisites for the domain rename operation are not trivial. The preparation phase will ensure that these are in place. Complete all of the preliminary steps in this section before beginning the rename procedure. If these prerequisites are not taken care of, the domain rename cannot be successfully performed.

Setting Windows Server 2003 Forest Functionality

The first step in preparing for a domain rename is to ensure that all DCs are running some edition of Windows Server 2003. This is a prerequisite to raising the forest to the Windows Server 2003 functional level, which is another preparatory step. See the section *Raising the Functional Level of a Domain and Forest* for additional information on functional levels.

Creating Shortcut Trust Relationships

Interaction between domains in your forest is based on the establishment of trusts among the domains. The Active Directory Installation Wizard creates most of these trusts automatically during the domain creation process. Through the manual creation of shortcut trusts, you can maintain that interaction after the domains are renamed. It is only necessary if the forest *structure* will change as result of the manipulation of the namespace. If you are renaming a domain in place without changing its relationship with other domains in the forest, then this step is not needed. Refer to Chapter 5 for the trust-creation procedures.

Pre-Creating a Parent-Child Trust Relationship

While repositioning domains, the necessary shortcut trust relationships must be created between the domain you want to reposition and its new parent domain. These pre-created trust relationships substitute for the required parent-child trust relationships that will be missing in the restructured forest.

For example, suppose you want to restructure the Zoo.net forest, shown in Figure 4.46, so that the Cat.fish.zoo.net domain becomes a child of the Zoo.net domain. You must create two one-way, transitive shortcut trust relationships between Cat.fish.zoo.net and Zoo.net before you can rename the child domain Cat.fish.zoo.net to the child domain

Catfish.zoo.net. This trust relationship pre-creates the two-way parent-child trust relationship required for the parent and child domains after the rename. Figure 4.46 shows the *before* structure, and Figure 4.47 shows the *after* structure, illustrating the needed shortcut trust relationships for the new structure.

Figure 4.46 Pre-Creating a Parent-Child Trust Relationship Before the Forest Restructure

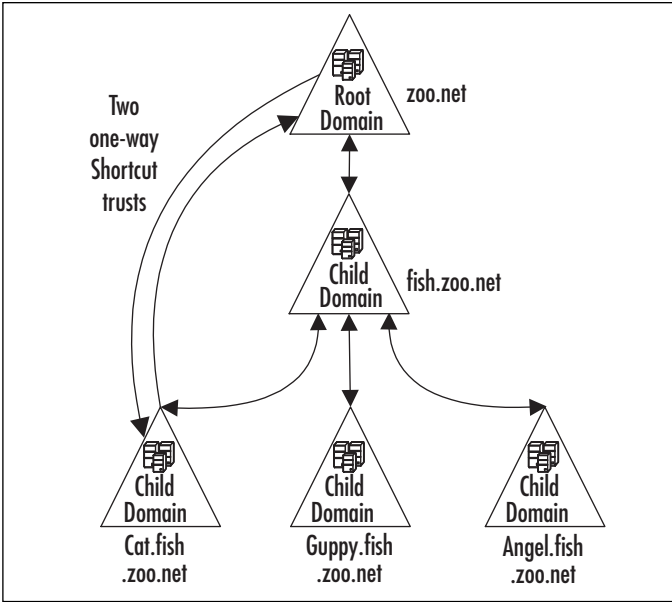
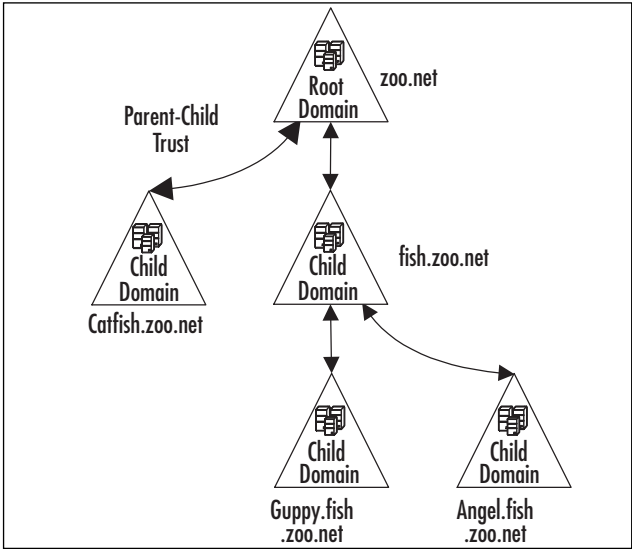


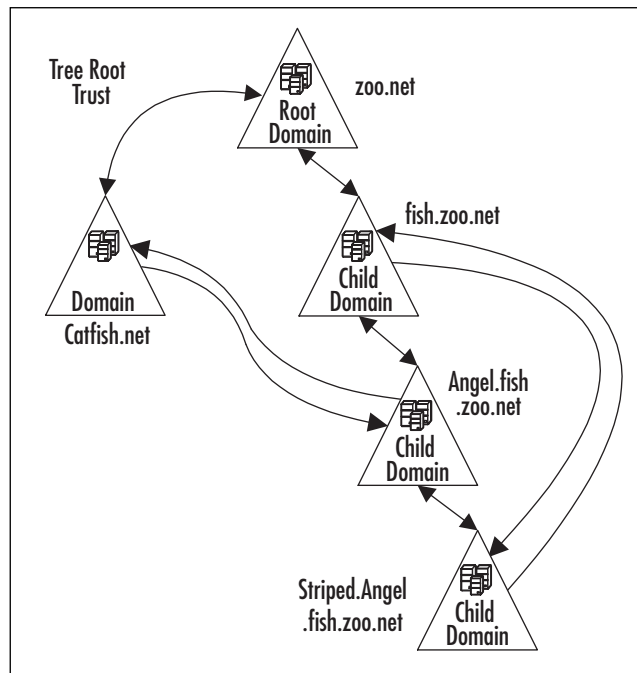
Figure 4.47 Parent and Child Trust After the Forest Restructure



Pre-Creating Multiple Parent-Child Trust Relationships

If you need to restructure a domain that is both a child domain and a parent domain, you will need to create shortcut trust relationships in two places. For example, suppose you want to restructure the Zoo.net forest, shown in Figure 4.48, so that the Striped.angel.fish.zoo.net domain becomes a direct child of Fish.zoo.net, and the Angel.fish.zoo.net domain becomes a child of Catfish.net. This restructure operation calls for four shortcut trusts that will become the two parent-child trust relationships for the new forest. Figure 4.48 shows the *before* structure, and Figure 4.49 shows the *after* structure, illustrating the needed shortcut trust relationships.

Figure 4.48 Pre-Creating Multiple Parent-Child Trust Relationships Before the Forest Restructure



Pre-Creating a Tree-Root Trust Relationship with the Forest Root Domain

When you restructure a domain to become a new tree root, you must pre-create two one-way, transitive trust relationships with the forest root domain. For example, suppose you have a three-level deep tree and you want to shorten it by creating a new tree. This will move the lowest domain to become a new tree-root domain. Figure 4.50 shows the two one-way shortcut trusts you create, and Figure 4.51 shows the tree-root trust relationship after the restructuring. Stripedangel.fish.zoo.net becomes the tree-root domain Angelfish.net.

Figure 4.49 Multiple Parent and Child Trusts After the Forest Restructure

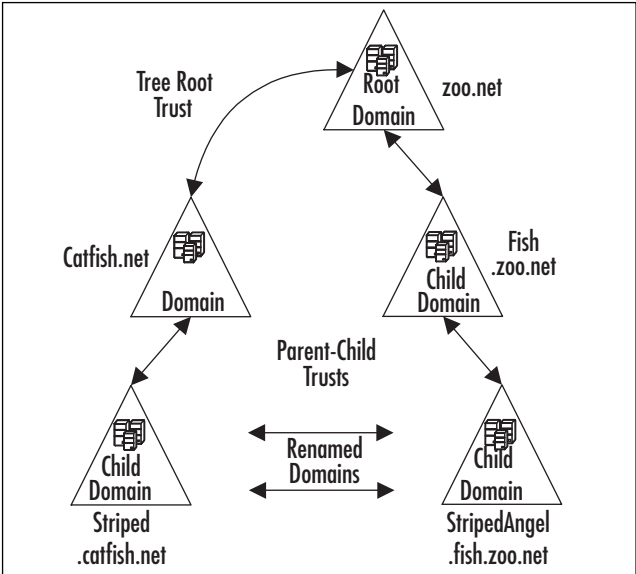


Figure 4.50 Pre-Creating a Tree-Root Trust Relationship Before the Forest Restructure

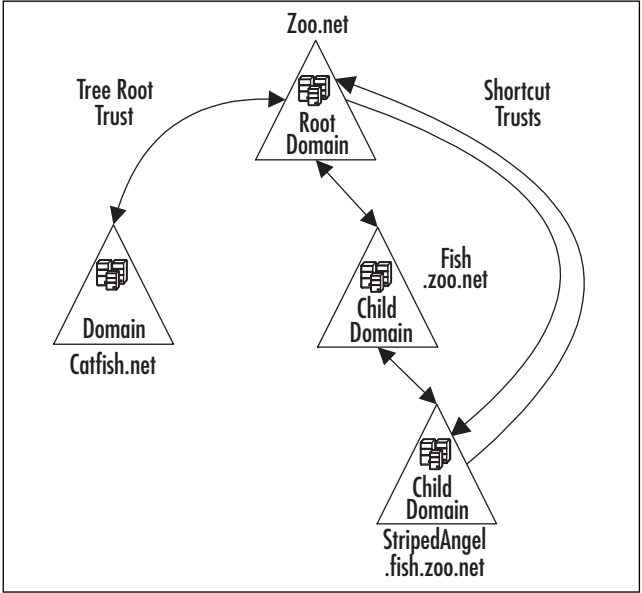



Figure 4.51 Tree-Root Trust Relationship After the Forest Restructure

Preparing DNS

Any time a client requires access to Active Directory, it activates an internal mechanism called the *DC locator* for locating DCs through DNS. It uses SRV records for this. If no SRV records are found in DNS, the access fails. To prevent this failure, before renaming an Active Directory domain you need to be sure that the appropriate zones exist for the forest and for each domain.

After you create the DNS zones for the new domain name, your DCs will populate each zone through dynamic update. This is one of the reasons for the reboot after the execution of the renaming script. Configure the zones to allow secure dynamic updates as a good security practice. Repeat the zone creation for each domain you plan to rename.

Everything needed to support your existing Active Directory domain must be recreated to support the domain after renaming. Usually, this is accomplished by mirroring your current DNS infrastructure. As an example, say you want to rename an existing domain called Labs.dog.com to Retrievers.dog.com. If the zone containing your current SRV resource records is called Labs.dog.com, you will need to create a new DNS zone called Retrievers.dog.com.

To analyze and prepare DNS zones for domain rename, first compile a list of DNS zones that you need to create. Second, create the forward lookup zones using the DNS tool and configure them to allow dynamic updates. The section *Configuring DNS Servers for Use with Active Directory* gives more detailed information.

What Happens to My Distributed File System When I Rename My Domain?

First, those of you who are not using DFS should think seriously about it. DFS allows you to redirect specific folders like *My Documents* out to a high-availability network location where each user's files can be backed up and protected. *Folder redirection* is a Group Policy extension that allows you to identify a connection between network servers or DFS roots and the local folders that you want to redirect.

What happens to DFS when you rename a domain all depends on how you have it configured. Think about it. If you use a domain-based DFS path like `\\domainName\DFSRoot`, then when the `domainName` goes away, what happens to the path?

It goes dead, and everyone's documents disappear, or become inaccessible. As far as the users know, all of their data is gone. Your telephone will ring by 5 a.m. the next day—guaranteed. What does it depend on, and how can you keep your telephone from ringing? If your Folder Redirection policy specifies the NetBIOS name of the domain in your domain-based DFS path, *and* you keep the NetBIOS name of your domain the same instead of changing it along with the DNS name, then you're okay.

What if you want to change your NetBIOS name *along with* your DNS name? You could push out a new group policy and move the files to another location. Temporarily, you could point your folder redirection to a stand-alone DFS path, or even to a simple server-based share. You should do that a couple of days before the rename just to be sure it works before shaking things up again—you'll be too busy renaming to worry about DFS at that point. Since `\\hostName\DFSRoot` stays rock solid through a domain rename, your documents should still be available the next morning. When things settle down, restore the user files back to your domain-based DFS root and push out the old DFS policy again. That isn't without risk, but it keeps things working.

What about home directories and roaming profiles? Same thing. Look at the pathname you specify in your policy to determine whether they'll break when you rename the domain. Make sure to fix those beforehand.

Configuring Member Computers for Host Name Changes

Because Active Directory is tightly integrated with DNS, member computers are designed to automatically change their primary DNS suffixes when the domain membership of the computer changes. If you rename the domain, this is treated like a membership change and the fully qualified DNS host name changes automatically to match. This is the default behavior, and you can check for it by following the steps in Exercise 4.18.

As an example, if you want to rename an existing domain called `Labs.dog.com` to `Retrievers.dog.com`, the full DNS host name of the member computers of this domain will also change from `host.Labs.dog.com` to `host.Retrievers.dog.com` *if* the default behavior is in effect.



NOTE

You should check to see if this default behavior has been changed in your domain, because your rename will fail if you are not using the default setting.

The full DNS name and therefore the primary DNS suffix of a member computer changes when the domain is renamed if both of the following conditions are true:

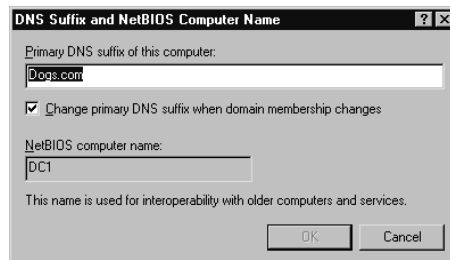
- The primary DNS suffix of the computer is configured to update when domain membership changes. See Exercise 4.18 for instructions on how to check this setting.
- The member computer has no group policy applied that specifies a primary DNS suffix. See Exercise 4.19 for instructions on how to check this setting.

EXERCISE 4.18

USING THE CONTROL PANEL TO CHECK FOR PRIMARY DNS SUFFIX CONFIGURATION

1. On a member computer, open the **System Control Panel**.
2. Click **Computer Name | Change**.
3. Click **More**, and verify if **Change primary domain suffix when domain membership changes** is selected (as shown in Figure 4.52). If it is, then the computer will automatically adjust to the new primary DNS suffix.
4. Click **OK** until all dialog boxes are closed.

Figure 4.52 The System Control Panel, General Tab, More Button



Determining Whether Group Policy Controls the Primary DNS Suffix for the Computer

There are a few ways to determine whether Group Policy controls the primary DNS suffix for the computer. Log on to a representative member computer and do one of the following:

- Open a command prompt and type **gpresult**. Look in the output to see if Primary DNS Suffix is listed under Applied Group Policy objects.
- Open **Active Directory Users and Computers**, right-click the computer object you want to check, and click **All Tasks | Resultant Set of Policy (Logging)**.
- Perform the steps in Exercise 4.19. If a value is present in step 4, then the primary DNS suffix group policy is applied to the computer.

EXERCISE 4.19

USING THE REGISTRY TO CHECK FOR PRIMARY DNS SUFFIX DOMAIN RENAME COMPUTER READINESS

1. Click **Start | Run**.
2. Type **regedit** and click **OK**.
3. Navigate to **HKEY_LOCAL_MACHINE\Software\Policies\Microsoft\System\DNSClient**.
4. If the **Primary DNS Suffix** key contains a value, then the computer will not automatically adjust to the new primary DNS suffix.
5. Navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters**.
6. Verify whether the value of **REG_RWORD SyncDomainWithMembership** is **0x1**. This value indicates that the primary DNS suffix changes when the domain membership changes. Any other value means that the computer will not automatically adjust to the new primary DNS suffix.

Because the replication effects of member computer names being updated is proportional to the number of computers within the renamed domain, large domains can generate a large amount of traffic. This replication “storm” is a problem for only the largest deployments. If you think that the resulting replication traffic might pose a problem to your infrastructure, then consult the section entitled *Avoiding Replication Effects of Domain Rename in Large Deployments* later in this chapter.

Preparing Certification Authorities

You can continue to support the management of enterprise certificates through a domain rename when the following requirements have been met beforehand:

- The Certificate Authority functionality (CA) is not installed on any DCs.
- All CAs should include both LDAP and HTTP URLs in their Authority Information Access (AIA) and Certificate Distribution Point (CDP) extensions.



NOTE

If the CA has issued any certificate with only one of these URL types, the certificate might not work. The steps covered in this chapter might not be sufficient for full management of your CAs after the domain is renamed, depending on the complexity of your domain configuration. Proceed with the domain renaming only if you have substantial expertise in handling Microsoft CAs.

If any of the following situations exist, CA management is not supported through the domain rename process:

- The CDP or AIA configuration includes only LDAP URLs. Certificates issued by the CA will no longer be valid because the old LDAP extensions will be wrong after the domain rename process. As a workaround, you will have to renew the existing CA hierarchy and all issued End Entity certificates.
- If you have established any interdomain trust relationships authenticated with cross-certificates, and they have name constraints, those names might not be valid after the domain rename operation. As a workaround, you will have to reissue cross-certificates containing the new name constraints.
- If RFC 822-style e-mail names (user@host.network) are a property of the user account, these e-mail names will be incorrect after the domain rename operation. Additionally, if any certificate template is configured to include RFC 822-type e-mail names within the certificate, those names will no longer be valid. Those Active Directory accounts should be updated prior to issuing any new certificates.

Finally, check to see if your certificate revocation lists (CRLs) and CA certificates are going to expire soon. If so, these should all be renewed, reissued, and replicated to all client machines prior to the rename process.

Avoiding Replication Effects of Domain Rename in Large Deployments

Thousands of computer names being changed at about the same time can cause a replication “storm.” If you are concerned about this replication traffic, this section is for you. If you can tolerate a significant period of network congestion or saturation, you can skip this

preparatory step. Note that versions of Windows before Windows XP do not fully support this approach.

The solution is to rename member computers in smaller batches to lessen the replication traffic problem. You must take steps to limit the number of computers that will be renamed following domain rename by reversing one of the conditions in *Configuring Member Computers for Host Name Changes* earlier in this chapter.

When a computer's name is changed, an update of the *dnsHostName* and *servicePrincipalName* attributes on its computer account in Active Directory is triggered. This happens during the reboot after the domain rename. In addition, it triggers an update of the host (A) and pointer (PTR) DNS resource records in the DNS database. The two events, if triggered by a large number of computers within a short period of time, might provoke replication activity that saturates the network. Using Group Policy, you should reconfigure the default behavior that changes the primary DNS suffix on member computers when a domain is renamed.

Group Policy triggers the same replication traffic as the domain rename procedure. For that reason, you must manage the Group Policy application in stages. This is accomplished by dividing computer objects among several OUs or sites in Active Directory. To temporarily separate some of the computers, you might want to create additional interim OUs. Generally speaking, only do this for workstations. Allow member servers to change automatically during the reboot, making them the last to change; otherwise, some services might be affected until the domain rename is complete. If you have a large number of servers, apply the DNS Suffix Search List policy to workstations first, and then to servers.

Before Applying Group Policy

The purpose of applying this group policy is to avoid replication and DNS update traffic caused by the automatic update of the primary DNS suffix on all member computers following a domain rename. Use Group Policy to revise the primary DNS suffix of all computers in stages to the new domain name before the procedure. That way, domain computers are manually updated and already have the correct primary DNS suffix at the time you perform the domain rename. After you apply the group policy, the DNS suffix of member computers will not match the DNS name of the domain for some period. To handle this problem, the first step is to configure the domain to accept the possible names that the DNS suffix can have. Make sure you do this before applying the policy.

The default primary DNS suffix is the name of the domain itself. To make a new one available to member computers, you must first make the DNS suffixes known at the domain level. You accomplish this by editing the *msDS-AllowedDNSSuffixes* attribute on the domain object to contain these additional DNS domain names. This multivalued attribute contains a list of valid DNS suffixes for member computers of the Active Directory domain. The *msDS-AllowedDNSSuffixes* attribute should not contain the same values in more than one domain.

At the same time, configure the DNS Suffix Search List group policy on all systems to contain the old primary DNS suffix, new primary DNS suffix, and any parent suffixes of

the old and new primary DNS suffixes. For example, if the old DNS suffix of a domain was `Cat.fish.zoo.com` and the new name will be `Catfish.naturepreserve.com`, the DNS Suffix Search List should contain the following suffixes:

- `Cat.fish.zoo.com`
- `Catfish.naturepreserve.com`
- `Fish.zoo.com`
- `Naturepreserve.com`

Note that the versions of Windows before Windows XP do not support the Group Policy setting for the DNS Suffix Search List. To perform this procedure, all DCs in the domain must be running Windows Server 2003, and a subdomain must exist in DNS for each new DNS suffix that you add. Use ADSI Edit, one of the Windows Server 2003 Support Tools, to modify the *msDS-AllowedDNSSuffixes* attribute value on the domain object in Active Directory as shown in Exercise 4.20. Do this for each domain whose name is going to change.

EXERCISE 4.20

USING ADSI EDIT TO ADD DNS SUFFIXES TO *MSDS-ALLOWEDDNSSUFFIXES*

1. Click **Start | Programs | Windows Server 2003 Support Tools | Tools | ADSI Edit**.
2. In the scope pane, right-click **ADSI Edit** and select **Connect to**.
3. Under **Computer**, click **Select** or type a domain or server name, and then click **OK**.
4. Double-click the *domain* directory partition for the domain you want to modify.
5. Right-click the *domain* container object, and select **Properties**.
6. In the **Attributes** box, on the **Attribute Editor** tab, double-click the **msDS-AllowedDNSSuffixes** attribute.
7. In the **Multi-valued String Editor** dialog box, in the **Value to add** field, type a DNS suffix and then click **Add**.
8. Repeat this process for all the DNS suffixes that you need for the domain, and click **OK**.
9. Click **OK** to close the **Properties** dialog box.
10. Repeat steps 2 through 9 for any additional domains that are being renamed.

Using Group Policy to Predefine the Primary DNS Suffix Prior to Domain Rename

To prepare for the application of Group Policy, you need to create groupings of member computers for incremental rollout. Perform the following steps for each domain to be renamed.

1. Estimate the largest number of computers that can be renamed in your environment without adverse affects. Microsoft's recommendation is to define groups of 1000 or less for a normal healthy LAN environment. Adjust this number for local conditions.
2. Define rollout groups of the chosen size.
3. Create a schedule, leaving sufficient time between applications of Group Policy to allow replication to occur. Make sure the updated *dnsHostName* and *servicePrincipalName* attributes on computer accounts and replication of the DNS records of the renamed computers are completed during each period before the next group begins.
4. Apply Group Policy to the rollout groups according to the schedule. Make sure all groups receive the policy before renaming the domain. Note that computers must be rebooted for the host name change to take effect.



NOTE

Do not apply this policy to DCs. They will be renamed later in a separate procedure.

5. After the domain rename procedure is complete, and after all computers have rebooted, disable the temporary DNS Suffix group policy.

Performing the Rename Procedure

This section presents step-by-step procedures for executing the domain rename operation in your forest. Be sure to review and complete the preliminary procedures before performing any steps in this section. There will be a short period of service interruption in your Active Directory forest during this procedure. For the most part, this occurs while all the DCs in the forest are automatically rebooting. The applicable section of the procedure tells you when to expect it. Except for this short interruption, the Active Directory service should continue to be available and function normally throughout the rest of the procedure.

A certain level of proficiency is expected from those performing this procedure. Do not attempt it unless you have considerable experience with Active Directory, CAs, domain and forest maintenance procedures and troubleshooting, and are comfortable with the administrative tools involved. Although not all steps require the same level of authority, you should

have access to the group memberships listed in Table 4.6 to perform the procedure as a whole. These include:

- Enterprise Admins group in the forest
- Domain Admins group in trusted domains
- Domain Admins group in trusting domains
- Local Administrators group on the control station

Each step should be completed before going on to the next. Many steps will fail or have unpredictable results if performed out of order. This especially applies to the *freezing* and *unfreezing* of the forest configuration. During this time, the forest must be *quiescent*. You must not make any changes to the Active Directory such as adding or removing domains, DCs, directory partitions, or trust relationships other than those called for in the procedure itself. When the rename process is complete, you can press forward with those types of activities.

STEP 1: Back Up All DCs

Simply put; make sure you accomplish a full backup of the system state of every DC in the entire forest. Get complete backups of CAs, DFS roots, FSMO masters, and other common-use and infrastructure items. Check the logs of every backup for any errors that might have occurred, and perform one or more restores on test systems to ensure your backup processes are sound.

STEP 2: Prepare the Control Station

Pick a conveniently located server that meets the following requirements:

- Must be a member of one of the domains that will be renamed.
- Must be a member server, *not a DC*, running Windows Server 2003 Standard Edition, Windows Server 2003 Enterprise Edition, or Windows Server 2003 Datacenter Edition.
- Must have available the installation CD for the version of Windows installed on the control station.

This server will function as the administrative *control station* for the entire domain rename operation. You will be running and controlling all rename procedures from this station. Special tools will need to be installed on the local disk, and scripts created and edited here. Although you do have to contact each DC in the forest, you will reach them remotely from this station.

Perform Exercise 4.21 to install Windows Server 2003 Support Tools from the OS CD.

EXERCISE 4.21

SETTING UP THE CONTROL STATION WITH THE REQUIRED TOOLS FOR THE DOMAIN RENAME OPERATION

1. Log on to the control station with at least local administrator rights.
2. Create a directory named `X:\RenameTools` on a local disk drive, where `X` is a local drive letter.



NOTE

You will be executing all rename tools from within this directory. Give it a name you will be comfortable with, but that will alert other administrators as to what it is used for.

3. Insert the Windows Server 2003 Standard Edition, Windows Server 2003 Enterprise Edition, or Windows Server 2003 Datacenter Edition operating system CD into the CD-ROM drive.
4. Open a command prompt and copy the files from the `valueadd` directory, shown here with the CD-ROM in drive letter `G`. Use the command **copy G:\valueadd\msft\mgmt\DomainRename*. * X:\RenameTools**.
5. Verify that `random.exe` and `gpfixup.exe` are present in your working directory on the control station.
6. Browse to the `G:\Support\Tools` directory, and double-click **Setup** to install the Windows Server 2003 Support tools.
7. Verify that `repadmin.exe` and `dfsutil.exe` are installed on the control station.

STEP 3: Document the Current Forest

In this step, you will use the `random` utility to generate a description of your current forest structure as an XML-encoded file. It will contain a list of all the domain directory partitions as well as your application directory partitions within the forest. This is the baseline that you start from, and you will modify this file later.

If you want to use authorization other than the account you are logged on as, use alternate credentials with the `/user` and `/pwd` command-line switches of `random.exe`. Perform Exercise 4.22 to generate the baseline forest description.

EXERCISE 4.22

GENERATING THE CURRENT FOREST DESCRIPTION FILE

1. Log on to the control station with Enterprise Administrator rights.
2. Open a **Command Prompt**, change to the RenameTools directory, and type **random /list**.
3. Type the following command to save a copy of the forest description file for future reference: **copy domainlist.xml domainlist-save.xml**

The */list* option of the *random.exe* creates an XML-encoded file named *domainlist.xml* in the current directory. It contains a textual description of your current forest structure, including a list of all the application directory partitions and domain directory partitions within your forest. You will find an entry for each of these domain and application directory partitions bounded by the `<Domain></Domain>` XML tags, as shown in Figure 4.53. Each entry contains naming data that includes the object GUID of the partition root object, the DNS name of the domain or application directory partition, and the NetBIOS name of the domain. An application directory partition does not have a NetBIOS name.

In the example, the *domainlist.xml* file shows the structure of a forest containing two domains called Zoo.net and Fish.zoo.net with NetBIOS names of ZOO and FISH, respectively. Three other entries appear, corresponding to the application directory partitions used by the Active Directory integrated DNS service. These application directory partitions must also be renamed:

- DomainDnsZones.fish.zoo.net
- DomainDnsZones.zoo.net
- ForestDnsZones.zoo.net

The entry for an application directory partition is annotated with an XML comment in the form:

```
<!-- PartitionType:Application -->
```

The entry for the root domain of the forest is also annotated with an XML comment in the form:

```
<!-- ForestRoot -->
```

For this procedure, the critical fields are those bounded by the `<DNSname></DNSname>` and `<NetBiosName></NetBiosName>` tags.

Figure 4.53 Forest Description File, domainlist.xml, Generated by Rendom.exe

```

<?xml version = "1.0"?>
<Forest>
  <Domain>
    <!-- PartitionType:Application -->
    <Guid>3dacd6ab-d4e8-b94e-63b9-8bacc5d5e10b</Guid>
    <DNSname>DomainDnsZones.fish.zoo.net</DNSname>
    <NetBiosName></NetBiosName>
    <DcName></DcName>
  </Domain>
  <Domain>
    <Guid>73cfaaec-faa6-4732-a851-c8050763fab0</Guid>
    <DNSname>fish.zoo.net</DNSname>
    <NetBiosName>FISH</NetBiosName>
    <DcName></DcName>
  </Domain>
  <Domain>
    <!-- PartitionType:Application -->
    <Guid>4d1c9e17-18bb-c642-1f4c-ac1be13d117</Guid>
    <DNSname>ForestDnsZones.zoo.net</DNSname>
    <NetBiosName></NetBiosName>
    <DcName></DcName>
  </Domain>
  <Domain>
    <!-- PartitionType:Application -->
    <Guid>4d1c9e17-18bb-c642-1f4c-ac1be13d117</Guid>
    <DNSname>DomainDnsZones.zoo.net</DNSname>
    <NetBiosName></NetBiosName>
    <DcName></DcName>
  </Domain>
  <Domain>
    <!-- ForestRoot -->
    <Guid>73cfaaec-faa6-4732-a851-c8050763fab0</Guid>
    <DNSname> zoo.net</DNSname>
    <NetBiosName>ZOO</NetBiosName>
    <DcName></DcName>
  </Domain>
</Forest>

```

STEP 4: Design the New Forest

From the preliminary procedures, your new forest structure should be well documented. In this step, you will overlay the new structure over the old one as described by the domainlist.xml file generated in step 3. You will use a text editor to edit the domainlist.xml file and replace your old domain names with new ones. You must also edit the names of application directory partitions in the same way. When you rename an Active Directory domain, the corresponding DNS-specific application directory partition must also be renamed, *if you are using Active Directory-integrated DNS*. If it is not, new DNS servers added to the network will not automatically load the DNS zones stored in the DNS-specific application directory partition and will not function properly.

Use a simple text editor such as Notepad to make the changes. Remember that at this time you can change the NetBIOS name of any domain, the DNS name of any domain, or both. In addition, take care to change the names of any child domains affected by a renamed parent. Review all name changes for *well-formed* characteristics as described in the *Domain Rename Limitations in a Windows 2003 Forest* section in this chapter.



WARNING

Double-check this step, preferably with more than one person to ensure that the changes are complete and correct. These are the actual names that will be implemented, and any typographical or hierarchical error could translate into a nonfunctional forest or domain. If your target structure is not what you intended, you must perform the entire domain rename procedure again.

Here are some guidelines for name changes:

- Change the DNS name from old to new, which is the field bounded by the tags `<DNSname></DNSname>`.
- Change the NetBIOS name from old to new, which is the field bounded by the tags `<NetBiosName></NetBiosName>`.
- Change the domain-level DNS Application Partition name from old to new, which is the `DomainDnsZones.<domain DNS name>` field bounded by the tags `<DNSname></DNSname>`.
- Change the forest-level DNS Application Partition name from old to new, which is the `ForestDnsZones.<forest DNS name>` field bounded by the tags `<DNSname></DNSname>`.
- *Do not change* the GUID represented in the field bounded by the `<Guid></Guid>` tags.

Review Figure 4.54 for a sample of what the domainlist.xml file would look like after the following restructuring design overlay. In this case, both DNS and NetBIOS names were changed:

- The top-level DNS name of Zoo.net has been changed to the more politically correct name of Naturepreserve.net, while its child domain of fish.zoo.net has been changed to Aquatics.naturepreserve.net.
- The NetBIOS name of ZOO has been changed to PRESERVE, while the NetBIOS name of its child domain has been changed from FISH to AQUATICS.
- The domain-level DNS application partition name has been changed from DomainDnsZones.fish.zoo.net to DomainDnsZones.aquatics.naturepreserve.net.
- The domain-level DNS application partition name has been changed from DomainDnsZones.zoo.net to DomainDnsZones.naturepreserve.net.
- The forest-level DNS Application Partition name has been changed from ForestDnsZones.zoo.net to ForestDnsZones.naturepreserve.net.
- The GUID fields have not been modified.
- No NetBIOS names have been assigned to any *PartitionType:Application*.



NOTE

If you have a Microsoft TAPI dynamic directory for an Active Directory domain, you might have application partitions for the TAPI application data. There is normally one TAPI-specific application directory partition for each domain. When you rename an Active Directory domain, the corresponding TAPI-specific application directory partition is not renamed automatically. If they exist, you should change those as well.

Figure 4.54 Forest Description File, domainlist.xml, Customized for the New Forest Design

```
<?xml version = "1.0"?>
<Forest>
  <Domain>
    <!-- PartitionType:Application -->
    <Guid>3dacd6ab-d4e8-b94e-63b9-8bacc5d5e10b</Guid>
    <DNSname>DomainDnsZones.aquatics.naturepreserve.net</DNSname>
    <NetBiosName></NetBiosName>
    <DcName></DcName>
```

Continued

Figure 4.54 Forest Description File, domainlist.xml, Customized for the New Forest Design

```

</Domain>
<Domain>
  <Guid>73cfaaec-faa6-4732-a851-c8050763fab0</Guid>
  <DNSname>aquatics.naturepreserve.net</DNSname>
  <NetBiosName>AQUATICS</NetBiosName>
  <DcName></DcName>
</Domain>
<Domain>
  <!-- PartitionType:Application -->
  <Guid>4d1c9e17-18bb-c642-1f4c-ac1be13d117</Guid>
  <DNSname>ForestDnsZones.naturepreserve.net</DNSname>
  <NetBiosName></NetBiosName>
  <DcName></DcName>
</Domain>
<Domain>
  <!-- PartitionType:Application -->
  <Guid>4d1c9e17-18bb-c642-1f4c-ac1be13d117</Guid>
  <DNSname>DomainDnsZones.naturepreserve.net</DNSname>
  <NetBiosName></NetBiosName>
  <DcName></DcName>
</Domain>
<Domain>
  <!-- ForestRoot -->
  <Guid>73cfaaec-faa6-4732-a851-c8050763fab0</Guid>
  <DNSname> naturepreserve.net</DNSname>
  <NetBiosName>PRESERVE</NetBiosName>
  <DcName></DcName>
</Domain>
</Forest>

```

Rendom.exe will display the new forest structure in the domainlist.xml file in a user-friendly format using text indentation to reflect the domain hierarchy using the /showforest option. Type the command **rendom /showforest** from a command prompt in the Renametools directory after each change to the domainlist.xml file to verify the forest structure. Recheck the contents of this file after every edit until you are sure it is correct.

STEP 5: Draft Domain Rename Instructions

Now that you have built a definition of your new forest structure, it is time to generate the domain rename instructions that will implement the change. These instructions will execute individually and remotely on each DC in the forest. Not surprisingly, the Domain Naming Master plays a central role in the renaming process. The `rendom.exe` utility creates the specially formatted scripts, containing a sequence of directory updates, and writes them to the `msDS-UpdateScript` attribute on the Partitions container object in the configuration directory partition on the Domain Naming Master for the forest.

The new forest description must be available as the XML-encoded file `domainlist.xml`, which you created by editing the original forest description file in the previous step. Follow these steps to create the domain rename instructions.

1. Open a command prompt, and change to the **RenameTools** directory.
2. Enter the command **rendom /upload**.
3. Verify the existence of the state file `dclist.xml` in the RenameTools directory and that it contains an entry for every DC in your forest.

In addition to generating the domain rename instructions and uploading them into the Active Directory, the `rendom /upload` command also generates a state file called `dclist.xml` and writes it to the current directory of RenameTools. Rendom uses this state file to track the progress and state of each DC in the forest. This tracking continues throughout the remaining steps of the domain rename procedure. Refer to Figure 4.55 for an example of the `dclist.xml` file. It contains the states for two DCs in the Zoo.net domain called DC1 and DC2, and two DCs in the Fish.zoo.net domain called DC3 and DC4.

Figure 4.55 Examining the `dclist.xml` State File Used for Tracking the Progress of Domain Rename

```
<?xml version = "1.0"?>
<DcList>
  <Hash>zzzzzzzz</Hash>
  <Signature>zzzzzzzz</Signature>

  <DC>
    <Name>DC1.zoo.net</Name>
    <State>Initial</State>
    <LastError>0</LastError>
    <Password />
    <LastErrorMsg />
    <FatalErrorMsg />
    <Retry></Retry>
```

Continued

Figure 4.55 Examining the dclist.xml State File Used for Tracking the Progress of Domain Rename

```

</DC>
<DC>
  <Name>DC2.zoo.net</Name>
  <State>Initial</State>
  <LastError>0</LastError>
  <Password />
  <LastErrorMsg />
  <FatalErrorMsg />
  <Retry></Retry>
</DC>
<DC>
  <Name>DC3.fish.zoo.net</Name>
  <State>Initial</State>
  <LastError>0</LastError>
  <Password />
  <LastErrorMsg />
  <FatalErrorMsg />
  <Retry></Retry>
</DC>
<DC>
  <Name>DC4.fish.zoo.net</Name>
  <State>Initial</State>
  <LastError>0</LastError>
  <Password />
  <LastErrorMsg />
  <FatalErrorMsg />
  <Retry></Retry>
</DC>
</DcList>

```

Every DC in the forest must have an entry in the dclist.xml state file. The state of each DC is delineated by the tags `<State></State>`. At this point in the rename process, all DCs are initialized to the *Initial* state. As the procedure progresses, these states will change.

STEP 6: Push Instructions to DCs

In Exercise 4.23, you will trigger a forced Active Directory replication. This pushes out the domain rename instructions that you previously uploaded to the Domain Naming Master to all DCs in the forest. Then, you should verify that the DC Locator SRV records registered in DNS by each DC for the new domain names have replicated to all DNS servers that are authoritative for those records.



NOTE

You do not have to force replication, but it will accelerate the replication of the changes to the Partitions container in the configuration directory partition to all DCs in the forest. Alternately, you can wait for replication to complete according to the delay characteristics and replication intervals of your forest.

EXERCISE 4.23

FORCING THE SYNCHRONIZATION OF CHANGES MADE TO THE DOMAIN NAMING MASTER

1. Open a command prompt, and change to the **RenameTools** directory.
 2. Type the command **repadmin /syncall /d /e /P /q DomainNamingMaster**. Note that **DomainNamingMaster** is the DNS host name of the current Domain Naming Master for the forest.
-



NOTE

Use the *dsquery* utility to determine which host is the Domain Naming Master, or perform Exercise 4.12. In addition, be aware that the *repadmin* command-line options are case sensitive.

It is critical that all DCs successfully replicate before continuing. If *repadmin* completes successfully, the Domain Naming Master DC will have replicated to every other DC in the forest. If you get an error for some of the DCs in the forest, you must try again until all DCs in the forest have successfully received the changes from the Domain Naming Master.

STEP 7: Verify DNS Readiness

During domain rename, the Net Logon service of each DC pre-publishes the SRV resource records to the authoritative DNS servers associated with their new domain name. The DC Locator will malfunction if these are not successfully published. Other records must also be in place for authentication and replication to take place. Refer to Exercise 4.24 to verify that the various DNS records for the new domain names were successfully created.



NOTE

Do not, under any circumstances, proceed with domain rename if any of the DNS records listed in Table 4.7 are missing. They are all required for normal operation of the forest.

EXERCISE 4.24

VERIFYING THE DNS SERVICE LOCATOR RECORDS

1. Click **Start | Programs | Administrative Tools | DNS** to start the DNS administrator console.
2. Expand the server name.
3. Expand the Forward Lookup Zones.
4. Expand the domain you want to verify.
5. Verify that the DNS records listed in Table 4.7 are present for all your DCs within each domain. These records are crucial to the operation of the domain.

Table 4.7 Required SRV Resource Records

Location of DNS Record	Type of Record	Purpose
<i>DsaGuid._msdcs.DnsForestName</i>	CNAME	There must be one CNAME record pertaining to every DC in all authoritative DNS servers. Without it, replication will not take place from that DC.

Continued

Table 4.7 Required SRV Resource Records

Location of DNS Record	Type of Record	Purpose
<code>_ldap._tcp.pdc._msdcs. DnsDomainName</code>	SRV	There must be one SRV record associated with the PDC FSMO on all authoritative DNS servers. Without it, the authentication of users and computers will not function correctly.
<code>_ldap._tcp.gc._msdcs. DnsForestName</code>	SRV	There must be at least one record pointing to at least one GC on all authoritative DNS servers in the forest. This record also ensures the proper functioning of authentication for users and computers. Minimally, at least one record of this type must be present on all authoritative DNS servers. Records for the other GCs should eventually replicate.
<code>_ldap._tcp.dc._msdcs. DnsDomainName</code>	SRV	There must be at least one record for at least one DC on all authoritative DNS servers for each domain. This is another record required for the functioning of authentication of users and computers. Minimally, at least one record of this type must be present on all authoritative DNS servers. Records for the other DCs should eventually replicate to the other authoritative DNS servers.

Now you need to check the status of every DC in the forest to verify that the Active Directory database is in a healthy state and ready to perform the domain rename instructions. Use Exercise 4.25 to accomplish this. The `Remom` tool will issue a Remote Procedure Call (RPC) individually to each DC in the forest and update the state file `dclist.xml`.

EXERCISE 4.25

VERIFYING THE READINESS OF ALL DOMAIN CONTROLLERS

1. Open a command prompt, and change to the **RenameTools** directory.
2. Type the command `remom /prepare`. The `remom` utility will check for the following:

Active Directory Container	Attribute	Is Replicated To
Partitions	<i>msDS-UpdateScript</i>	Every DC in the forest
Partitions	<i>msDS-DnsRootAlias</i>	Every DC in the forest
Partitions	<i>servicePrincipalName</i>	Every DC in a domain and the GC servers

3. Verify the existence of the state file `dclist.xml` in the `RenameTools` directory and that it contains an entry for every DC in your forest showing a state of *Prepared*.

The control station computer issues an RPC to every DC in the forest when you run the `random /prepare` command. The results are tracked by the state file `dclist.xml`. This RPC executes on each DC to verify that its directory replica is in a healthy state and that it is ready to run the domain rename instructions. `Random` updates the state field in each DC section of the `dclist.xml` file to the *Prepared* status as shown by (`<State>Prepared</State>`) if it is successful. If not, check the command execution log called `random.log` in the current working directory. It contains valuable information about the actual tasks performed by the tool, and at what stage or on which DC the problem occurred.

During this time, the forest must be *quiescent*. You must not make any changes to the Active Directory such as adding or removing domains, DCs, directory partitions, or trust relationships. `Random` will detect this type of activity and require you to execute `random /end`. If this happens, you will have to start over again from step 3 and reconfigure the forest configuration.

Do not proceed until all DCs are in the *Prepared* state.

STEP 8: Execute Instructions

The script you uploaded to the *msDS-UpdateScript* attribute on the Partitions container on every DC in the forest will now be executed. To execute the domain rename instructions, you will run the `random` utility, which causes the control station computer to issue an RPC to each DC in the forest individually. As the DCs receive their execute commands, they run the rename instructions that they have already received, and then reboot automatically. At the end of step 8, every DC tracked by the state file `dclist.xml` will be in one of two final states:

- **Done** The DC has successfully completed the domain rename operation.
- **Error** The DC has encountered an irrecoverable error and can never complete the domain rename operation.



NOTE

There will be a temporary disruption in service while the DCs respond to the *rendom* utility and reboot.

Executing the Domain Rename Instructions on All DCs

Now you need to use the **rendom** command again. This time, all of your preparation and planning will come to fruition. Think of this stage as “flipping the switch.” Follow these steps to run the *Renderom* utility from a command prompt and then check the state file *dclist.xml* for the status of each DC. The *Renderom* command must be repeated until all DCs have either successfully executed the domain rename, or you have established that one or more DCs are unreachable and will be removed from the forest.

1. Open a command prompt, and change to the **RenameTools** directory.
2. Type the command *rendom /execute*.
3. Verify the existence of the state file *dclist.xml* in the **RenameTools** directory and that it contains an entry for every DC in your forest showing a state of Done or Error.
4. If *dclist.xml* shows any DCs still in the Prepared state, repeat step 2 as many times as needed until the stopping criterion is met, which is the Done or Error state.



NOTE

Don't worry about running the command multiple times, because *rendom /execute* skips any DC in the Done or Error state no matter how many times you run it, and only retries those in the Prepared state.

The control station issues an RPC to every DC in the forest that is known to the state file *dclist.xml*, which commands it to execute the rename instructions that it already has, and then reboot. The DC's entry in the state file will be updated to read `<State>Done</State>`. If it fails with a fatal or irrecoverable error, the DC's state file entry will be updated to read `<State>Error</State>`. When it reaches the Error state, the error code is written to the last error field `<LastError></LastError>` and a corresponding error message is written to the `<FatalErrorMsg></FatalErrorMsg>` field.

If a DC reached the Error state in the *dclist.xml* file, but you believe it is a recoverable error that you have since corrected, you can force the *rendom /execute* command to retry issuing the RPC to that DC as described in the following steps.

Forcing `random /execute` to Re-Issue the RPC to a DC in the Error State

To force `Random / execute` to re-issue the RPC to a DC in the Error state follow these steps:

1. Open a command prompt, and change to the **RenameTools** directory.
2. Locate the `<Retry></Retry>` field in the `dclist.xml` file for the DC that you want to retry.
3. Change the entry to `<Retry>yes</Retry>` for that DC, and resave the file. The next execution of the `random /execute` command will re-issue the execute-specific RPC to that DC.
4. Type the command `random /execute`.
5. Check the state file `dclist.xml` to see if it contains an entry for the retried DC. If it says the state is Done, then the RPC was successful; if it shows a state of Prepared, then retry step 4; if it shows a state of Error, you can start this exercise over at step 2 or proceed to step 6.
6. Declaring the Execute Instructions stage complete is at your discretion. If all the DCs show a state of Done or Error, you can stop. You can retry execution attempts repeatedly if you have an Error, but when you decide to stop trying, you must remove Active Directory from any DC still in the Error state.



NOTE

Don't worry about running the command multiple times, because `random /execute` skips any DC in the Done or Error state no matter how many times you run it, and only retries those in the Prepared or Error/Retry state.

One final warning: The DNS host names of the DCs in the renamed domains do not change automatically in the manner of the member servers. In other words, the fully qualified DNS host name of a DC in the renamed domain will continue to reflect the old domain name in its DNS suffix. A different procedure is required for that. See the section *Renaming a Domain Controller*.

STEP 9: Unfreeze the Forest Configuration

During this procedure, the forest should have been quiescent. You should not have attempted any changes to the Active Directory such as adding or removing domains, DCs, directory partitions, or trust relationships other than those called for in the procedure. Perform the steps in Exercise 4.26 to unfreeze the forest configuration. Now that the core portion of the rename process is complete, you can press forward with those types of activities, although the best

practice is to finish the rename procedure first, along with any post-procedure tasks that are needed. When the forest is stable and healthy, then proceed normally.

EXERCISE 4.26

UNFREEZING THE FOREST CONFIGURATION

1. Reboot the control station twice. This will ensure that all local services learn of the new DNS and/or NetBIOS name of the control station's domain.
 2. Open a command prompt, and change to the **RenameTools** directory.
 3. Type the command *random /end*. Note that this command removes the *msDS-UpdateScript* attribute from the Partitions container of the Domain Naming Master FSMO.
-

STEP 10: Re-establish Trusts

The intraforest shortcut trusts created in the preliminary steps are automatically adjusted during the domain rename operation so they will continue to work. Unfortunately, your external trusts, including cross-forest trusts, were not protected and must be re-established. You must be a member of the Domain Admins group in the target domain of the external trust to delete and recreate inter-forest and external trusts. Refer to Chapter 5 for the trust-creation procedures.

STEP 11: Repair DFS Topology

In this step, you will repair references to a renamed domain in the DFS topology data. The *dfsutil.exe* command-line tool will accomplish this. *Dfsutil* scans the entire topology for a given DFS root including the root name, root replica servers, and link target servers, and fixes any occurrences of the *oldname* with the *newname* as specified on the command line. *Dfsutil* also connects to DFS root replica servers and changes the topology information held in its local Registry there as well.

To perform Exercise 4.27, you must be a member of the Domain Admins group in the target domain of the DFS fix-up. The DFS utility might need to be run more than once to fix the topology for every DFS root. All DFS root servers in a renamed domain must be running Windows 2000 with Service Pack 3 or higher.

EXERCISE 4.27

REPAIRING THE DFS TOPOLOGY

1. Use the DFS MMC snap-in or the *dfsutil.exe* utility to examine the DFS topology. The first step is to prepare a list of DFS roots where a root path, root replica server name, or a link target server name needs to be fixed as a result of renaming the domain. Subsequent steps describe when each topology component needs to be repaired. Note that the *DFSutil.exe* utility can only repair domain-based DFS root topologies, not stand-alone DFS roots.
2. Examine your DFS topology for any domain-based DFS root paths. They need to be changed in the topology when the domain name changes. For example, if the name of the domain zoo.net changed to naturepreserve.net, then a domain-based DFS root named \\ zoo.net \public would need to be changed to \\naturepreserve.net\public. Remember that if the root path uses the NetBIOS name of the domain, and the NetBIOS name of the domain was not changed, then it does not need to be repaired.
3. Any domain-based DFS root *replica* host name or *link target* host name needs to be changed in the DFS topology during the domain rename process if it is specified as a fully qualified DNS name. For example, the DNS host name of a DFS replica server named guppy.zoo.net might change to guppy.naturepreserve.net because of the domain name change. Likewise, a DFS link might need to be changed from \\poodle.zoo.net\good places to bury a bone to \\poodle.naturepreserve.net\bones. Again, NetBIOS names only change if they were instructed to change during the rename process.
4. On the control station, open a command prompt and change to the **RenameTools** directory. For every DFS root that requires a repair of any topology component described previously, type the following command (the entire command must be typed on a single line): **DFSutil /RenameFtRoot /Root:DFSRootPath /OldDomain:OldName/ NewDomain:NewName /Verbose**. Note that:
 - *DFSRootPath* is the DFS root to operate on; for example, \\naturepreserve.net\public.
 - *OldName* is the exact old name to be replaced in the topology for the DFS root.
 - *NewName* is the exact new name to replace the old name in the topology.

5. Check your fixed DFS topology again to confirm that it reflects the renamed domain.
 6. Reboot all DFS root replica servers twice to refresh the DFS service topology information.
 7. Repeat steps 1 through 6 for every renamed domain. You can enter the commands in sequence.
-

STEP 12: Repair Group Policy Objects and Links

GPOs and GPO references in each renamed domain need to be repaired with the *gpfixup.exe* command-line tool as shown in Exercise 4.28. These GPOs and their links still have the old domain name embedded in their properties, and will not function normally until repaired. Managed software deployment is also impaired because Group Policy-based software installation and maintenance data such as software distribution point network paths can also be based on the domain name. *Gpfixup.exe* will repair these for you as well, and it needs to be run once in each renamed domain. Moreover, since GPOs cannot reference application directory partitions, there is no repair required on those. This step completes the core domain rename procedure; however, many other steps could be necessary depending on your configuration as shown in the section *Steps to Take After the Domain Rename Procedure*.



NOTE

The GPO/link fix-up procedure does not repair interdomain GPO links. If you have any of these in your forest, they will have to be repaired manually. This is a matter of breaking and reestablishing each link. It also does not repair network paths for software distribution points that are external to the domain.

EXERCISE 4.28

REPAIRING GPOS AND LINKS

1. Open a command prompt. Click Start | All Programs | Accessories | Command Prompt.
2. Change to the **RenameTools** directory.
3. Type this entire command on a single line: **gpfixup /olddns:OldDomainDnsName /newdns:NewDomainDNSName /oldnb:OldDomainNetBIOSName**

`/newnb:NewDomainNetBIOSName/dc:DcDnsName 2>&1`
`>gpfixup.log`. In this case:

- *OldDomainDnsName* is the old DNS name of the renamed domain.
- *NewDomainDnsName* is the new DNS name of the renamed domain.
- *OldDomainNetBIOSName* is the old NetBIOS name of the renamed domain.
- *NewDomainNetBIOSName* is the new NetBIOS name of the renamed domain.
- *DcDnsName* is the DNS host name of a DC in the renamed domain, preferably the PDC Emulator. Pick one that successfully completed the rename operation with a final Done state in the `dclist.xml` state file in step 8.



NOTE

The command line parameters `/oldnb` and `/newnb` are only required if the NetBIOS name of the domain changed; otherwise, these parameters can be omitted from the command line for `Gpfixup`. In addition, the redirected output—both status and errors—is saved to the file `gpfixup.log`, which can be periodically displayed to monitor progress of the command.

4. To force replication of the Group Policy repair changes to the rest of the DCs in the renamed domain, type **`repadmin /syncall /d /e /P /q DcDnsName NewDomainDN`** and then press **Enter**. In this case:
 - *DcDnsName* is the DNS host name of the DC that was targeted by the `gpfixup` command.
 - *NewDomainDN* is the DN corresponding to the new DNS name of the renamed domain.



NOTE

Remember, the DNS host name of a DC in a renamed domain does not change automatically when the domain name changes. Use the old name unless you have changed it manually to the new one at this point.

5. Repeat steps 3 and 4 in this procedure for every renamed domain. You can do them in sequence. For two domains, execute **`gpfixup`** twice and

repadmin twice. *Do not* run *gpfixup* more than once for each renamed domain, and do not run it at all for renamed application directory partitions.

Steps to Take After the Domain Rename Procedure

Follow the instructions in this section to be sure that all functionality relying on an accurate domain name has been addressed. One of the main concerns involves your Enterprise CA.

Verifying Certificate Security

If you use enterprise certificates, perform all of the following procedures after domain rename is complete. Your enterprise CA can be configured with both LDAP and HTTP URLs pointing to your CRL. Determine your certificate attributes before continuing.



NOTE

If you only have LDAP URLs in your enterprise certificates, then all previously issued certificates will stop working when you rename the domain. The only available workaround for correcting the LDAP CDP and AIA pointers is to renew your entire CA hierarchy and reissue all End Entity certificates. This will result in PKI downtime until you resolve these issues.

Preparing URLs for CDP and AIA Extensions After Domain Rename

To ensure that your old enterprise certificates operate properly after the domain rename, make a CNAME DNS record redirecting the old HTTP server name to the new DNS name for the server, as shown in Exercise 4.29. This refers to the host servicing the CRLs of your CA. With this redirection, the HTTP URLs in the old certificates will continue to be valid. This is needed so that the client machines will be able to obtain CRLs and CA certificates for verification purposes.

EXERCISE 4.29

CONFIGURING DNS FOR THE REDIRECTING ALIAS ENTRY

1. Click **Start | Administrative Tools | DNS**.
2. Expand the *DNS server* entry.
3. Right-click the old DNS zone.

4. Select **New Alias (CNAME)** .
5. In the **Alias name** box, type the original FQDN of the HTTP server.
6. In the **Fully qualified domain name for target host** box, type the new FQDN of the HTTP server, and then click **OK**.
7. Test the new mapping by pinging the FQDN of the old HTTP server from a command prompt. The ping should be redirected automatically to the new FQDN of the HTTP server.



NOTE

You can remove the CNAME record when you are assured that all existing certificates have been renewed.

Verifying the Use of User Principal Names

Smart card logons via Kerberos require that the UPN in the user certificate match the UPN in the user account in Active Directory. The two types of UPNs are *implicit* and *explicit*.

- **Implicit UPN** The absence of an explicitly assigned value for its UPN attribute means that a user account is assumed to have an *implicit UPN* for authentication purposes based on the DNS name of the domain in which the account exists. If the DNS name changes, the implicit UPNs of all user accounts in the domain also change. Both the old and the new implicit UPNs will be accepted for authentication until the cleanup step described later. After that step, only the new implicit UPN will be accepted.
- **Explicit UPN** A domain user account is said to have an *explicit UPN* if it has an explicitly assigned value for its UPN attribute. If the DNS name changes, the explicit UPNs of user accounts are *not* impacted, and no additional action is needed to maintain certificate functionality.

Enabling Certificate Enrollment in the Renamed Domain

In the new domain, you need to enable certificate enrollment using either autoenrollment or the Certificates MMC snap-in. To accomplish this, a small change has to be made in Active Directory to the Enrollment Services container in the configuration directory partition (cn=*Enrollment Services*,cn=*Public Key Services*, cn=*Services*, cn=*Configuration*, dc=*ForestRootDomain*). This container holds a CA object with a *dNSHostName* attribute containing the old DNS name of the CA machine. For convenience, you can use the Visual

Basic script shown here to change the value of this attribute. This is an advanced procedure, and should only be performed by experienced systems administrators and systems engineers. Do not attempt it if you are unfamiliar with scripting, Registry editing, and the dangers involved with both.

Perform the following procedure for each CA in your domain. In addition, note that you will have to set the container name, CA name, CA host name, and DNS names to match your domain configuration.

Execute the following commands within a Visual Basic script, or use this information with any other LDAP administration tool. It references your CA and changes the *dnshostname* object attribute:

```
name = "LDAP://CN=CAName,CN=Enrollment Services,CN=Public Key Services,
      CN=Services,CN=Configuration,DC=SubDomain,DC=Domain,DC=RootDomain"
Set obj = GetObject(name)
Obj.dnshostname = "TheCAMachine'sNewDNSName"
Obj.setinfo
```

Change the Registry on each CA machine to reflect its new DNS name. On the CA machine, open the Registry editor and locate the entry **CAServerName** under HKLM\System\CurrentControlSet\CertSvc\Configuration**YourCAName**. Change the value in **CAServerName** to correspond to your new DNS host name.

To enable proper Web enrollment, update the file used by your Web enrollment ASP pages. On each CA machine, search for the file named certdat.inc. On a default install, it is located in the %windir%\system32\certsrv folder. Save a copy of this file before editing, open it using Notepad, change the DNS name of the CA machine in the file, and resave it as certdat.inc. Note that the *sServerType* value can be either “Enterprise” or “StandAlone.” The value you should change is marked in the following code in bold, (*sServerConfig*). The certdat.inc file looks something like this:

```
<% CODEPAGE=65001 UTF-8%>
<% certdat.inc - (CERT)srv web - global (DAT)a
   ' Copyright (C) Microsoft Corporation, 1998 - 1999 %>
<%   ' default values for the certificate request
      sDefaultCompany=""
      sDefaultOrgUnit=""
      sDefaultLocality=""
      sDefaultState=""
      sDefaultCountry=""

      ' global state
      sServerType="Enterprise"
      sServerConfig="DNSName\CAName"
```

```
sServerDisplayName="CAName"
nPendingTimeoutDays=10

' control versions
sXEnrollVersion="5,131,2510,0"
sScrdEnrlVersion="5,131,2474,0"

%>
```

Update the Shared Folder option. You also need to edit the `certsrv.txt` file to reflect the new DNS name of the CA machine if the CA was installed with the shared folder option. This file is located in your shared folder.

Verify the validity of CDP and AIA extensions. Your CDP and AIA extensions are sometimes hard coded. If they are, you must change the extension URLs to reflect the new DNS name of the CA machine. This is done in the following manner:

1. In the Certification Authority MMC snap-in, right-click the **CA name** and select **Properties**.
2. On the **Extensions** tab, check the CDP and AIA extensions, paying particular attention to the portion bolded in the following examples:
 - Flexible extensions have the following format:
`http://<ServerDNSName>/CertEnroll/<CAName>
 <CRLNameSuffix><DeltaCRLAllowed>.crl.`
 - Hard-coded extensions have the following format: `http://dnsname.compa-
 nyname.com/certenroll/<CAName><CRLNameSuffix><DeltaCRLAllow
 ed>.crl`
3. If the CDP and AIA extensions are flexible, then take no action.
4. If the CDP and AIA extensions are *not* flexible, or hard coded, change the extension URLs to reflect the new DNS name of the CA machine.

Only if the CA is running on Windows 2000, change the Registry to reflect the new domain name for the LDAP extension for the CDP as shown in the following steps:

1. On the CA machine, open the registry editor and locate the entry `LDAPRevocationDN` under `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration\YOURCANAME\ExitModules\CertificateAuthority_MicrosoftDefault.Exit`.
2. Change the value in `LDAPRevocationDN` to reflect the new domain name for the LDAP extension.

That covers the procedure necessary to complete the renaming process for the CA server itself, but other servers also need changes to complete the process. You will need to

perform the following steps on servers hosting functions such as Web proxy, revocation, and enrollment within your CA hierarchy. As before, be careful with changes to the Registry and ensure that you have good backups before performing these steps.

First, update the CA Web proxy. If you have a Web proxy machine *for CA web pages* whose DNS host name has changed, then you need to change the following Registry key:

1. Open the registry editor and locate the entry `WebClientCAMachine` under `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration`.
2. Change the value in `WebClientCAMachine` to correspond to your new CA Web Proxy DNS host name.

Now, update the Netscape revocation check mechanism:

1. On all machines where Web pages for the CA reside, such as on the Web proxy itself and on the CA servers, look for a file named `nsrev_CANAME.asp` containing the DNS host name of the CA machine that is used by the Netscape revocation checking mechanism. Under the default installation settings, `nsrev_CANAME.asp` will be in the folder `%Windir%\system32\certsrv\certenroll`, and its contents will be similar to:

```
<%
Response.ContentType = "application/x-netscape-revocation"
serialnumber = Request.QueryString
set Admin = Server.CreateObject("CertificateAuthority.Admin")
stat = Admin.IsValidCertificate("CAMachineDnsHostName\CANAME",
serialnumber)
if stat = 3 then Response.Write("0") else Response.Write("1") end if
%>
```

2. Open `nsrev_CANAME.asp` with Notepad and change the `CAMachineDnsHostName`, as bolded in the preceding entry, to correspond to the new DNS host name.

Finally, change the User Identity for SCEP Add-on: If the Simple Certificate Enrollment Protocol (SCEP) Add-on for Microsoft Certificate Services is installed, check the user context where the MSCEP process runs. Since the IIS metabase is not altered during domain rename, you might need to change this username to reflect the new name `NewDomainname\UserName`. To change the user identity for SCEP in IIS:

1. In the IIS MMC snap-in, browse to **Application pools**.
2. Under **Application pools**, right-click the folder for **SCEP** and select **Properties**.
3. On the **Identity** tab, change the username as described previously.

Renewing Subordinate and Issuing CA Certificates

After you have performed Exercises 4.29, 4.43, and 4.44 on all CA or CA-related servers as called for, you should renew all certificates to update the CDP and AIA locations. Start with all subordinate and issuing CAs certificates in hierarchical order, starting from the top. After that is complete, update the group policy on all machines to ensure that the new root CA certificates have full distribution.

Publishing new Certificate Revocation Lists

On all CA machines in the renamed domain, publish new Delta and Base CRLs by running the **certutil.exe -crl** command.

Updating Domain Controller Certificates

Any authentication mechanism based on certificates, such as replication and smart cards, requires an update to the DC certificates. If template-based autoenrollment was set before the domain rename procedure, these certificates can be updated by incrementing the version number of the Domain Controller Authentication and Directory Email Replication Certificate templates to force re-enrollment. If autoenrollment was not already set, roll out a Group Policy setting Machine-Based Autoenrollment. When that takes effect, the DC machines will re-enroll and update the existing V1 Domain Controller Certificate. If you increase the version number on other templates, especially those related to authentication, they will also trigger autoenrollment for users and their machines.

Miscellaneous Tasks

Microsoft recommends a number of other follow-up tasks to the domain rename procedure. These tasks should be performed, but no particular order is necessary.

Publishing Service Connection Points for Renamed TAPI-Specific Application Directory Partitions

During the domain rename procedure, you might have renamed TAPI-specific application directory partitions. If so, then you need to republish service connection points for the new name of the application directory partition. This will allow TAPI clients to locate the newly named partition. At the same time, you need to remove the service connection points for the old name.

This is a simple matter of removing the old service connection point by executing the command

```
tapicfg removescp /directory:mstapi.olddomainname /domain:newdomainname
```

You then need to publish a new service connection point for the new application directory partition name by executing the command:

```
tapicfg publishscp /directory:mstapi.newdomainname /domain:newdomainname /forcedefault
```

You should execute both commands from the control station.

Orchestrating a Password Reset for Digest Authentication

If you are using a Digest authentication mechanism using the DNS domain name as the *realm*, then Digest authentication cannot be used after a domain rename until a given user's password is changed. You will have to ensure that all users reset their passwords. Digest authentication does the same thing as Basic authentication, but it provides a security improvement in the way in which a user's credentials are sent across the network. Using Digest authentication, credentials are transmitted across the network as an MD5 (message digest) hash. When this is done, the original username and password cannot be easily deciphered using packet-sniffing tools. Here are some suggestions on how to accomplish this:

1. Expire all user passwords by changing your domain password policy in the renamed domain.
2. Send out an e-mail warning users that they must change their passwords immediately after they reboot their machines for the second time after the rename.

Users change their passwords by using **Ctrl+Alt+Del** and clicking the **Change Password** button.

Remove Any Redundant Inter-Domain Trusts within Your Forest

During forest restructuring, as opposed to simple domain renaming, you created additional shortcut trusts to preserve complete trust between all domains in your new forest. Sometimes, there will be old trust relationships that are no longer needed due to the new structure. Review all trusts for obsolete or incorrect relationships, and use the Active Directory Domains and Trusts MMC snap-in to remove them.

Repairing Start Menu Shortcuts for the Security Policy MMC Snap-Ins

The shortcuts to the Domain Security Policy and Domain Controller Security Policy MMC snap-ins in the Start menu are broken by the name change as well. The following steps show how to repair these snap-ins. Perform them on every DC in every renamed domain as needed.

To repair the shortcut for the Domain Security Policy snap-in:

1. Click **Start | Programs | Administrative Tools**.
2. Right-click **Domain Security Policy** and select **Properties**.
3. Edit the **Target** field to replace the old *domainname* that appears as part of the **/gobject:** parameter with the new *domainname*, click **OK**.

To repair the shortcut for the Domain Controller Security Policy snap-in:

1. Click **Start | Programs | Administrative Tools**.
2. Right-click **Domain Controller Security Policy** and select **Properties**.

3. Edit the **Target** field to replace the old *domainname* that appears as part of the **/gobject:** parameter with the new *domainname*, and click **OK**.

Removing Group Policy Setting the Primary DNS Suffix of Member Computers in Renamed Domains

If you performed computer renames in groups to avoid excess replication traffic, you configured and applied a group policy to apply the primary DNS suffix setting to member computers in your renamed domains. This policy can now be removed.

Removing DNS Zones that Are No Longer Needed

With name changes occurring, some of the DNS zones in your DNS infrastructure might be obsolete. Old records might still exist. Use the DNS administration tool to remove the old DNS zones.

Backing Up Domain Controllers

The content of the Active Directory database, system Registry, and GPOs on the DCs will have changed as a result of the domain rename operation. For this reason, your pre-procedure backups are no longer up to date, and new ones should be performed right away.

As a post-domain rename procedure, perform the following steps:

1. Back up the system state: Accomplish a full system state backup of all DCs in the forest so that you have a recoverable backup state.
2. Back up GPOs: Download and install the Group Policy Management Console (GPMC) from the Microsoft Web site. GPMC makes Group Policy easier to use, and adds functional improvements such as the ability to back up GPOs independently of the rest of Active Directory. If you backed up GPOs with GPMC prior to the domain rename operation, those objects cannot be restored afterward from those backups.
3. Back up all data volumes.

You should back up the system state and all data volumes at the same time. This precaution prepares you in case you need to rebuild the server. Better yet, create an Automated System Recovery (ASR) backup set. You should do this whenever the OS changes, such as when you install new hardware, drivers, or a service pack. It especially holds true after a domain rename. With an ASR backup set, you can recover from a system failure with more confidence and ease. You need to back up all data volumes at the same time because ASR protects only the system areas, not the data.

The system state is a collection of system-specific data within the OS that must be backed up as a unit. It does not provide a replica of the entire system. The system state data includes the following: the Registry, the COM+ Class Registration database, system files, boot files, and all files covered by Windows File Protection. If the server is a certificate

server, the system state backup also includes the Certificate Services database. If the server is a DC, it includes the SYSVOL directory and the Active Directory database. Finally, if the server is a node in a cluster or a Web server, the system state includes the cluster database information or the IIS Metabase, respectively.

Restart Member Computers

All member computers in the renamed domains in your forest must be restarted to trigger the domain membership changes. Some tips for performing this step include:

- Reboot member computers twice.
- Member computers on a wireless LAN (WLAN) should be connected to a wired network while performing the two required reboots. Otherwise, eject the wireless network card and then reinsert it after logon prior to each reboot.
- Disjoin and rejoin any remote computers that connect to the renamed domain through a remote connection such as dial-up and VPN.
- Disjoin and rejoin any Windows NT 4.0–based computers.

Attribute Clean-Up After Domain Rename

As the last mandatory procedure, this step removes all values of the msDS-DnsRootAlias and msDS-UpdateScript attributes from Active Directory that were written during the domain rename operation by connecting to the Domain Naming Master DC. Simply execute the `rdm /clean` command from a command prompt at the control station.

Rename Domain Controllers

As previously noted, the DNS host names of DCs in the renamed domains do not change automatically. This host name can be changed using a special procedure. The time it takes for the name change to replicate throughout the DNS and Active Directory databases might cause a temporary inability of clients to locate or authenticate to the renamed computer. For that reason, renaming a DC requires that you follow a computer rename preparation procedure. This precludes service interruptions.

In Windows 2000, the process of renaming a DC involves demoting, renaming, and promoting again. Since the process of re-promoting the DC requires a replication of the Active Directory, this can take a long time. Renaming a DC in Windows Server 2003 is much easier than it was in 2000, but that does not mean it is a simple procedure. Before renaming, you need to transfer operational master roles to another computer. If it is a GC server, you should move that role as well. Remember that the first DC that you installed in the forest is the root DC. This DC is responsible for the GC and for all FSMO roles unless you have spread them out manually. In this case, you need to install a second DC and transfer all of these functions to it before you rename the first server.

The way in which Windows 2003 implements this new capability is through a new Computername option in Netdom.exe. This is a relatively quick and safe procedure, requiring only a reboot. The first step is to install the Support Tools from the Windows Server 2003 distribution CD on every DC that you want to rename. If you are performing a domain rename, then you are already at the Windows Server 2003 functional level. Domain controller rename is only supported at that level.

There are several cautions when changing a DC's computer name:

- It might be tempting to go directly to **My Computer | Properties | Computer Name** and click the **Change** button just as you would for any other computer. Windows 2000 grayed out the option, but in Windows Server 2003, the button is functional. Windows allows this simpler version of renaming with only a warning, but replication latency renders the DC unavailable for a period of time even though the DNS SRV records are still offering the DC's services to the network. The safe method is to use the netdom utility.
- Don't be confused by the command syntax. Microsoft calls the option to change the computername "computername." This can easily be confused with the actual name of the computer that you need to type into the command line.
- Most command-line utilities in Windows Server 2003 call for a computername argument. Typically, this requires a DNS name, but in the case of the domain controller rename, it expects the NetBIOS name instead. If you are unsure of the distinction, use the *NBTSTAT-n* command from a command line on your DC. The output shows only NetBIOS names and their identifying suffixes. The name registered for *Domain <00>* is the NetBIOS *domain name*, while the *Computername <20>* record, identifying the *Server* service, is also the computer's NetBIOS *host name*.



NOTE

The *NBTStat* command gives you access to a lot of information. Microsoft Knowledge Base article 163409 describes NetBIOS suffixes, (which are identified by the 16th character of the NetBIOS name), and article 119495 lists and describes some of the names that register with the WINS Service. *NBTStat*, within its various options, gives you NetBIOS services, sessions, status, and statistics. It can also remove, correct, and reload NetBIOS name cache entries, and display the primary Media Access Control (MAC) address of a remote NetBIOS-compatible host.

Follow these steps to rename your Windows Server 2003 DC:

1. Log in as a domain administrator and open a command prompt.
2. Execute the rename command:

**Netdom Computername OldComputerNetBIOSname /add:
NewComputerFQDN**

3. Verify the secondary name with the following command:

Netdom Computername OldComputerNetBIOSname /enumerate

The command will report old and new DNS names. Allow some time for the computer account to be replicated throughout the domain, and the DNS resource records to be distributed to the authoritative DNS servers.

4. Select the new computername as the primary one:

**Netdom Computername OldComputerNetBIOSname /makeprimary:
NewComputerFQDN**

The */enumerate* option should have displayed both computer names. Type the new name exactly as shown by the command in step 3. Now that you have changed the primary name, the */enumerate* command will not work again until you reboot.

5. Reboot and log on as a domain admin.
6. Check the names again:

Netdom Computername NewComputerNetBIOSname /enumerate

7. Delete the old computer name:

**Netdom Computername NewComputerNetBIOSname /remove:
OldComputerFQDN**

Type the old name exactly as shown by the */enumerate* command in step 6.

8. Confirm the functionality of the new name. Using the DNS administrator, expand the DNS server icon. Click the forward lookup zone and check for an (A) record for the new computer name. Verify that it points to the correct IP address. There should not be a host record for the old name. If you find one, delete it.



WARNING

Netdom doesn't change any DNS delegation records that are based on the old computer name. If you have them, they will need to be manually corrected or replication will fail. Specifically, check for delegation records in the *_msdcs* zone. Consult Microsoft Knowledge Base Article 321046 for instructions on how to use the *DNSLint* utility to troubleshoot active directory replication issues.

Implementing DNS in the Active Directory Network Environment

DNS is a service that takes the user-friendly names of the Internet and looks up their associated IP addresses. Active Directory uses DNS in the same way, looking up servers and services for users and applications. Because DNS is a hierarchical naming system, all Internet names exist within one of many available namespaces. One particularly familiar root namespace is the .com domain. The purpose of this hierarchy is to ensure that all names within a given namespace are unique. This way, when you connect to Microsoft.com, you know that you are connecting to the Microsoft.com owned by Microsoft, and not a Microsoft.com owned by some other company. Another function of DNS is to send you an IP address based on the service you want, rather than the machine name you want. Say that you send e-mail to an address in another company. Your mail server doesn't have to know the host name of the destination mail server to get its IP address. Instead, it asks DNS, "Who is the mail server for that domain?" DNS answers with an IP address, and your e-mail is sent.

Within your enterprise, DNS operates in much the same way; looking up IPs and helping people get access to the resources they need. On your network, there are many different servers and services. Fortunately, dynamic update allows your servers to update DNS tables on their own, advertising their services automatically as you install them and move them around. For example, when a user tries to log on, his or her station asks DNS, "Where is a logon server?" DNS looks up what is called a service location resource record (SRV) for logon servers and sends it back to the user's station, which connects to a DC for logon.

Whatever implementation of DNS you use with your Active Directory, it must support these SRV records. Using the Windows implementation of DNS gives you additional features that make it easier to use. For example, when you install a new domain, the DNS zones are automatically created and configured for you, significantly reducing the time you must spend manually configuring each DNS server. If you do have problems, the Windows Server 2003 DNS service has configuration enhancements that simplify the debugging and logging of incorrect DNS configurations. This helps you solve problems faster by suggesting troubleshooting steps for you to take. Another way in which Active Directory helps you configure DNS is through Group Policy for DNS clients. This greatly simplifies the implementation of DNS changes.

Windows 2000, Windows XP, and Windows Server 2003 are similar in their user interfaces for DNS, but Windows NT is significantly different. Table 4.8 shows a comparison of the administrative tools used in NT versus more recent versions.

Table 4.8 Comparison of Administrative Tools in Windows NT vs. Windows Server 2003

DNS Task	Windows NT 4.0 uses	Windows 2000, XP, and Server 2003 use
Installing the DNS server service	Network control panel	Windows Components wizard
Starting the DNS Manager	Start Administrative Tools DNS Manager	Start Administrative Tools DNS
Starting, stopping, or restarting the DNS service	Services control panel	Start Administrative Tools DNS, then right-click Computername All Tasks and select Start, Stop, Pause, Resume, or Restart
Adding a remote server to DNS Manager	Server menu in DNS Manager	In the DNS tool, right-click DNS Connect to DNS Server..., The following computer: type Computername, click OK

DNS and Active Directory Namespaces

Active Directory uses the DNS standard for naming objects. Like DNS, Active Directory is a hierarchical arrangement of objects within objects. Active Directory uses the same rules and procedures as DNS to resolve domain names, computer names, and service names. An enterprise with Internet connectivity can register its internal Active Directory forest-root domain as an Internet domain with one of the Internet naming registries. With an external namespace like this, the domains and trees within your forest become subdomains of an Internet root domain.

For security reasons, the confidentiality, integrity, and availability of your company's information assets is usually protected by a number of methods. One method is to implement an internal namespace that is hidden from the publicly accessible Internet. Technically speaking, you can use the same name on your internal and external namespaces while maintaining them as isolated entities. However, your users and system administrators will be easily confused by the parallel names, and conflicts can arise. Keeping your external namespace separate and distinct from your internal ones makes it easier to maintain firewall and router configurations, especially when using domain name filters and exclusion lists. If you publish your Active Directory namespace externally, then you supply potential hackers with a portion of your naming system. This is part of what they need to compromise your user accounts and computer services. In addition, internal namespaces are not constrained by worldwide Internet uniqueness. If you want, instead of limiting yourself to .com, .net, or .org, you could name your internal domain llove.golf or anything else you choose within the limitations of RFC 1123 and UTF-8 standards.



NOTE

In Active Directory, NetBIOS computer and service names are simply generated by using the first 15 characters of the DNS host name. That means the first 15 characters before the first dot in the fully qualified DNS hostname—anything after the first dot doesn't count. Active Directory has no dependence on NetBIOS names, although they still work to provide backward compatibility with pre-2000 versions of Windows.

In Windows Server 2003, Active Directory supports multiple discontinuous interforest namespaces through the implementation of GCs, and multiple discontinuous extraforest namespaces through the use of routing hints. This means that multiple namespaces can coexist within the same Active Directory. You should, however, be aware that *name collisions* are possible using cross-forest trusts if namespaces overlap within the federated forest.

DNS Zones and Active Directory Integration

Standard DNS zones are stored in text files in the `%systemroot/System32/Dns` folder. After DNS is integrated into the directory, it exists in a *dnsZone* container object identified by the name of the zone. DNS and Active Directory can use identical names for different namespaces. For this reason, it is important to understand that they are *not* the same namespace. DNS contains zones and records, while Active Directory contains domains and domain objects.

Windows Server 2003 brings in a new feature that blurs the line between DNS and Active Directory, called *application partitions*. With this release of Windows, DNS zones and records can be contained within the Active Directory itself, and are subject to the same replication and authentication parameters. Some advantages of this include a reduction in the number of objects stored in the GC, and a finely tuned replication domain. Regular Active Directory integrated DNS zones are replicated in their entirety to the domain partition of every DC whether it needs it or not and to the GC. By contrast, application partition-integrated DNS zones only replicate to DNS servers, reducing replication traffic and unused replicas. There are two automatically configured application partitions. One has a forestwide scope and resides on DCs running DNS in the forest root. The other application partition has a more limited domain scope, and resides on DCs running DNS in each domain.



EXAM WARNING

Be aware that each computer in the domain has representative DNS records and Active Directory objects. Objects, records, services, and attributes can use the same hostname, and refer to the same computer, but exist in different namespaces. This is true even for Active Directory-integrated and application partition-integrated DNS zones.

**NOTE**

Only primary DNS zones can be stored in the Active Directory. Secondary zones must be stored in the old standard text format. This might seem odd at first, but secondary DNS zones are essentially obsolete in light of the multi-master replication model of the Active Directory-integrated DNS zone. Secondary zones might still be needed if some zones will not be stored in the Active Directory, or will be maintained during the migration period.

Some of the benefits to be gained by integrating DNS with Active Directory include:

- An upgrade from the standard DNS single-master update model to the multi-master model. Updates can take place at any DNS server, not just the one that is authoritative for each zone. The multi-master model eliminates the single point of failure for dynamic updates.
- Access control lists (ACLs) on directory-integrated zones, allowing you to specify who can delete, modify, or even who can read records within the zone.
- Secure updates, which protect the integrity of your DNS zones by protecting against DNS poisoning and other malicious attacks.
- The automatic replication and synchronization of DNS zones whenever you install a new DC.
- A common replication topology for DNS zones and Active Directory domains. The nonintegrated DNS requires the design, implementation, testing, and administration of two different replication topologies.
- Added replication efficiency, since Active Directory replication is faster and more efficient than standard DNS replication.

Active Directory Integrated DNS Supports Multidomain DNS Conditional Forwarding

Multidomain DNS Conditional Forwarding sounds complex, but it couldn't be much easier. You learned in Windows 2000 that Active Directory needs a sturdy and secure DNS infrastructure. One of the secure aspects of DNS comes from a design called *split DNS* where you basically keep two sets of records, one for the outside of your network and one for the inside. The *Internet* DNS server holds the address information for your Web site, e-mail, and FTP servers. The *inside* server holds address information for your intranet and serves the needs of your Active Directory.

Continued

Split DNS works fine, except when joined with a very useful feature of Windows called Active Directory integrated zones. Active Directory integrated zones let you secure a zone for a DNS domain such as Dogs.com with one limitation: the DNS servers for Dogs.com must also be DCs for an Active Directory domain whose name is Dogs.com.

If you want to run more than one Active Directory domain in your intranet, then that presents a problem. Let's review some facts:

- Each Active Directory domain requires its own DNS zone.
- A DC can only serve one domain at a time.
- If you want to use Active Directory integrated zones, then you must have a separate set of DNS servers for each domain, since you must have a separate set of DCs.

That's where the problem lies. It is easy to keep two separate sets of records on just one DNS domain. You divide the world up into two simple areas: outside and inside. In this case, your inside users get full unfiltered DNS information about the outside world, but not the other way around. It isn't so easy with multiple domains.

Now let's add the second internal domain called Fish.com. To make Dogs.com users see the correct set of records, you must point all of their servers and workstations to the internal DNS servers that contain the internal-only version of the Dogs.com information. To support the users in Fish.com, you will set up a different DNS server for the internal-only version of Fish.com records, and point all those users to that server. When you use split-DNS, users in Dogs.com get the internal-only DNS information about their own domain and the public DNS information about the Fish.com domain. From the other side, users in Fish.com get their own internal DNS records, but only the public DNS information for any other domain.

Here's the problem: What if a Dogs.com user wants to access a resource on Fish.com? That means the user must locate a DC in Fish.com to do the lookup on the IP address of the resource in Fish.com.

Unfortunately, Dogs.com DNS servers only know the internal structure of their own domain, while Fish.com DNS servers only know the internal structure of Fish.com. The user tries to look up that information in the local DNS, but it doesn't know the answer.

There is more than one approach to solving this problem, but Windows Server 2003 offers a good one called *conditional DNS forwarding*. It lets you set up the Dogs.com DNS server so that whenever it needs to lookup something in a zone called Fish.com, it will know to go straight over to the Fish.com server for the answer, rather than out to the Internet root servers hosting the .com domain.

Configuring DNS Servers for Use with Active Directory

It is a good practice to add additional DNS servers to eliminate concerns over a single point of failure. Traditionally, you would install a standard secondary DNS server. With Active Directory integration, all you need to do is convert your standard primary DNS server to an Active Directory integrated primary DNS server. Once that is done, simply configure additional DCs to take on the DNS role for redundancy.

There are a few things to be acquainted with when you integrate your DNS into the directory. For one thing, there are no more secondary DNS servers. Once integrated into Active Directory, all DNS servers are primary. Zone transfers no longer take place; instead, Active Directory replication is used to distribute changes as they occur. With legacy DNS systems, Standard DNS zones are administered using a text editor. DNS zones stored in Active Directory are administered using the DNS console or the *dnscmd* command-line tool only—no more text editing.

The following sections take you through the Windows Server 2003 versions of common DNS administrative procedures related to Active Directory and application partition integration. By default, the DNS Server service will attempt to discover and build the standard DNS application directory partitions in Active Directory. Depending on how DNS was originally implemented, these default partitions might already exist. If necessary, you can manually create them as shown in the following steps. Some DNS procedures require the Windows support tool's *dnscmd* utility, which you can install by double-clicking **supportools.msi** on the Windows Server 2003 CD in the `\Support\Tools` folder.

Integrating an Existing Primary DNS Server with Active Directory

To integrate an existing primary DNS server with Active Directory, follow these steps:

1. On the current DNS server, click **Start** | **Programs** | **Administrative Tools** | **DNS** to start the DNS Administrator console.
2. Expand the server name.
3. Right-click your primary DNS zone, click **Properties**, click the **General** tab, and note the *Type* value. This will be *Primary zone*, *Secondary zone*, or *Stub zone*.
4. Click **Change**.
5. In the **Change Zone Type** box, click the check box for **Store the zone in Active Directory**.

NOTE

This check box is only available if your DNS server is also a DC.



6. Click **Yes** to verify, and then click **OK**.
7. In the Domain properties, the type should now read **Active Directory-Integrated**. You can add as many additional DNS servers as you want.
8. To force replication to occur immediately instead of waiting for the regular replication cycle, follow these steps:
 1. Click **Start | Administrative Tools | Active Directory Sites and Services**.
 2. Expand the sites. If no additional sites are configured, you will use the one called *default-first-site-name*.
 3. Expand the following folders: your *site*, Servers, *your Computer*, NTDS Settings. One or more DC objects are listed in the right pane. Right-click each entry to see its “friendly” name. Right-click an entry, and select Replicate Now to begin replication immediately. The time it takes to update the target controller depends on network performance and the amount of data replicated.

Creating the Default DNS Application Directory Partitions

To create the default DNS application directory partitions, follow these steps:

1. Log on to your DNS server as an Enterprise Administrator.
2. To open DNS, click **Start | Administrative Tools**, and double-click **DNS**.
3. In the console tree, expand and right-click the DNS server and select **Create Default Application Directory Partitions**. Follow the instructions to create the DNS application directory partitions. The options are:
 - **Would you like to create a single partition that stores DNS zone data and replicates that data to all DNS servers in the Active Directory domain *DnsDomainName*. Yes, or No.** This option creates one DNS application directory partition for each domain in the forest. DNS zones stored in this partition are replicated to all Active Directory-integrated DNS servers in the domain. Depending on your domain structure and the context of the command, you might get this question multiple times for different domains.
 - **Would you like to create a single partition that stores DNS zone data and replicates that data to all DNS servers in the Active Directory forest *DnsForestName*? Yes, or No.** This option creates one DNS partition named for your forest. It contains all the DNS servers running on the DCs in the forest, and replicates the DNS data to all DNS servers. DNS zones stored in this application directory partition are replicated to all Active Directory-integrated DNS servers in the forest.

**NOTE**

The option to create the default application directory partitions in the DNS console will not be available if they are already present in the Active Directory.

Using dnscmd to Administer Application Directory Partitions

There are some differences between standard DNS and the Active Directory-integrated version of DNS. For example, when you uninstall a DNS server hosting Active Directory-integrated zones, these zones will either be saved or deleted. Since the zone data is stored on other DNS servers, it will not be deleted unless the DNS server that you uninstall is the last one hosting that zone. Windows gives you a warning if this is the case.

Only Enterprise Admins can create a DNS application directory partition. Most other DNS tasks can be handled by the DnsAdmins or Domain Admins group.

1. Log on to your DNS server with the credentials needed for the given task.
2. Open a command prompt. Click **Start | All Programs | Accessories | Command Prompt**, or click **Start | Run** and type **cmd**.
3. Use the following dnscmd.exe options. See the example following for assistance:
 - Type the following command as an Enterprise Administrator to create a DNS application directory partition: **dnscmd *ServerName* /CreateDirectoryPartition *FQDN***.
 - Type the following command as a member of the DnsAdmins or DomainAdmins group to enlist a DNS server in a DNS application directory partition: **dnscmd *ServerName* /EnlistDirectoryPartition *FQDN***.
 - Type the following command as a member of the DnsAdmins or DomainAdmins group to n-enlist a DNS server in a DNS application directory partition: **dnscmd *ServerName* /UnenlistDirectoryPartition *FQDN***.

In this case, *ServerName* specifies the DNS host name or IP address of the DNS server, and *FQDN* specifies the name of the target DNS application directory partition.

**NOTE**

An FQDN has a trailing period to qualify its position at the root of the namespace, such as "Fish.com." shown in the following example. To see the full syntax of these commands, type:

```
dnscmd /CreateDirectoryPartition /?
dnscmd /EnlistDirectoryPartition /?
dnscmd /UnenlistDirectoryPartition /?
```

Here is an example. Note the “.” at the end of the FQDN:

```
D:\SupportTools>dnscmd DC4.Fish.com /CreateDirectoryPartition Fish.com.  
DNS Server DC4.Fish.com created directory partition: Fish.com.  
Command completed successfully.
```

Securing Your DNS Deployment

DNS is full of information. It helps users find services, and services find resources. Unfortunately, it sometimes provides malicious users with that same wealth of information about your network. To help keep this from happening, use the following guidelines as a minimum approach to your DNS security architecture.

- Use a split DNS design with internal DNS servers protected by your firewall and external DNS servers on the outside. Your internal namespace can be a child domain of your external namespace, or be completely different.
- Use your internal DNS servers to host your internal namespace and your external DNS servers to host your external namespace. The external servers should not be able to forward name lookups from the Internet to your internal network, but internal servers can forward queries to the outside.
- Use a packet-filtering firewall to lock down DNS port 53 so that only external DNS servers under your control can communicate with your internal DNS servers.
- Configure secure dynamic updates. With this setting, only computers joined to the Active Directory can authenticate with DNS; hence, register their service locator records. Computers that are unable to authenticate cannot make changes to DNS data.
- Carefully monitor and control who has the ability to control DNS zones through the DACL in Active Directory.

Summary of Exam Objectives

The logical structure of the network is defined by forests and domains, with domains organized into domain trees in which subdomains (called child domains) can be created under parent domains in a branching structure. Domains are logical units that hold users, groups, computers, and OUs (which in turn can contain users, groups, computers, and other OUs). Forests are collections of domain trees that have trust relationships with one another, but each domain tree has its own separate namespace. Aspects of the physical structure include sites, servers, roles, and links.

An Active Directory always begins with a forest root domain, which is automatically the first domain that you install. This root domain becomes the foundation for additional directory components. The domain is the starting point of Active Directory. It is the most basic component that can functionally host the directory. Simply put, Active Directory uses the domain as a container of computers, users, groups, and other object containers. Objects within the domain share a common directory database partition, replication boundaries and characteristics, security policies, and security relationships with other domains. The process of creating the forest and domain structure is centered on the use of the Active Directory Installation Wizard, which is also known as the *dcpromo* utility.

The first step is to install Windows Server 2003 as a member server or a stand-alone server. Next comes the decision process leading to the installation of a DC. Essentially, there are two reasons to install a DC: to create a new domain, or to add an additional DC to an existing domain. Depending on your current forest structure, you will end up with one of four results: A new forest, a new domain tree in an existing forest, a new child domain in an existing domain, or a new DC in an existing domain. Since a domain cannot exist without a DC, you must create at least one for each domain. The process of creating the first DC also creates the domain itself. Here are the four main reasons to create a new DC: creating the first domain in your network, creating a new domain in your forest, improving a domain's reliability, and improving network performance between sites. To improve a domain's reliability, you should always create at least a second DC in each domain. That way, if the first DC fails, you will still be able to use the second DC.

Application data partitions are another DC enhancement that allows for the creation of application-specific Active Directory partitions, also known as *naming contexts*. An instance of an application directory partition on another DC is called a *replica*. Site topologies and replication schedules are observed in this type of partition, and the application objects are not replicated to the GC. Conveniently, application partitions can leverage DNS for location and naming. Windows Server 2003 Web Edition cannot host application partitions because they do not support the DC role.

Before raising the functional level of a domain, all DCs must be upgraded to the minimum OS level. Remember that when you raise the domain functional level, it can never be changed back except for the case of Windows Server 2003 native. It can be temporarily reverted to interim level for the purpose of joining an NT domain during an upgrade directly to Windows Server 2003. New features and better security become available with

each subsequent higher level, but compatibility issues might be sufficient reason to delay the level-raising procedure.

In Windows NT 4.0, the domain had only one authoritative source for domain-related information, the primary domain controller, or PDC. The implementation of Active Directory brought the multi-master model, where objects and their properties could be modified on any DC and become authoritative through replication conflict resolution measures. The problem with the multi-master architecture is that some domain and enterprisewide operations are not well suited for it. The best design placed those functions on a single DC within the domain or forest, and Microsoft created the Active Directory Flexible Single-Master Operations (FSMO) roles. The Active Directory supports five operational master roles: the Schema Master, Domain Master, RID Master, PDC Emulator, and the Infrastructure Master. Two of these operate at the forest level only, the Schema Master and the Domain Naming Master. Conversely, the RID Master, PDC Emulator, and Infrastructure Master operate at the domain level. You can use the *ntdsutil.exe* command-line utility to transfer FSMO roles, or you can use an MMC snap-in tool. Depending on which role you want to transfer, you need to use one of the following three MMC snap-in tools: Active Directory Schema, Active Directory Domains and Trusts, or Active Directory Users and Computers. To seize a role, you must use the *ntdsutil* utility. If a computer cannot be contacted due to a hardware malfunction or long-term network failure, the role must be seized. After you seize a master role, the old DC that hosted it should never be brought back online. This is especially true of the Schema Master, Domain Naming Master, and RID Master roles. Rebuild the system as a member server and then promote it as a newly installed DC if needed.

DNS and Active Directory use identical names for different namespaces. For this reason, it is meaningful to understand that they are not the same namespace. DNS contains zones and records, while Active Directory contains domains and domain objects. Windows Server 2003 brings in a new feature that blurs the line between DNS and Active Directory, called *application partitions*. With this release of Windows, DNS zones and records can be contained within the Active Directory itself, and are subject to the same replication and authentication parameters. Some advantages of this include a reduction in the number of objects stored in the GC, and a finely tuned replication domain. DNS zones stored in Active Directory are administered using the DNS console or the *dnscmd* command-line tool only—no more text editing. When you deinstall a DNS server hosting Active Directory-integrated zones, these zones will be saved or deleted. Since the zone data is stored on other DNS servers, it will not be deleted unless the DNS server you uninstall is the last one hosting that zone. Windows gives you a warning if this is the case.

Because forests and domains are the foundation of the Windows Server 2003 Active Directory, this topic is an important one for all network administrators to master. The knowledge and skills imparted in this chapter are designed to help you understand and answer the questions on Exam 70-294, and to assist you in implementing and maintaining an Active Directory network on the job.

Exam Objectives Fast Track

Understanding Forest and Domain Functionality

- ☑ The new forest features of domain renaming and forest restructuring are only available after all DCs in the entire forest have been upgraded to Windows Server 2003 and the forest functional level has been raised to the Windows Server 2003 level.
- ☑ Cross-forest trusts are another new feature of Windows Server 2003 that when used, result in a configuration called *federated forests*. Without the forest trust, Windows-based Kerberos authentication between forests would not work.
- ☑ *Linked value replication* provides a solution to Windows 2000's limit of 5000 direct group members by replicating only those members that changed instead of replicating the entire group.
- ☑ Functional levels are a mechanism that Microsoft uses to remove backward compatibility within the Active Directory once it is no longer needed. In Windows Server 2003, there are two additional functional levels to consider in both domains and forests, as compared to Windows 2000 (where the functional levels were called "modes").

Creating the Forest and Domain Structure

- ☑ You should know what type of domain you want to install before you begin, and the namespace it will use.
- ☑ To improve a domain's reliability, you should always create at least two DCs in each domain.
- ☑ Application Director Partitions allow the fine tuning of the replication of DNS Active Directory Integrated zones.
- ☑ The Schema Master and Domain Naming Master roles are only present at the forest root.
- ☑ The first DC that you install in the forest is the root DC. It is responsible for the GC and for all five FSMO roles. Some roles can later be transferred to other DCs for performance and diversification.

Implementing DNS in the Active Directory Network Environment

- ☑ The domain is the primary boundary defining DNS and NetBIOS namespaces.
- ☑ NetBIOS computer and service names are generated by using the first 15 characters of the DNS host name. Active Directory has no dependence on NetBIOS names, but they continue to work, providing backward compatibility with pre-2000 versions of Windows.
- ☑ There are three basic configurations of Windows DNS: Standard DNS, Active Directory-integrated DNS, and application partition-integrated DNS.
- ☑ Active Directory-integrated DNS is an upgrade from the standard DNS single-master update model to the multi-master model, removing single points of failure.
- ☑ Every domain requires a DNS server. If one is not available during the installation of the first DC, Windows Server 2003 will install and configure one for you.

Exam Objectives Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

- Q:** What is the big deal about raising the functional levels of my domains and forests? Shouldn't I raise the levels as soon as they meet the prerequisites?
- A:** No. Remember that functional levels, once raised, cannot practically be lowered again. In addition, some situations are better suited to skipping a level, rather than raising to one level and then the other. In this case, known future restructuring and upgrade activities should be considered before raising functional levels.
- Q:** Now that cross-forest trusts are available, wouldn't it be easier to create a different forest for each division of the company instead of putting them in different domains within a single forest?
- A:** No. The enterprise approach to IT management includes gains in efficiency, functionality, data sharing, resource management, and business-unit collaboration, while reducing manpower and IT expenses. Localized ways of thinking are not competitive in a business sense, create inefficiencies, reduce functionality, and foster an environment

where workarounds and patches end up lowering information security in order to make things work in the enterprise. Cross-forest trusts can be an advantage in the right situations, but dividing a single enterprise into subunits is not one of them.

- Q:** The differences between functional levels of domains and forests are complicated and detailed. Couldn't I skip this section and focus on the larger issues and concepts when studying for the test, since I can always look up the information on the job?
- A:** No. Microsoft tests are scenario-based and rich with configuration details. Many test questions will have different answers based on the functional level context of the scenario. In other words, you won't completely understand the question if you don't understand the functional level involved. If you don't understand the question, you are likely to answer it incorrectly.
- Q:** How much of the Active Directory design stage should be complete before I install my first DC?
- A:** Primarily, the DNS design should be complete, and the decision should be made about how the forest-root domain will be used. Additional DCs and domains can be added later. FSMO roles and GCs can be shifted as needed, and trusts with other forests and external domains can be added later. Essentially, the first DC that you install should be in a lab environment. From that perspective, you should install your first DC for testing and training purposes as soon as possible.
- Q:** If my DNS is already integrated with the Active Directory, why would I go through the trouble of converting it to an application directory?
- A:** Application partition-integrated DNS zones only replicate to DNS servers, reducing replication traffic and increasing directory efficiency by removing unused replicas of the data. Additionally, the number of objects stored in the GC is reduced, and you no longer have to maintain the replication topology of DNS separate from the replication topology of the Active Directory.
- Q:** If every FSMO role can be seized by another DC upon failure, why would I want to spread the roles out between different machines?
- A:** There are several reasons. Chief among these are the associated risks of seizing roles. Lost or corrupted directory data can result from FSMO failures, especially if the malfunctioning machine ever comes back online. Seizing roles should not be considered a routine operation. Another consideration is performance. Each role exacts a certain amount of CPU and memory overhead, and your servers might perform better if roles are spread among multiple systems. If that weren't enough, some roles and functions should not coexist on the same DC, such as the Infrastructure Master and the GC. FSMO placement should not be ignored, and this knowledge will be important on the test.

Q: I read a lot about the concept of *namespaces*. Isn't this just an abstract concept that has no practical meaning to the real world?

A: No. Namespaces are well-defined boundaries that are critical for you to understand, whether you are in the planning, designing, implementing, or maintenance stage of your Active Directory deployment. Often, two namespaces will have the same or similar names. If a name or service resolution failure occurs, you must understand namespaces to intelligently troubleshoot the problem.

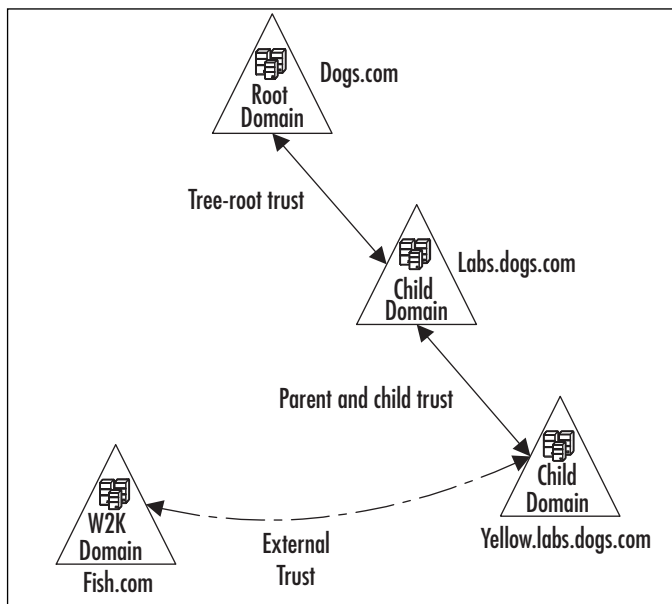
Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Understanding Forest and Domain Functionality

1. Your Yellow.labs.dogs.com Information Assurance department has just used John the Ripper, a password-cracking program, against your Active Directory user accounts. They report that 40 percent of your passwords were compromised within an eight-hour period. After some research, you determine that removing the LM authentication for down-level clients will make password guessing much more difficult. Further, you decide to require Kerberos authentication at all of your DCs. After some telephone calls, you assure your manager and CIO that all computers in all domains within your forest are loaded with Windows 2000, 2003, or XP. No legacy workstations are present that require backward LM or NTLM compatibility. With their permission, you apply the security setting using Group Policy. Almost immediately, authentication fails between Fish.com and the Yellow.labs.dogs.com domain, although ping still works. Previously, all users in both domains had complete access to all resources in the two domains, but now every authenticated access fails. Refer to Figure 4.56. Why is this problem happening?

Figure 4.56 Question #1 Diagram



- A. Network problems.
 - B. Selective authentication problems.
 - C. NT LAN Manager authentication problems.
 - D. Trust transitivity problems.
2. Over a weekend, you upgrade your Windows NT 4.0 domain through the Windows Server 2003 interim functional level up to the Windows Server 2003 functional level. Most client PCs are running Windows NT 4.0. Your test NT workstation continues to operate, but on Monday morning you start getting trouble tickets from several workstations that cannot log on to the new domain. As a temporary workaround, those users are able to log on from other workstations. What is the most likely cause?
 - A. The malfunctioning workstations are running Windows NT 4.0 SP2.
 - B. The users' old passwords do not meet the new complexity requirements of Windows Server 2003.
 - C. Some workstations did not pick up the new trust relationships, and need to be removed and re-added to the domain.
 - D. Some user accounts have not replicated to all DCs yet. Choose the **Synchronize all domain controllers** option from the **Active Directory Domains and Trusts** MMC snap-in.

3. A large company has just merged with yours. They have a corporate 24-hour Enterprise Administrators group running a forest at the Windows Server 2003 functional level. Your company has NT, 2000, and XP clients on a single NT domain. Your PDC and four of your five BDCs are still running Windows NT. Two of your permissions groups have more than 5000 members. The new conglomerate has given you a budget for upgrading the local infrastructure so that you can participate in the corporate forest structure through a cross-forest trust. What are the minimum requirements you must meet before joining the corporate forest? (Choose all that apply.)
 - A. Ask the Enterprise Administrators to temporarily set their forest to interim level using manual LDAP administrative tools.
 - B. Install new servers and create a pristine local forest root using your namespace.
 - C. Upgrade your PDC to Windows Server 2003, and join the corporate forest during the upgrade.
 - D. Upgrade your four Windows NT BDCs to Windows Server 2003.

4. Two forests are joined with a cross-forest trust, linking the two for resource sharing. Both forests synchronize their time with the Internet, they both contain some Windows NT 4.0 workstations, and they reside in different time zones. One day, the cross-forest trust fails, and users cannot authenticate across their forest boundaries, although the network is healthy at layer 3. What might have happened?
 - A. The two forests are using different time sources, or one forest-root PDC emulator has malfunctioned.
 - B. The Windows NT systems need to reboot, or use the scheduler service to run the `net time \\<timecomputer> /set /yes` command on each NT machine.
 - C. DCs in both forests should be set up to automatically run the `w32tm /config /syncfromflags:manual /manualpeerlist:Peerlist` command, listing only each other as their time source.
 - D. The PDC Emulator needs its time reset, since Windows NT 4.0 workstations and servers automatically synchronize with their PDC.
 - E. The scenario is impossible, since Windows NT cannot function in a Windows Server 2003 functional level domain, and cross-forest trusts can only be created in such a domain.

5. Your IT manager wants you to link four divisions of the company through a ring of eight unidirectional cross-forest trusts. He uses this reasoning: If multiple forest trusts are established, then authentication requests made in any domain of any forest can pass through multiple forest trusts, hence multiple Kerberos domains, on their way to their destination. Why is he wrong?

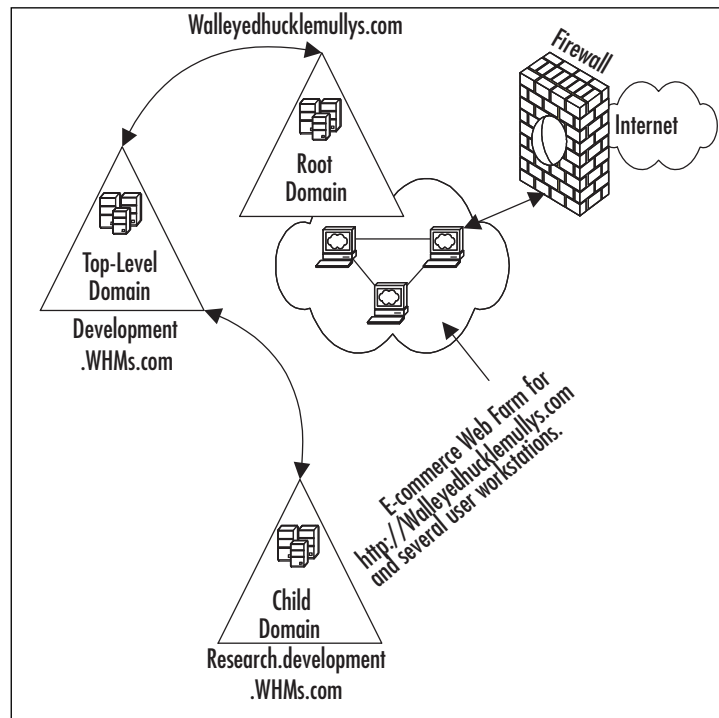
- A. While each cross-forest trust is transitive at the forest level, where all domains in both forests can authenticate, they are not transitive at the federated forest level as he suggests. The trust path cannot include more than one cross-forest trust.
 - B. Cross-forest trusts are not transitive, and will not allow pass-through authentication.
 - C. To create a ring of trust around four forests, you only need four cross-forest trusts.
 - D. Cross-forest trusts are bidirectional, so only three trusts are needed to link all four forests. Completing the “ring” is not necessary.
6. You are the database administrator for a large enterprise, and you are trying to install an Active Directory-integrated application. The product literature says it implements new classes within the directory. The installation fails with an obscure message. What is the most likely reason for the failure?
- A. The application uses a service account that is not a member of the Enterprise Administrators group.
 - B. You are not a member of the Domain Admins account in the domain where you are trying to install the application.
 - C. You are not in the Schema Administrators group.
 - D. You do not have permissions to create GPOs.
7. What FSMO roles should exist in a child domain in a Windows Server 2003 forest? (Choose all that apply.)
- A. Schema Master
 - B. Domain Naming Master
 - C. PDC Emulator
 - D. RID Master
 - E. GC
 - F. Infrastructure Master

Creating the Forest and Domain Structure

8. Your company is expanding its single location with five new offices, all in different states. Each location will have 10 marketing employees or less to supplement the 25 already employed at the main office. Your security engineer says that all employees will abide by the same company security policy. They will hire another systems administrator at your office to handle the increased workload. When asked about how the company’s single-domain Windows Server 2003 Active Directory will be affected by the expansion, you reply that new servers will have to be installed at the remote locations. Your manager wants to know what server hardware and software to budget for. What do you tell him? (Choose one.)


- A. Five servers and five copies of Windows Server 2003 Datacenter Edition
- B. Five servers and five copies of Windows Server 2003 Standard
- C. Ten servers and ten copies of Windows Server 2003 Enterprise Edition
- D. Ten servers and ten copies of Windows Server 2003 Standard
9. You have just been hired to install an Active Directory for a new startup company. An e-commerce Web farm has already been set up, waiting to join your root domain. Walleyedhuckle Mullys.com employs seven chemists who appear to split their time evenly between research and development activities, two Web developers, one network engineer, and the owner of the company. The previous MCSE who ordered the computers and software has already left the company, and all you have is a pile of boxes and a crude sketch of a logical Active Directory forest. The owner of the company has already approved the design and wants you to work all weekend to make it work. After you clean up the drawing somewhat, it looks like Figure 4.57. Among the boxes you find six computers, three copies of Windows Server 2003 Standard, and three copies of Windows Server 2003 Web Edition. Under the pressure of time, you create the forest structure with two computers in each of the three domains. On Monday morning, you approach your new manager to say that the Active Directory design is weak. What is wrong with it? (Choose all that apply.)

Figure 4.57 Question #9 Diagram



- A. DNS should be split, and the design is full of bad namespace and computer security practices.
 - B. Too many domains.
 - C. You should have installed Windows Server 2003 Enterprise instead of Standard.
 - D. Firewall is in the wrong place, because Web servers should be directly on the Internet.
 - E. Not enough DCs for the number of domains.
10. The name of your company has changed from Fish to Aquatics. In her zeal to embrace the new company image, your CEO directs you to change the Active Directory domain name to reflect the new company name. Your efforts are unsuccessful. Which of the following would cause the domain rename tool to fail? (Choose all that apply.)
- A. Your DCs are running Windows Server 2000 Enterprise Edition.
 - B. The Active Directory forest functional level is set to Windows Server 2003 interim.
 - C. You are logged in as a Domain Admin.
 - D. Your current mail server is running Exchange 2000.
 - E. Your DFS root server is running Windows 2000 with Service Pack 3.
 - F. You are logged on to a Windows 2003 Server DC as an Enterprise Admin.
11. Your network operations center has identified excessive bandwidth utilization caused by authentication traffic in the root domain subnet, especially between Calico.cats.com and Labs.dogs.com. Your logical network is set up as shown in Figure 4.58. What type of trust or trusts would you set up to alleviate the situation?

Figure 4.58 Question #11 Diagram



- A. Set up a bidirectional transitive parent and child trust between Calico.cats.com and Labs.dogs.com.
 - B. Set up a shortcut trust between Calico.cats.com and the forest root, and set up a second shortcut trust between Labs.dogs.com and the forest root.
 - C. Set up a shortcut trust between Calico.cats.com and Labs.dogs.com.
 - D. Set up two shortcut trusts between Calico.cats.com and Labs.dogs.com.
 - E. Set up a realm trust between Calico.cats.com and Labs.dogs.com.
12. Which of the following factors should have the most influence on the number of domains that you install? (Choose all that apply.)
- A. The need for different security policies for different groups of users.
 - B. The need for different schemas for each user group.
 - C. The possibility that one company division might eventually be sold.
 - D. The separation of IT management between different major physical sites.
 - E. The structure of your organizational chart.
13. Your company has just elected a new chairperson who believes in a sound information infrastructure. Currently, your network is a product of years of least-cost and least-effort development. Having the opportunity to build a new enterprise, you sit down and begin your design. What is the most critical design question to answer before beginning the installation of your new Windows Server 2003 forest?
- A. The name of the root-level domain.
 - B. The number of domains that you need to support the enterprise.
 - C. The locations of the operational master roles and GCs.
 - D. Whether to use application directory partitions for your DNS servers.

Implementing DNS in the Active Directory Network Environment

14. Your DNS expert is not an MCSE, and has mostly UNIX experience. Your enterprise runs on Windows Server 2003 Active Directory application partition enabled DNS servers. You want your DNS expert to have access to the DNS servers, but you do not want to grant the ability to configure other DCs.

To which group should you add your expert?

- A. Enterprise Admins
 - B. DNSAdmins
 - C. Administrators
 - D. Domain Admins
15. A help desk ticket is routed to you with the diagnosis that your DNS server is failing to renew the lease on a user's IP address. After reviewing the roles that DNS plays in the Active Directory, you think you know where the problem lies. What are the roles of DNS? (Choose all that apply.)
- A. Locating DCs
 - B. Assigning IP addresses
 - C. Defining the Active Directory namespace
 - D. Name resolution

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|-------------------|--------------------------|
| 1. C | 9. A, B, E |
| 2. A | 10. A, B, C, D, F |
| 3. A, C | 11. D |
| 4. A | 12. A, D |
| 5. A | 13. A |
| 6. C | 14. B |
| 7. C, D, F | 15. A, C, D |
| 8. B | |

MCSA/MCSE 70-294

Working with Trusts and Organizational Units

Exam Objectives in this Chapter:

- 1.3.6 Establish trust relationships. Types of trust relationships might include external trusts, shortcut trusts, and cross-forest trusts.
- 2.1.1 Manage trust relationships.
- 3.3.1 Analyze the administrative requirements for an OU.
- 3.4.3 Move objects within an OU hierarchy.
- 3.4 Implement an OU structure.
 - 3.4.1 Create an OU.
 - 3.4.2 Delegate permissions for an OU to a user or to a security group.
- 1.5 Plan an administrative delegation strategy.
 - 1.5.1 Plan an organizational unit (OU) structure based on delegation requirements.
 - 1.5.2 Plan a security group hierarchy based on delegation requirements.
- 3.3 Plan an OU structure.
 - 3.3.2 Analyze the Group Policy requirements for an OU structure.

Introduction

Trust relationships define the ways in which users can access network resources across domains and forests. Without a trust between the domain to which a user belongs and the domain in which a resource resides, the user won't be able to access that file, folder, printer, or other resource. Hence, it is important for network administrators to understand how the built-in (implicit) trusts in the Active Directory network function, and how to create explicit trusts to provide access (or faster access) between domains.

Organizational units (OUs) are container objects within the directory structure that can be used, as the name implies, to organize resources, including (but not limited to) users, groups, and computers. Group policies can be applied to OUs, and administration of an OU can be delegated, making it easy to perform tasks that need to apply to only select objects.

This chapter addresses these two important components of Active Directory: trust relationships and OUs. You'll learn about the different types of trusts that exist in the Active Directory environment, both implicit and explicit, and you'll learn to create shortcut, external, realm, and cross-forest trusts. You'll also learn to verify and remove trusts, and how to secure trusts using SID filtering.

Next, we discuss the creation and management of OUs and you learn to apply group policy to OUs and how to delegate control of an OU. We show you how to plan an OU structure and strategy for your organization, considering delegation requirements and the security group hierarchy.

EXAM
70-294

OBJECTIVE
1.3.6
2.1.1

Working with Active Directory Trusts

One of the many issues that need to be dealt with in any computer organization is how to protect resources. The main difficulty that administrators face is the dilemma of how to ensure that the resources of the company are not accessible by those who do not need access. The other side of that coin, equally important, is how to ensure that people who do need access are granted access with the least amount of hassle. In small companies, the issues are simpler, because multiple domains rarely exist. In today's larger corporations and conglomerates, the issues of security are compounded. What administrators need is an easy tool to manage access across multiple domains and, often, across forests.

The tool is Active Directory Domains and Trusts. With Active Directory Domains and Trusts, an administrator can establish relationships between domains that will allow users in one domain to access the resources in another. This way, the administrator can ensure that all users who need access can have it without the hassles involved in having user accounts in multiple domains.

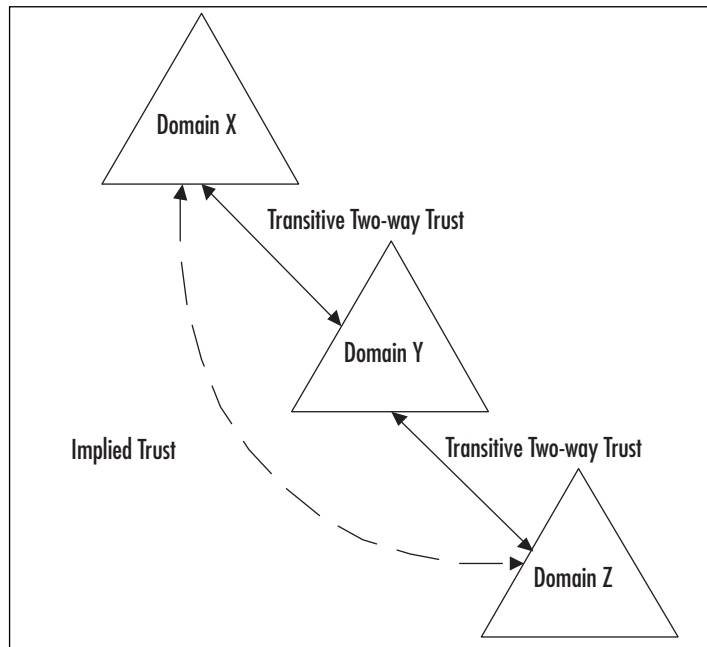
Several terms need to be defined in order to understand how trusts work. First, you need to understand the differences between two types of trusts:

- **Transitive trusts** The trust relationship passes through each trusted domain so that if A trusts B and B trusts C, A trusts C.

- **Non-transitive trusts** The trust relationship stops with the two domains between which it is created.


In addition to being transitive or non-transitive, a trust can be either one-way or two-way. A *transitive two-way* trust allows all domains to share resources to all users regardless of to which domain they belong. For example, you create a trust between two domains (Domain X and Domain Y). User accounts in Domain X have access to resources in Domain Y. The reverse is also true, user accounts in Domain Y have access to domain X's resources. This is a two-way trust; the trust works in both directions. Now let's add a new Domain Z and create a trust relationship to Domain Y. The transitive trust allows the user accounts in Domain X to access resources in Domain Z and vice versa without having to create an additional trust between Domain X and Domain Z (see Figure 5.1). This reduces the numbers of trusts that an administrator needs to create and maintain.

Figure 5.1 Transitive Trust



A *non-transitive* trust restricts the relationship to domains outside of the trust. It does not inherently allow other domains to pass through their authentication information to access resources outside of the trust. Once again, let's use the example of Domain X and Domain Y. If a non-transitive trust relationship is established between Domain X and Domain Y, the user accounts in the two domains have access to resources in the other domain. If we now add a new Domain Z and create a trust between Domain Y and Domain Z, users in Domain X are not automatically allowed access to resources in Domain Z (see Figure 5.2).


Figure 5.2 Non-Transitive Trust



We've been talking about two-way (bidirectional) trusts; but a trust can also be one-way (unidirectional). One-way trusts are created to allow more restrictive control over which users are allowed access to resources. For example, in Figure 5.3, a one-way trust is created between Domain X and Domain Y. Users in Domain X have access to resources in Domain Y. However, users in Domain Y do not have access to resources in Domain X. In this definition, Domain X is referred to as the trusted domain, and Domain Y is the trusting domain. A two-way trust allows users in either domain to have access to resources in the other domain.

One-way trusts must specify the *direction* of the trust. One-way trusts can be either *incoming* or *outgoing*, depending on whether the trust is created from the trusting or the trusted domain. Incoming trusts permit the users in the domain where the trust is created (the *trusted* domain) to access resources in the specified domain (the *trusting* domain). Users in the trusting domain do not have access, through this trust, to the resources in the trusted domain. (You can, however, create a second trust that goes the other way, to accomplish the same effect as a two-way trust.)

Outgoing trusts allow the users in the specified domain (the trusted domain) to have access to resources in the originating domain (the trusting domain). Users in the originating domain do not have access to resources in the specified domain.

Figure 5.3 One-Way Trust

Another concept and set of terms to understand in regard to trusts is:


- Implicit
- Explicit

Implicit trusts are trusts that are created automatically by the nature of the built-in relationships between domains within a forest. These implicit trusts are two-way and transitive. Implicit trusts automatically exist between each domain that is created and its child domain(s). An implicit trust also exists between the root domain of each domain tree and the root domains of every other domain tree in the forest.

An *explicit* trust is one that is created by an administrator; it does not exist automatically, but has to be explicitly created. For example, an administrator can create an explicit trust (in this case, called a *shortcut trust*) between any two child domains in different domain trees to provide for a direct trust (and faster authentication) between them.

Explicit trusts are also used to enable authenticate across forests. When a forest trust is created, a transitive trust is created between the forest root domains in both forests. This allows all the members in the forest to exchange authentication information with the other forest. The forest trust is also called an explicit trust between the two forests. If an additional forest trust is created between one of the original forests and a third forest, an implicit trust with the other original forest is not established to the third forest. In order for the third forest to have a trust relationship with the other forest, an explicit forest trust must be created between the two (see Figure 5.4).

Figure 5.4 Implicit Trust



Trust Types

Although the concept behind trust relationships is not a new item in the Windows NT family, the new trust types and flexibility offered is new to Windows Server 2003 domains. The Active Directory Domains and Trusts tool in the Administrative Tools menu gives you the necessary tools to allow users in Active Directory domains to easily gain access to resources in other domains, even if the other domain is a non-Windows Kerberos domain.

As with any new features, especially those related to security and access, you can expect several new questions focused on the new terms and concepts. You will be required to have a thorough understanding of the concepts that follow.

Types of Trust Relationships

Two or more Active Directory domains are implicitly or explicitly connected using trust relationships. The authentication requests made from one domain to the other domains use these relationships. The trusts provide a seamless coexistence of resources within the forest structure. Users are granted access to the resources in the other domain(s) after being

authenticated in their own domain first. Once authenticated in their own domain, they can traverse the other domains to gain access to their resources.



TEST DAY TIP

On the day of the test you will want to review the types of trusts as well as when to use each of the different trusts. On the exam, you might be given a scenario that will require you to determine the type of trust that will best meet the requirements in the scenario.

Because there are several new types of trust, you should ensure that you know when it is appropriate to use the different trusts.

The primary advantage of these relationships is that administrators no longer need to create multiple user accounts for each user who needs access to resources within each domain. Administrators can now add the users of the other domains to their access control lists (ACLs) to control access to a resource. To take full advantage of these relationships, the administrator must know about the various types of trust that exist, and when to use them.

Default Trusts

When the Active Directory Installation Wizard is used to create a new domain within an existing forest, two default trusts are created: a parent and child trust, and the tree-root trust. Four additional types of trusts can be created using the **New Trust Wizard** or the command-line utility **netdom**. The default trust relationships inside a Windows 2000 and Windows Server 2003 forest are transitive, two-way trusts.

A parent and child trust is a transitive, two-way trust relationship. It allows authentication requests made in the child domain to be validated in the parent domain. Because the trusts are transitive, these requests pass upwards from child to parent until they reach the root of the domain namespace. This relationship will allow any user in the domain to have access to any resource in the domain if the user has the proper permissions granted.


An additional transitive, two-way trust is created to simplify the navigation, the tree-root trust. This is especially needed in large organizations that might have multiple levels of child domains. The tree-root trust is a trust that is created between any child domain and the root domain. This provides a shortcut to the root. This trust relationship is also automatically created when a new domain tree is created.

Shortcut Trust

Shortcut trusts are transitive in nature and can either be one-way or two-way. These are explicit trusts that you create when the need exists to optimize (“shortcut”) the authentication process. Without shortcut trusts in place, authentication travels up and down the domain tree using the default parent and child trusts, or by using the tree-root trusts. In large complex organizations that use multiple trees, this path can become a bottleneck

when authenticating users. To optimize access, the network administrator can create an explicit shortcut trust directly to the target domain (see Figure 5.5).

Figure 5.5 Shortcut Trust



These trusts are used when user accounts in one domain need regular access to the resources in another domain. Shortcut trusts can be either one- or two-way.

One way shortcut trusts should be established when the users in one domain need access to resources in the other domain, but those in the second domain do not need access to resources in the first domain.

Two-way trusts should be created when the users in both domains need access to the resources in the other domain. The shortcut trust will effectively shorten the authentication path, especially if the domains belong to two separate trees in the forest.

Realm Trust

Realm trusts are explicit trusts that are created to join a Windows Server 2003 domain to a non-Windows Kerberos v5 realm. This allows you the flexibility of creating a trust for your non-Windows networks to interoperate with the security services based on other Kerberos v5 implementations, such as with UNIX. This extension of security can be switched from one-way or two-way trusts and from transitive to non-transitive.

External Trust

An external trust is used when you need to create a trust between domains outside of your forest. These trusts can be one- or two-way trusts. They are always non-transitive in nature. This means that you have created an explicit trust between the two domains, and domains outside this trust are not affected. You can create an external trust to access resources in a domain in a different forest that is not already covered by a forest trust (see Figure 5.6).

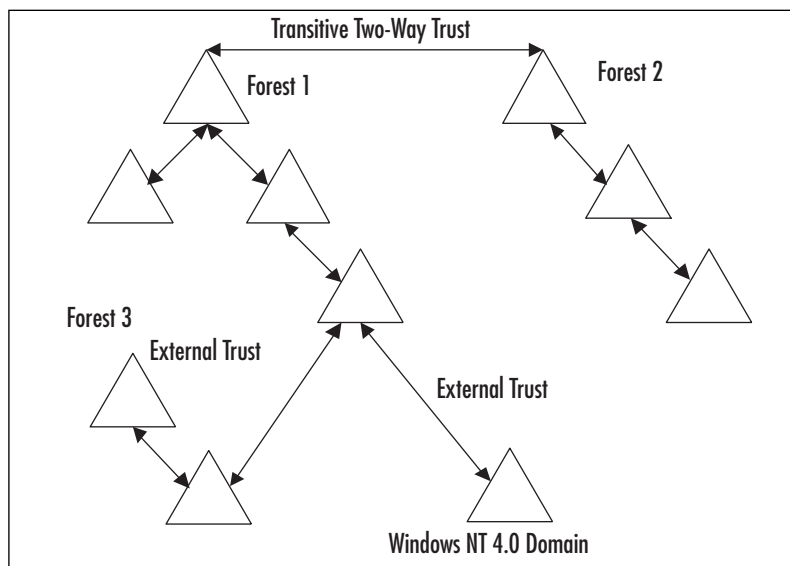


EXAM WARNING

You will always need to create an external trust when connecting to a Windows NT 4.0 or earlier domain. These domains are not eligible to participate in Active Directory. These trusts must be one-way trusts. If you have worked with Windows

NT 4.0, you will remember that the only trusts allowed were non-transitive one-way trusts.

Figure 5.6 External Trust



After the trust has been established between a domain in a forest and a domain outside the forest, the security principals from the domain outside the forests will be able to access the resources in the domain inside the forest. Security principals can be the users, groups, computers, or services from the external domain. They are account holders that are each assigned a security identifier (SID) automatically to control access to the resources in the domain.

The Active Directory in the domain inside the forest will then create foreign security principal objects representing each security principal from the trusted external domain. You can use these foreign security principals in the domain local groups. This means that the domain local groups can have members from the trusted external domain. You use these groups to control access to the resources of the domain.

The foreign security principals are seen in Active Directory Users and Computers. Since the Active Directory automatically creates them, you should not attempt to modify them.

Forest Trust

A forest trust can only be created between the root domains in two forests. Both forests must be Windows Server 2003 forests. These trusts can be one- or two-way trusts. They are considered transitive trusts because the child domains inside the forest can authenticate themselves across the forest to access resources in the other forest.



EXAM WARNING

Although the trust relationship is considered transitive, this only applies to the child domains within forests. The transitive nature of the trust exists only within the two forests explicitly joined by a forest trust. The transitivity does not extend to a third forest unless you create another explicit trust (see Figure 5.4).

Forest trusts help manage the Active Directory infrastructure. They do this by simplifying the management of resources between two forests by reducing the required number of external trusts. Instead of needing multiple external trusts, a two-way forest trust between the two root domains will allow full access between all the affected domains. Additionally, the administrator can take advantage of both the Kerberos and NTLM authentication protocols to transfer authorization data between forests.

Forest trusts can provide complete two-way trusts with every domain within the two forests. This is useful if you have created multiple forests to secure data within the forest or to help isolate directory replication within each forest.

Creating, Verifying, and Removing Trusts

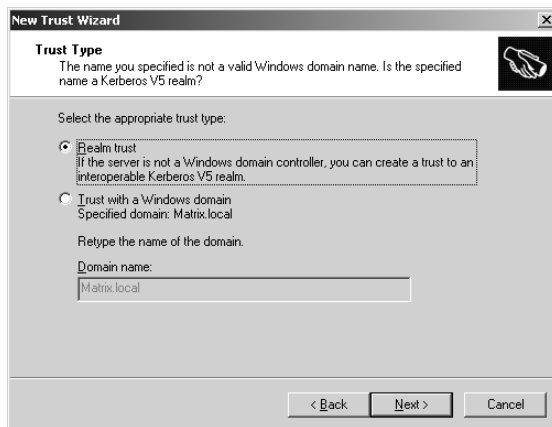
Trust relationships are created and managed using the **Active Directory Domains and Trusts** utility in the **Administrative Tools** menu. To create or manage trusts, you must be a member of the Domain Admins group or the Enterprise Admins group in the Active Directory, or have the appropriate authority delegated to you.

Most administrators will use the **RunAs** command to manage trusts. This is generally accepted as a security best practice.

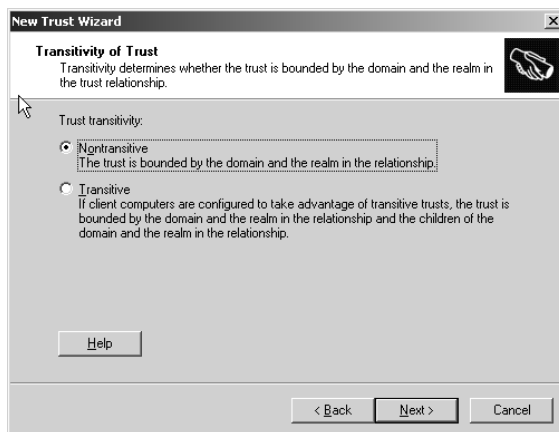
EXERCISE 5.01

CREATING A TRANSITIVE, ONE-WAY INCOMING REALM TRUST

1. Open **Active Directory Domains and Trusts** by clicking **Start | Programs | Administrative Tools**, and then selecting **Active Directory Domains and Trusts**.
2. In the console tree, right-click the domain node. Select **Properties** in the context menu.
3. On the **Trusts** tab, click the **New Trust** button.
4. When the **New Trust Wizard** opens, click **Next**.
5. On the **Trust Name** page, enter the target realm's name and click **Next**.
6. On the **Trust Type** page, select **Realm Trust** and click **Next** (see Figure 5.7).

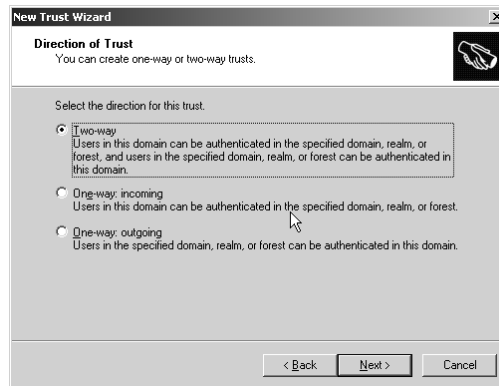
Figure 5.7 Trust Types

7. On the **Transitivity of the Trust** page, click **Transitive**, and then click **Next** (see Figure 5.8).

Figure 5.8 Transitivity of Trust

8. On the **Direction of Trust** page, click **One-way: incoming**, and then click **Next** (see Figure 5.9).

Figure 5.9 Direction of Trust



9. On the **Summary** page, review the information, and then click **Finish**.

This wizard will allow you to also create non-transitive trusts and two-way and one-way outgoing realm trusts. Alternatively, you can use the **netdom** command to create a realm trust.

Securing Trusts Using SID Filtering

One security concern when using trusts is a malicious user who has administrative credentials in the trusted domain sniffing the trusting domain to obtain the credentials of an administrator account. With the credentials of the trusting domain administrator, the malicious administrator could add a spoofed SID to allow full access to the trusting domain's resources. This type of threat is called an *elevation of privilege attack*.

The security mechanism used by Windows Server 2003 to counter an elevation of privilege attack is *SID filtering*. SID filtering is used to verify that an authentication request coming in from the trusted domain only contains the domain SIDs of the trusted domain. It does this by using the SIDHistory attribute on a security principal.



NOTE

Security principal is a term used to describe any account that has a SID automatically assigned. Examples of security principals are users, groups, services, or computers. Part of each security principal is the domain SID to identify the domain in which the account was created.

SID filtering uses the domain SID to verify each security principal. If a security principal includes a domain SID other than one from trusted domains, the SID filtering process

removes the SID in question. This is done to protect the integrity of the trusting domain. This will prevent the malicious user from being able to elevate his or her own privileges or those of other users.

There are some potential problems associated with SID filtering. It is possible for a user whose SID contains SID information from a domain that is not trusted to be denied access to the resources in the trusting domain. This can be a problem when universal groups are used. Universal groups should be verified to contain only users that belong to the trusted domain.

SID filtering can be disabled if there is a high level of trust for all administrators in the affected domains, there are strict requirements to verify all universal group memberships, and any migrated users have their SIDHistories preserved. To disable SID filtering, use the **netdom** command.

EXAM
70-294

OBJECTIVE
3.3.1
3.4.3

Working with Organizational Units

An OU is a frequently misunderstood unit inside Active Directory. An OU is simply a container in Active Directory that can contain any of the following:

- Users
- Groups of users
- Printers
- Shared folders
- Computers
- Other OUs

OUs are security *boundaries*. Many people confuse an OU with a security group. The key to understanding the difference begins with understanding what each is designed to accomplish. A security group is used to control access to a resource. In other words, *permissions* to a file, folder, or printer are set on security groups to manage the users as they try to access the object. *Policies* are set on OUs, to control what users can see and do on their computers, such as changing their wallpaper, accessing the Control Panel, what applications they can use, or whether they can shut down the computer. Policies can control either users or computers.

The means by which policies are applied is the Group Policy Object (GPO). GPOs can be applied to domains, sites, and OUs. The smallest scope that can have a GPO applied to it is the OU. It is also the last unit to which the GPO is applied, as it is the closest unit to the security object (the user account or computer account). The GPO gives us the ultimate in control over users and computers.

Creating OUs inside a domain allows for two different types of hierarchies. One hierarchy is the structure of the domain and child domains; the other hierarchy is the structure of the OU and its child OUs. The two hierarchies give you flexibility in how to manage

the organization. The concept of placing one OU inside another is called *nesting*. Although there are no limits to the number of nested OUs, Microsoft recommends that you not exceed 10 levels of nesting.

Understanding the Role of Container Objects

OUs are not security principals. Security principals are user accounts, group accounts, and computer accounts. OUs are containers that are used to organize the Active Directory. The purpose of creating OUs is to allow the administrator to create a container that can be used to implement security policies, run scripts, deploy applications, and delegate authority for granular administrative control.



TEST DAY TIP

Domain local groups and global groups are used to manage users and control access to resources. OUs are created to reflect the organizational structure, to manage security policies, and to delegate authority.

EXAM
70-294
OBJECTIVE
3.4
3.4.1

Creating and Managing Organizational Units

OUs are created and managed in the **Active Directory Users and Computers** tool in the **Administrative Tools**. This tool allows you to add OUs to the domain. After adding an OU, you have the ability to delegate control, add members, and move the OU. All of these activities can be accomplished by right-clicking on the OU that you want to manage and selecting the appropriate action from the context menu. The context menu will give you options to delete, rename, and enter the properties of the OU as well.

Configuring and Implementing Organizational Units

Active Directory provides us the potential to create multiple structures to represent our organizations. The goal is to provide enough flexibility to accomplish our tasks in a fashion that makes the most sense. The physical layout of the network is encompassed using sites, subnets, and site links. The logical layout of the Active Directory uses OUs. Remember to always note the differences between the physical and the logical layouts when designing the structure.

Often, administrators create the OU structure based on the departmental hierarchy. This is not always the best practice. You should create OUs for applying GPOs, hiding Active Directory objects or part of the Active Directory tree from a part of the organization, or for the delegation of authority. Remember, if you have not clearly defined the purpose of the OU, then you probably shouldn't create it.

Continued

For each OU that you plan to create, you should be able to document its purpose, a list of the users who will have control over it, and how much and the type of control they will have.

The Properties window of an OU has three tabs:

- General
- Managed By
- Group Policy

The **General** tab allows you to enter a description of the OU, street, city, state/providence, zip/postal code, and country/region information.

The **Managed By** tab allows you to change the user account that manages the OU. When a user account has been selected, the tab will display information about the account, such as office, address, and telephone numbers. This is read for the corresponding section of the user information stored about that user account. The **Managed By** tab has three buttons to manage this section of the OU Properties: **Change**, **View**, and **Clear**. The **Change** button opens a user window so you can select the account that will be used to manage the OU. The **View** button lets you see the user's account Properties window. You have the opportunity of making any necessary changes to the user's account. The **Clear** button removes the user account from the Managed By tab.

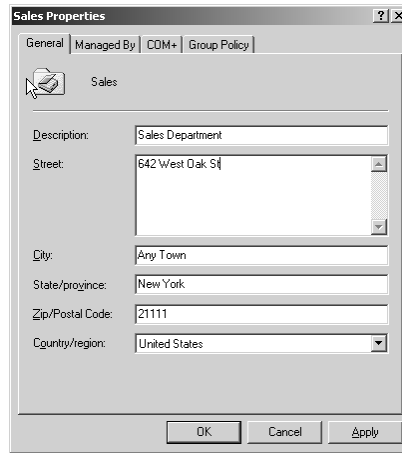
The **Group Policy** tab is discussed in the next section.

EXERCISE 5.02

CREATING AN ORGANIZATIONAL UNIT

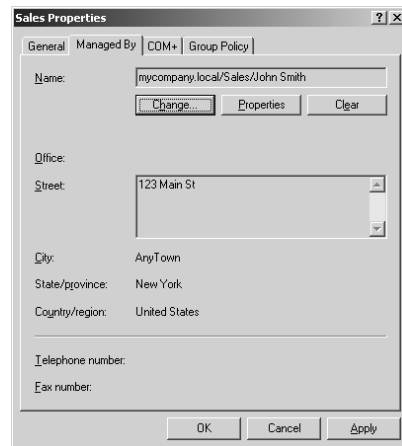
1. Open **Active Directory Users and Groups** by clicking **Start | Control Panel | Performance and Maintenance | Administrative Tools**, and then double-click **Active Directory Users and Groups**.
2. In the console tree, right-click on the domain node. Select **New** in the context menu, and then select **Organizational Unit**.
3. In the **New Organizational Unit** window, type the name of the OU.
4. Click **OK** to create the OU.
5. Right-click on the new OU and select **Properties** from the menu.
6. On the **General** tab, enter a description to explain the purpose of the OU (see Figure 5.10).

Figure 5.10 OU Properties



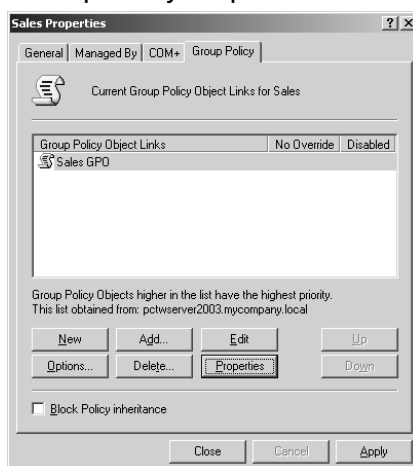
7. Click on the **Managed By** tab (see Figure 5.11). Click the **Change** button. Select a user account to manage the OU from the **Users and Groups** window.

Figure 5.11 Managed By Properties




8. Click the **Group Policy** tab (see Figure 5.12). Click the **New** button to create a new GPO.

Figure 5.12 Group Policy Properties



9. Rename the GPO by typing the new name.
10. Right-click on the GPO, and select **No Override** from the menu. Notice the check mark by the GPO in the **No Override** column (see Figure 5.13).

Figure 5.13 No Override Option



11. Click the **Edit** button. From the GPO window, double-click **User Configuration | Administrative Templates | Start Menu & Taskbar**. Double-click **Remove Favorites menu from Start Menu**. In the window that opens, click **Enable** to remove favorites from the Start menu.

12. Click the **Explain** tab. This defines what the impact of your actions will be.
 13. Click **OK** to close the **Remove Favorites menu from Start Menu** window.
 14. Close the GPO window.
 15. Check the **Block Inheritance** option.
 16. Right-click on the GPO and select **Disable** from the menu.
 17. Close all windows.
-

Applying Group Policy to OUs

One of the fundamental reasons for creating an OU is to apply a GPO to it. After creating the OU, you can then create a new GPO or apply an existing GPO. The **Group Policy** tab found in the **OU Properties** window is the most important tab of the OU properties. This is where you create, associate, and edit the GPOs that will affect the OU. This tab has the following buttons:

- New
- Add
- Edit
- Options
- Delete
- Properties

The **New** button will create a new GPO. When it is clicked, you need to supply the name for the new GPO.

After the GPO is created, use the **Edit** button to edit its configuration settings. The GPO is broken into two sets of configuration settings, *Computer* and *User*. Each of these settings is further defined by three categories of settings: *Software*, *Windows*, and *Administrative Templates*. These are the settings you use to control the OU (or other unit to which the GPO is applied).

The **Add** button lets you create a Group Policy Object Link. The link lets you apply an existing GPO to the OU. You will have the ability to navigate through the domain to locate the existing GPO and link it to the new OU.

The **Options** button gives you two options: *No Override* and *Disable*. **Disable** is very intuitive; it will disable the GPO. The **No Override** option is used by a parent OU's GPO to ensure that the settings in the GPO are not overridden by a child OU's GPO. These options can be accessed by right-clicking the GPO and selecting the option from the context menu.

The **Properties** button opens the **Properties** window for the GPO. The **Properties** window has three tabs: **General**, **Links**, and **Security**. The **General** tab displays a Summary section and a Disable section. The Summary section displays GPO information such as the date created, date last modified, revision versions, domain name, and the unique name of the GPO. The Disable section allows you to disable either or both sets of configuration settings. You can disable the Computer Configurations Settings and/or the User Configuration Settings. Disabling unused parts of the GPO increases performance. The **Links** tab displays all of the sites, domains, or OUs found that use the GPO. It has a **Find** button to assist you in locating where the GPO has been applied. The **Security** tab sets the permissions for the GPO.

The permissions that are set via the **Security** tab control the level of access that a user or group of users has over the GPO. The levels of permissions are:

- Full Control
- Read
- Write
- Create Child Objects
- Delete Child Objects
- Apply Group Policy

The last button is **Delete**, which is used to delete a GPO.

At the bottom of the **Group Policy** tab is the option to **Block Inheritance**. **Block Inheritance** will block settings from the GPOs that would otherwise be inherited from a parent OU. This gives the child OU the ability to control which settings to accept from the parent OUs. However, if the parent has set the **No Override** and the child sets **Block Inheritance**, the **No Override** setting takes precedence.



TEST DAY TIP

The relationship between GPOs and OUs is one that makes for easy test questions. Pay particular attention to the effects of the **No Override** setting and the **Block Inheritance** setting.

EXAM
70-294
OBJECTIVE
3.4.2

Delegating Control of OUs

Delegation of control over an OU is done to alleviate the tasks of the network administrators from performing the routine functions of an OU. Often, a manager or supervisor whose account is in the OU will have a better understanding of the daily tasks associated with the users and computers that belong to the OU, and is thus well positioned to take care of the OU. Delegation is a simple process. A wizard will walk you through the process. The Delegation of Control Wizard is discussed later in the chapter.

After you have decided to whom you want to delegate control, decide on which tasks to delegate. You have the ability to delegate management control over users and groups as well as the Group Policy Links. You can pass control of different activities to different people in the organization.

Specifically, the levels of delegations are:

- Create, delete, and manage user accounts
- Reset passwords on user accounts
- Read all user information
- Create, delete, and manage groups
- Modify the membership of a group
- Manage Group Policy Links

As you can see, delegation can reduce the amount of daily management tasks required by the network administrator.



TEST DAY TIP

The administrative task of delegating control to others is one that is likely to be covered on the exam. It is likely to be a straightforward scenario that will ask you to delegate control to another user or group of users. Pay attention to the levels of control that can be delegated.

EXAM
70-294

OBJECTIVE
1.5
1.5.1
1.5.2
3.3
3.3.2

Planning an OU Structure and Strategy for Your Organization

The OU structure can make your life easier—or it can do the opposite. If you spent time planning the structure and the implementation, the chances improve that your life will become easier and that you will be able to focus on the many facets of network administration without having to perform daily maintenance on the OUs and user issues such as resetting passwords. Your strategy should include the following:

- What OUs to create
- What policies need to be applied to cover the security requirements of the OU
- Who needs to be in charge of the OU (so you can delegate control to that user)

As with any structure, you will be faced with many decisions that need to be addressed; for example, whether a domain or OU is more appropriate for a given scenario. When making these decisions, remember to factor in the ease with which growth and changes can be accommodated.

Domains or Organizational Units

One of the challenges that you will face when planning your domain and OU structure is whether to create an OU or create a new domain when you need to split off part of the organization. There are some basic guidelines to help you determine which is the appropriate choice for a given scenario.

You will want to create new domains when the organization is decentralized and uses administrators for each of the sets of users and resources. Another reason to create multiple domains is when you need to create a GPO that will require different Password or Account Lockout Policies. That's because account policies can only be applied at the local or domain level, not at the OU level.

You should create OUs for everything else. Specifically, create OUs to reflect the organization's structure, especially if the organizational structure is likely to change. Create OUs so you can delegate control over users, groups, and resources. Domains are not easy to modify after they are created, so it is best to create OUs instead of domains, except when you have a specific reason to create separate domains that cannot be satisfied by the creation of OUs.



EXAM WARNING

You might have a question where you will need to apply GPOs to containers that have different Password or Account Lockout Policies requirements. Remember that Password and Account Lockout Policies can only be applied to domains. Otherwise, you would create OUs for the different GPOs.

Delegation Requirements

Delegation of control over an OU is frequently a necessity for many organizations. The delegation allows a local manager or IT staff member to control the OU. To delegate the control, you must be a member of the Enterprise Admins or Domain Admins global groups, or you must have been granted the privilege of delegating control.

From a security standpoint, when you delegate control, you should first determine the level of control that you want to grant. Just because you delegate basic administrative control over the OU does not mean that you fully relinquish control of the OU.



TEST DAY TIP

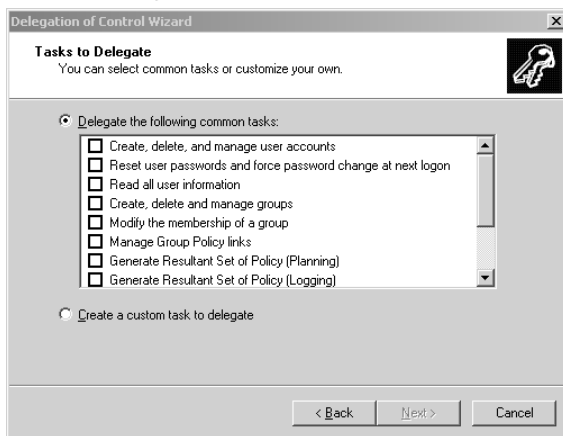
As you perform Exercise 5.03, pay attention to the levels of authority that you can delegate. The exam might ask questions concerning the various levels of control that can be granted.

EXERCISE 5.03

DELEGATING AUTHORITY

1. Click **Start | Programs | Administrative Tools | Active Directory Users and Computers**.
2. In the console tree, right-click the OU to be delegated. Select **Delegate Control** in the context menu. This invokes the **Delegation of Control Wizard**.
3. In the **Delegation of Control Wizard**, click **Next** to continue.
4. In the **User or Groups** window, click **Add** and then select the user who will receive the delegated control. Click **Add** and then click **OK**. Click **Next** to continue.
5. In the **Tasks to Delegate** window (see Figure 5.14), select the tasks that you want to delegate. Click **Next** to continue.

Figure 5.14 Delegation Tasks



6. In the **Completing Delegation of Control** window, review the information and click **Finish**.

Security Group Hierarchy

One of the issues that will need to be evaluated as you deploy Active Directory is the answer to this question: What is the *effective policy* that will be applied to a specific user? Because it is

possible for a user to have several layers of GPOs applied, it is very possible to have conflicting policies. This section discusses how to evaluate which policy will ultimately apply.

The first concept that needs to be covered is the order in which policies are applied. The first rule to remember is that a policy always overrides a profile setting. This becomes a factor as users might be moved from one OU where they use roaming profiles that allow the user a lot of liberty to configure their own settings. As these users are moved to another OU where the users' privileges are more controlled, they might notice that the user profile settings are overwritten by the OU policies.

The next concept is the order of the application of policies. Group policy is applied in this order:

- Local computer policy
- Site policy
- Domain policy
- OU policies, starting with the parent OU and working inward toward the security object through the child OUs

As an administrator, you still have further control over the application of policies. Windows Server 2003 Active Directory has two settings that help you with this control: **No Override** and **Block Inheritance**. The **No Override** setting is set to prevent a child OU policy setting from overwriting the policy setting of the parent. It does not apply if the policy setting is not set in the parent GPO.

The **Block Inheritance** setting allows you to control the inheritance of a policy setting in the parent by blocking it from being applied to the child. Even though you can set **Block Inheritance**, if the **No Override** option is set, **No Override** will be the setting that takes effect.



TEST DAY TIP

You might encounter questions on the exam that require you to evaluate a number of different GPOs applied at site, domain, and OU level and determine the effective policy for a particular user, computer, or OU. It is helpful, in these situations, when there are multiple nested OUs, to draw a diagram of the OU structure to help you see the relationships between parent and child containers.

Summary of Exam Objectives

In this chapter, we covered several of the Microsoft exam objectives. The first of these objectives is to establish trust relationships. Trust relationships are the relationships established between domains, trees, and forests so users in one domain can access the resources in another domain. This could be accomplished by creating new user accounts for the people who need to access the resources, but doing so would add to the administrative overhead of the domain. Microsoft developed a better solution: trust relationships.

Trusts come in many flavors to meet the needs of the situation where users in one domain need access to the resources in another domain. First, there are the default trusts created between parent and child domains. These trusts are automatically created to simplify usage of resources in a tree. The network administrator can create additional types of trusts such as external, shortcut, realm, and forest trusts. External trusts link two external domains. Shortcut trusts simplify the authentication paths needed to authenticate users. Realm trusts are created to connect a non-Windows network to a Windows Server 2003 domain. Forest trusts link forests together in the enterprise.

As you create these additional trust types, you can determine whether the trust will work in one direction only, or if it can work in both directions. When the trust works in both directions, it is called a two-way or bidirectional trust, and users in both domains have access to resources in both domains.

Another issue is whether the trust is transitive. A transitive trust "passes" through one trusted domain to another. A transitive trust implies a trust relationship when more than two domains are involved. If Domain A trusts Domain B, and Domain B trusts Domain C, then Domain A trusts Domain C. This is sometimes not the effect you want when creating trusts. The administrator has control over the transitive nature of the trust. As a further protection, SID filtering helps to prevent against elevation of privilege attacks that could potentially be launched by rogue users who have administrative access in the trusted domain.

The second part of this chapter covered working with organizational units (OUs). An OU is a container used to organize the resources and users of the domain. OUs can contain computers, users, groups of users, printers, shared directories, and other OUs. As the corporate infrastructure shifts, it is easy to move objects inside the Active Directory structure from one OU to another.

One of the major reasons for creating an OU is to apply policy settings that affect the Windows environment, security, and applications to the members of the OU. This is accomplished using Group Policy Objects (GPOs). Another major reason for creating OUs is to be able to delegate control to a local manager or supervisor. This empowers local supervisors with the ability to manage the users and computers within their realm of control.

Trusts and OUs are both important components of a Windows Server 2003 network, and thus it is important to understand both, not only to master the objectives of Exam 70-294, but to perform the duties of a network administrator.

Exam Objectives Fast Track

Working with Active Directory Trusts

- ☑ Trusts allow users in one domain to access resources in another domain without having to create additional accounts in the domain with the resources.
- ☑ Whenever a child domain is created, two-way transitive trusts are automatically created between the parent and the child.
- ☑ Realm trusts are created to join a Windows Server 2003 domain to a non-Windows Kerberos realm.
- ☑ Forest trusts are created between the root domains of two forests to allow users in one forest to access resources in the other forest.
- ☑ SID filtering is a security device that uses the domain SID to verify each security principal.

Working with Organizational Units

- ☑ OUs are Active Directory containers that can have users, groups, printers, shared folders, computers, and other OUs as members.
- ☑ OUs are created to help organize objects in the Active Directory; they are not security principals.
- ☑ The smallest scope to which a GPO can be assigned is an OU.
- ☑ Control of the OU can be delegated to other users to simplify the task of administration.

Planning an OU Structure and Strategy for Your Organization

- ☑ Create separate domains when you need decentralization of administrative functions and for GPOs that use different Password and Account Lockout Policies.
- ☑ You must delegate control over an OU for others to be able to manage the OU.
- ☑ GPOs are applied first to the local computer, then to the site, then to the domain, then to parent OUs, and finally to child OUs.
- ☑ You can control application of GPOs to child domains by using **Block Inheritance** or by setting **No Override**.

Exam Objectives Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

- Q:** What are the differences between external, realm, and shortcut trusts?
- A:** An external trust is created to establish a relationship with a domain outside your tree or forest. A realm trust is created to establish a relationship with a non-Microsoft network using Kerberos authentication. A shortcut trust is used to optimize the authentication process.
- Q:** What type of trust is needed to have users in a non-Windows Kerberos realm use resources in a Windows 2003 domain?
- A:** A realm trust will allow users in the non-Windows Kerberos realm to have access to the resources in a Windows 2003 domain.
- Q:** What type of trust needs to be created between the root domain and a domain that is several layers deep inside the same tree?
- A:** None. Transitive two-way trusts are automatically created between the layers of the tree structure.
- Q:** What is the difference between implied, implicit, and explicit trusts?
- A:** An implicit trust is one that is automatically created by the system. An example is the trusts created between parent and child domains. An explicit trust is one that is manually created. An example is a forest trust between two trees. An implied trust is one that is implied because of the transitive nature of trusts. An example is the trust between two child domains that are in different trees, and a tree-root trust was created between the roots of the trees.
- Q:** What exactly does SID filtering accomplish?
- A:** SID filtering is used to secure a trust relationship where the possibility exists that someone in the trusted domain might try to elevate his or her own or someone else's privileges.

Q: What is the difference between an OU, site, and domain.

A: All three are containers to which a GPO can be assigned. The domain is the basic building block of the organization. It can contain the other container types, site and OUs. The site is a container that will represent the physical layout of the organization. An OU is a logical container that can be used to implement security policies, run scripts, deploy applications, and delegate authority for granular administrative control.

Q: What is the difference between an OU and a security principal?

A: A security principal is a user, group, computer, or service that holds an account and can be given access to resources. An OU is a container that is used to organize objects in the Active Directory. OUs are also boundary units that are used to apply the security settings from a GPO.

Q: How and why is control of an OU delegated?

A: Control over a GPO is delegated to put the responsibility for the OU in the appropriate hands. Control is often delegated to the manager or supervisor responsible for the users and computers in the OU. You delegate control by right-clicking on the OU in **Active Directory Users and Computers** and selecting **Delegate Control** from the menu. This launches the **Delegation of Control** wizard. You can also set the user account that has management responsibilities from the **Managed By** tab in the OU's properties.

Q: How are GPOs applied?

A: GPOs applied to user configuration are applied as part of the logon process, whereas GPOs applied to computer configuration are applied as part of the boot process. First, any GPOs linked to the local computer are applied, followed by the site, then the domain, and finally the OUs. GPOs linked to the parent OU are applied first followed by the GPOs linked to the child. If a conflict exists in the settings of the various GPOs, the one applied last takes precedence.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Working with Active Directory Trusts

1. You are administering two domains, `mycompany.com` and `denver.hr.mycompany.com`. Users in `denver.hr.mycompany.com` need to access resources in `mycompany.com`. You want to optimize the trust relationships. What type of trust should you create to allow this?
 - A. Cross-domain trust
 - B. Shortcut trust
 - C. External trust
 - D. None
2. Your company, `mycompany.com`, is merging with the `yourcompany.com` company. The details of the merger are not yet complete. You need to gain access to the resources in the `yourcompany.com` company before the merger is completed. What type of trust relationship should you create?
 - A. Forest trust
 - B. Shortcut trust
 - C. External trust
 - D. Tree Root trust
3. Your boss just informed you that your company will be participating in a joint venture with a partner company. He is very concerned about the fact that a trust relationship needs to be established with the partner company. He fears that an administrator in the other company might be able to masquerade as one of your administrators and grant himself privileges to resources. You assure him that your network and its resources can be protected from an elevated privilege attack. Along with the other security precautions that you will take, what will you tell your boss that will help him rest easy about the upcoming scenario?
 - A. The permissions set on the Security Accounts Manager (SAM) database will prevent the other administrators from being able to make changes.

- B. The SIDHistory attribute tracks all access from other domains. Their activities can be tracked in the System Monitor.
 - C. The SIDHistory attribute from the partner's domain attaches the domain SID for identification. If an account from the other domain tries to elevate its own or another user's privilege, the SID filtering removes the SID in question.
 - D. SID filtering tracks the domain of every user who accesses resources. The SIDHistory records this information and reports the attempts to the Security log in Event Viewer.
4. You recently completed a merger with yourcompany.com. Corporate decisions have been made to keep the integrity of both of the original companies; however, management has decided to centralize the IT departments. You are now responsible for ensuring that users in both companies have access to the resources in the other company. What type of trust should you create to solve the requirements?
- A. Forest trust
 - B. Shortcut trust
 - C. External trust
 - D. Tree Root trust
5. You recently created a trust relationship with a partner company for collaboration on a joint project. This partner company has many such joint projects and has many trust relationships with other companies. You created a share containing all the files needed for the joint project. You worked with the partner company's administrator and added your project members to one of his existing universal groups that contains all of the members in his domain who need access to the project files. You added them to the permissions on the folder and the permissions on the share. You granted the universal group Read access to the share permission and Read & Execute access to the folder via NTFS permissions. SID Filtering has been enabled. The users in the universal group are now complaining that they cannot gain access to the project's files. What do you need to do to fix the problem?
- A. You need to upgrade the level of permissions on the folder to **Modify** so that the universal group can have access.
 - B. You need to upgrade the level of permissions to **Change** on the share so that the universal group can have access.
 - C. You need to break the trust relationship and recreate it; it has a corrupted file.
 - D. You need to have the domain administrator from the partner domain verify that only members from his domain are in the universal group

Working with Organizational Units

6. The development team of your company has started a new research project. They want to ensure that only the members of their project team are allowed to see the new directories that they create. You created a new OU that contains the user accounts of the development team, the computers they will be using, a shared folder where they are going to place their research documents, and several printers that are to be isolated from the rest of the company. They are concerned about who will have access to the new directories. How will you protect the directories from unauthorized access?
 - A. Create a GPO that will limit access to the directories. Apply the GPO to the new OU.
 - B. Create a GPO that will limit access to the directories. Apply the GPO to the domain.
 - C. Create a security group that contains the members of the research group. Remove the **Everyone** group from the ACL. Add the new group to the ACL and grant it the appropriate permissions.
 - D. Do nothing. Since the directories and files are part of an OU, no one outside the OU can access them.

7. You created three OUs for your domain: one called *Corp*, and two child OUs called *Sales* and *Tech*. You create two GPOs, one called *Desktop* the other called *Network*. The Desktop GPO specifies the desktop settings for all users. The Network GPO specifies the network and Registry policies. The Registry policy prohibits users from being able to edit the Registry. You first apply the Desktop GPO to the Corp OU and then apply the Network GPO to the Corp OU. You want the members of the Tech OU to be able to modify Registry settings. What should you do?
 - A. Nothing; because the GPOs were not applied to the Tech OU, they will not affect the users.
 - B. Nothing; because you applied the Desktop GPO first, the Desktop GPO will not take effect.
 - C. You should set **No Override** on the Tech OU so that its settings are not overridden.
 - D. You should set **Block Inheritance** on the Tech OU so that the settings from the parent OU are not applied to the child OU.

8. Your Active Directory domain has one site and five OUs. *Marketing* and *Technical* are child OUs to the *Corp* OU. The *Marketing* OU is a parent to the *Sales* and *PR* OUs. You are using GPOs to configure environment and security policies on the network. The following restrictions are in place:

- **Corp OU** Disable Registry editing tools for all users
- **Marketing OU** Disable modification of network connections for all users
- **Technical OU** Corporate logo as desktop wallpaper for all users
- **Sales OU** 3D Pipes screensaver for all users
- **PR OU** High Contrast #1 color scheme for all users

Which restriction or restrictions will be in place for users in the *Sales OU*? (Choose all that apply.)

- A. Disable Registry editing tools for all users.
 - B. Disable modification of network connections for all users.
 - C. Corporate logo as desktop wallpaper for all users.
 - D. 3D Pipes screensaver for all users.
 - E. High Contrast #1 color scheme for all users.
9. You have an OU called *Support*. You have a GPO called *RegEdit*. The only setting in the *RegEdit* GPO is that the use of the Registry editing tools has been disabled in the *User Configuration* node. For performance reasons, the decision has been made to limit the numbers of GPOs that are processed at logon. The decision has been made to remove the requirement to disable the use of the Registry editing tools. What should your course of action be to implement the new decisions?
- A. Remove the *RegEdit* GPO from the *Support OU*.
 - B. Create a new GPO that enables the use of the Registry editing tools. Apply the new GPO to the *Support OU*.
 - C. Edit the Registry on the computers used by the *Support OU* that will allow for use of the Registry editing tools.
 - D. Configure a local GPO to allow the use of the Registry editing tools. Set the *No Override* option to this policy.
10. You created three OUs for your domain: one called *Corp*, and two child OUs called *Sales* and *Tech*. You create two GPOs, one called *Desktop* and the other called *Network*. The *Desktop* GPO specifies the desktop settings for all users. The *Network* GPO specifies the network and Registry policies. The *Desktop* policy prohibits users from being able to change their wallpaper. You first apply the *Desktop* GPO to the *Corp OU*, and then apply the *Network* GPO to the *Corp OU*. You delegated control of the OU to the senior member of the *Tech* group. Later, the *Tech OU* manager modifies the *Desktop* GPO to allow his users to change their wallpaper. What should you do to ensure that their changes will not take effect?

- A. Nothing, since the GPOs were not applied to the Tech OU, they will not affect the users.
 - B. You should set **No Override** on the Tech OU so that its settings are not overridden.
 - C. You should set **No Override** on the Corp OU so that its settings are not overridden.
 - D. You should set **Block Inheritance** on the Tech OU so that the settings from the parent OU are not applied to the child OU.
11. Your network consists of a single domain and five OUs. The parent OU is named *Corp*. Corp has two child OUs, *First Floor* and *Second Floor*. The First Floor OU has one child OU, *Sales*. The Second Floor OU has one child OU, *Administration*. All of the company's DCs are members of the Corp OU. The First Floor and Second Floor OUs contain the resources that belong to their respective floors. The Sales OU has nonadministrative computers, users, and groups. The Administration OU has the administration computers, users, and groups. You need to design a domainwide security policy that will accomplish the following goals:
- All users need to have the same password and lockout policy.
 - Audit policies are required for only the DCs.
 - The nonadministrative computers do not need the same level of security applied to them as is required for the administrative computers.
 - The number of group policies to be processed at logon needs to be minimized.

You take the following actions:

- Create a single GPO.
- Import a security template for the DCs.
- Link the GPO to the domain.

Which of the desired results are achieved by your actions?

- A. All users have the same password and lockout policy.
- B. Audit policies implemented only on the DCs.
- C. The nonadministrative computers have the same level of security applied to them as is required for the administrative computers.
- D. The number of group policies to be processed at logon is minimized.

Planning an OU Structure and Strategy for Your Organization

12. Your Active Directory domain consists of one site. You have three OUs. The Corp OU is a parent OU to the Sales OU and Training OU. You have specified restrictions in various group policies and included them in GPOs. On the Corp OU, there is a linked GPO, which prevents users from using Registry editing tools. The Sales OU has a linked GPO that specifies a company logo as the desktop for all users. The Training OU has a linked GPO that disables users from modifying network connections. All other group policy settings are set to defaults. What restrictions (if any) will users in the Sales OU be under when they log on to the network? (Choose all that apply.)
- A. They cannot edit the Registry.
 - B. They have the company logo as their desktops.
 - C. They cannot modify network connections.
 - D. They will have no restrictions.
13. You have been tasked to ensure that network security policies are in place, and standards are implemented for users' configurations. The network is a single Active Directory domain network. There are five OUs: Corp, Sales, Marketing, Development, and Technical. The Corp OU is a parent OU to all other OUs. You are given the following list of objectives to meet:
- All users must be prohibited from editing their Registries.
 - All users must have a password of at least eight characters.
 - Users in the Sales and Marketing OUs must not be able to store more than 50MB of data on any server.
 - Users in the Development OU must change their passwords every 30 days.
 - All policy settings should only affect their intended targets.

Which of the following solutions will accomplish all of your objectives?

- A. Create a GPO called Policy, with settings prohibiting users from using Regedit, and requiring passwords of at least eight characters. Link Policy to the Corp OU. Create a GPO called Data, with disk quotas set at 50MB. Link Data to the Sales OU and to the Marketing OU. Create a GPO called Password, making users change their passwords every 30 days. Link Password to the Development OU.

- B. Create a GPO called Policy, with settings prohibiting users from using Regedit, and requiring passwords of at least eight characters. Link Policy to the domain. Create a GPO called Data, with disk quotas set at 50MB. Link Data to the Corp OU. Create a GPO called Password, making users change their passwords every 30 days. Link Password to the Development OU.
 - C. Create a GPO called Policy, with settings prohibiting users from using Regedit, and requiring passwords of at least eight characters. Link Policy to the Corp OU. Create a GPO called Data, with disk quotas set at 50MB. Link Data to the Corp OU. Create a GPO called Password, making users change their passwords every 30 days. Link Password to the Corp OU.
 - D. Create a GPO called Policy. In Policy, define settings prohibiting users from using Regedit, requiring passwords of at least eight characters, setting disk quotas at 50MB, and a maximum password age of 30 days. Link Policy to the Corp OU.
14. Your Active Directory domain has two OUs. The Corp OU is a parent OU to the Technical OU. You have implemented a GPO linked to the Corp OU. You do not want those settings affecting the users in the Technical OU. How can you accomplish this with minimal effort?
- A. On the GPO linked to the Technical OU, select **Block Policy** inheritance.
 - B. On the GPO linked to the Corp OU, select **Block Policy** inheritance.
 - C. On the GPO linked to the Technical OU, negate any options set in the Corp OU by choosing **Disabled** for those options.
 - D. On the GPO linked to the Technical OU, select **No Override**.
15. John Smith is a junior network administrator for your company. His user account is JSmith. You want him to take charge of linking all network group policies to the appropriate OUs. Because of his experience level, you do not want him to have additional controls over the OUs. What is the easiest way to accomplish this?
- A. Use the Delegation of Control Wizard. Select JSmith, and check **Create, delete, and manage groups**.
 - B. Use the Delegation of Control Wizard. Select JSmith, and check **Manage Group Policy links**.
 - C. Use the Delegation of Control Wizard. Select JSmith, and check **Create and Modify Group Policy**.
 - D. Use the Delegation of Control Wizard. Select JSmith, and check **Apply Group Policy**.

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|-------------------|-----------------|
| 1. D | 9. A |
| 2. C | 10. C |
| 3. C | 11. A, D |
| 4. A | 12. A, B |
| 5. D | 13. A |
| 6. C | 14. A |
| 7. D | 15. B |
| 8. A, B, D | |

MCSA/MCSE 70-294

Working with Active Directory Sites

Exam Objectives in this Chapter:

- 1.4 Implement an Active Directory site topology.
- 2.2 Manage an Active Directory site.
 - 2.2.3 Configure site boundaries.
 - 1.4.1 Configure site links.
 - 2.2.2 Configure site link costs.
 - 2.2.1 Configure replication schedules.
 - 2.5.1 Diagnose and resolve issues related to Active Directory replication.
 - 1.4.1 Configure site links.
- 2.3 Monitor Active Directory replication failures. Tools might include Replication Monitor, Event Viewer, and support tools.
 - 2.3.1 Monitor Active Directory replication.
 - 2.3.2 Monitor File Replication service (FRS) replication.

Introduction

In the previous chapter, we saw the logical structure of the network as defined by forests and domains. Sites and the subnets, of which sites are comprised, define the physical structure of an Active Directory network. Sites are important in an enterprise-level multiple location network for creating a *topology* that optimizes the process of replicating Active Directory information between domain controllers (DCs). Sites are used for replication and for optimizing the authentication process by reducing authentication traffic across slow, high-cost WAN links. Site and subnet information is also used by Active Directory-enabled services to help clients find the nearest service providers.

In this chapter, we discuss the role of sites in the Active Directory infrastructure, and how replication, authentication, and distribution of services information work within and across sites. We explain the relationship of sites with *domains* and *subnets*, and how to create sites and site links.

You'll also learn about site *replication* and how to plan, create, and manage a replication topology. We'll walk you through the steps of configuring replication between sites, and discuss how to troubleshoot replication failures.

Understanding the Role of Sites

In today's distributed network environment, the communication must always be rapid and reliable. Geographical and other restrictions resulted in the need to create smaller networks, known as *subnets*. These subnets provide rapid and reliable communication between locations, which can also be attained in larger networks by using Microsoft Windows Server 2003 Active Directory Sites. They ensure rapid and reliable communication by using the methods offered by Microsoft Windows Server 2003 Active Directory Sites to regulate inter-subnet traffic.

A *site* defines the network structure of a Windows Server 2003 Active Directory. A site consists of multiple *Internet Protocol* (IP) subnets linked together by rapid and reliable connections. The primary role of sites is to increase the performance of a network by economic and rapid transmission of data. The other roles of sites are replication and authentication. The Active Directory physical structure manages when and how the authentication and replication must take place. The Active Directory physical structure allows the management of Active Directory replication scheduling between sites. The performance of a network is also based on the location of objects and *logon authentication* as users log on to the network.



TEST DAY TIP

As a network administrator, you must be familiar with the various roles and services offered by the Active Directory Sites. You needn't worry about memorizing every detail for this particular exam. What you do have to know are the basics of how each role and services of Active Directory Sites works, and how Active Directory Sites can be used efficiently in terms of data transmission as part of a large network.

Replication

Replication is defined as the practice of transferring data from a data store present on a source computer to an identical data store present on a destination computer to *synchronize* the data. In a network, the directory data must live in one or more places on the network to be equally available to all users. The Active Directory directory service manages a replica of directory data on one or more DCs, ensuring the availability of directory data to all users. The Active Directory works on the concept of sites to perform replication efficiently, and uses the *Knowledge Consistency Checker* (KCC) to choose the best replication topology for the network automatically.

NOTE



The KCC is a process that runs on a DC, and identifies the most efficient replication topology for the network automatically, based on the data provided by the network in Active Directory Sites and Services.

Authentication

Authentication is a process by which a system validates users, using the logon information provided. The authentication process includes the confirmation of the source and integrity of information, such as verifying the identity of a user or computer. The information such as user's name and password are verified with the data available in the system. If the system finds a match, access is granted and an access token is generated that is used to subsequently determine the user's level of access to objects according to the DACLs on those objects. The granting of the level of access based on permissions is called *authorization*.

An important characteristic of authentication in the Windows Server 2003 family is its support for *single sign-on*. The single sign-on feature allows a user to log on to the network once, using a single password, and authenticate to any computer in a network.

The single sign-on feature offers the following security advantages:

- For a user, the use of a single password reduces ambiguity and increases the work efficiency of the system.
- For administrators, the level of administrative support needed for authenticating the domain users is reduced, since the administrator needs only to maintain one account per user.



EXAM WARNING

Make sure you are familiar with the advantages of the single sign-on feature and how it works.

Windows Server 2003 uses two methods to carry out authentication:

- Interactive logon authentication
- Network authentication



TEST DAY TIP

As a network administrator, you must be familiar with the various authentication mechanisms offered by Active Directory Sites. You needn't worry about memorizing every detail for this particular exam. What you do have to know are the basics of how each of the authentication mechanisms of the Active Directory Sites works, and how Active Directory Sites can be used efficiently in terms of user authentication in a network.

Interactive Logon Authentication

Interactive logon authentication verifies the user's logon information to either a domain account or to a local computer. This process of authentication is based on the type of user account, such as a domain account or a local computer account:

- With a domain account, a user logs on to the network by providing logon information such as a password or smart card, using single sign-on data stored in the Active Directory directory service. When a user logs on to the network with a domain account, the user can access resources both in the domain to which he or she logs on and any other trusted domains.
- With a local computer account, a user logs on to a local computer by providing logon information stored in the *Security Accounts Manager* (SAM) on the local machine.



NOTE

SAM is a local security account database for local computer accounts. Local user accounts are usually stored on workstations or servers, and can only be used to access the local computer, not resources on any other computer on the network.

Network Authentication

Network authentication verifies the user's identification to a network service to which the user tries to gain access. To offer this type of authentication, the security system of Windows Server 2003 supports authentication mechanisms:

- Kerberos V5
- Secure Socket Layer/Transport Layer Security (SSL/TLS)

When a domain account is used, network authentication occurs transparently and in the background via Kerberos or TLS/SSL. Users who use a local computer account must give user credentials such as a username and password while trying to gain access to a network resource.



EXAM WARNING


Make sure you know the differences between interactive logon authentication and network authentication in Windows Server 2003.

Distribution of Services Information

Active Directory distributes a wide range of service information. The DCs are also used to distribute directory information and generate responses for each service request. The Active Directory distributes service-centric information such as *configurations* and *bindings*. The distribution of this type of information enables the services to be more accessible by clients and is easily manageable for administrators.

The distribution of services information in Active Directory enables the client and applications to get information from the directory. This information is then used to access the services offered by the servers present on the network. Figure 6.1 shows how the services information is accessed between the client, server, and a DC in a network.

Figure 6.1 Services Information Shared between a Client, Server, and a Domain Controller



In Figure 6.1, the client shares the services information between a client, server, and a DC in three steps:

1. The client makes a request.
2. The client receives the services information from a DC as a response.
3. The clients available on the network server then use the services information.



TEST DAY TIP

Make sure you know the wide range of services information offered by the Active Directory Sites. Be aware of how the services information is accessed between the client, server, and a DC on a network.

Certain sets of services are distributed by the directories by default, including file and print services, storage management, Active Directory, and management services. These sets of services can be modified in the directories to meet the needs of your network environment. The distribution of services to the directory provides the following benefits:

- **Resource availability** This Active Directory model is a service-centric model that enables the client to provide access to the distributed network services. Since the services information is distributed to the directory, clients needn't store the resource's location.
- **Administration** Distributing services in Active Directory enables the administrator to resolve configuration-related problems in a network centrally, instead of having to visit individual computers. This feature ensures that all the services employ the latest configuration information.
- **Publishing services** This process enables the data or operations available to the network users. Publishing a service in Active Directory enables users and administrators to move from a machine-centric view of the network to a service-centric view.



EXAM WARNING

Make sure you are familiar with the benefits of distribution of services to the directory, and how it works to provide them for you.

Relationship of Sites to Other Active Directory Components

A site is as a collection of inter-connected computers that operates over IP subnets. A site is also a place on a network having high bandwidth connectivity. The relationship of sites to Active Directory components is based on the following network operations performed by sites:

- Control of replication occurrences
- Changes made with the sites
- How efficiently DCs within a domain can communicate

Relationship of Sites and Domains

A site can contain one or more domains, and a domain can be part of one or more sites. Sites and domains do not have to maintain the same *namespace*. Sites and domains are inter-related to each other because sites control replication of the domain information.

Head of the Class...

The Relationship of Sites and Domains

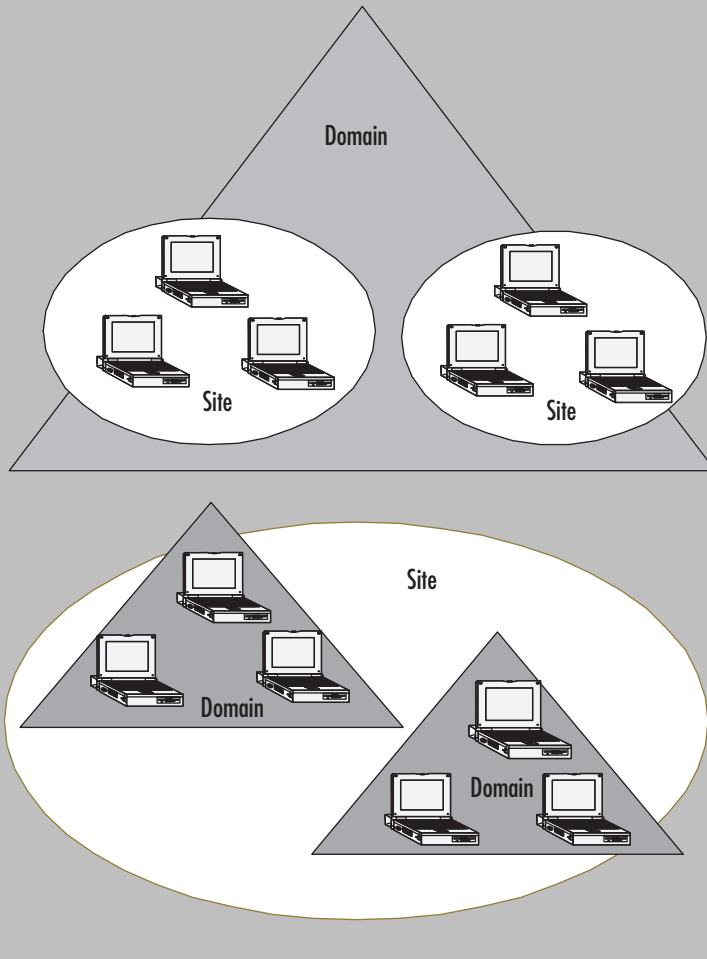
Domains are also defined as *units of replication*. Through the use of SRV records, the DNS server provides information regarding the location of domain controllers in various sites. A Domain Name System (DNS) server recognizes each domain that is present in a particular site. If your network requires more than one domain, you can easily create multiple domains. Figure 6.2 illustrates the relationship between sites and domains in a network, and helps us to understand that a site can have one or more domains, and a domain can have one or more sites.

In Figure 6.2, we see how multiple sites reside in a single domain, and how a single site can consist of multiple domains. A domain provides the following benefits:

- Organizing domain objects.
- Publishing of resources and information about domain objects.
- Applying Group Policy Objects (GPOs) to the domain to perform resource and security management
- Delegating authority eliminates the need for administrators with broad administrative authority.
- Security policies and settings such as user rights and password policies do not change from one domain to another.
- Each domain stores only the information about the objects located in that domain.

Continued

Figure 6.2 The Relationship of the Sites and Domains Present in a Network



EXAM WARNING

Make sure you are familiar with the benefits provided by a domain, and how it works to provide them for you.

For more information on the working of domains, see Chapter 4, “Working with Forests and Domains.”

Physical vs. Logical Structure of the Network

The sites present in an Active Directory denote the *physical structure* of a network. The physical structure information is available as site and site link objects in the directory. This information is used to build the most efficient replication topology. Generally, Active Directory Sites and Services are used to define sites and site links.

Sites represent the *physical structure* of the network, and domains represent the *logical structure* of the organization. In Active Directory, sites map the physical structure of a network, while domains map the logical or *administrative* structure of an organization. This partitioning of physical and logical structure offers the following advantages:

- You can develop and manage the logical and physical structures of your network independently.
- You do not have to base domain namespaces on your physical network.
- You can deploy DCs for multiple domains within the same site.
- You can deploy DCs for the same domain in multiple sites.



TEST DAY TIP

Make sure you know and understand the differences between the physical and the logical structure of the network. Be aware of how each is used to build the most efficient replication topology.


The Relationship of Sites and Subnets

In Active Directory, a site consists of a set of computers that are inter-connected in a local area network (LAN). Computers within the same site typically exist in the same building, or on the same campus network. A single site consists of one or more IP subnets. These subnets are a section of an IP network, with each subnet having a unique network address.

A subnet address consists of a cluster of neighboring computers in much the same way as the postal codes group neighboring postal addresses. Figure 6.3 shows one or more clients residing within a subnet that defines an Active Directory site.

The subnet created through Active Directory Sites and Services are sections of an IP network, with each subnet having a unique network address. In Figure 6.3, 172.16.224.0/19 is a unique network address of the Active Directory site.

Sites and subnets are represented in Active Directory by site and subnet objects, which we create through the Active Directory Sites and Services administrative tool. Each site object is associated with one or more subnet objects.

Figure 6.3 Active Directory Site with One or More Client Computers

EXAM
70-294
OBJECTIVE
1.4
2.2
2.2.3

Creating Sites and Site Links

In the previous sections, we discussed the concepts of sites and subnets. To review, sites and the subnets define the physical structure of an Active Directory network. A site is a collection of inter-connected computers that operate over subnets, sharing a network with high bandwidth connections. The high bandwidth connection is represented by the difference between the highest and lowest frequencies in a given range. Site links represent physical connections between sites, which enables communication between sites.



NOTE

The Windows Server 2003 Active Directory consists of the default site link, named DEFAULTIPSITELINK, which is created automatically when the first domain in the network is created. This link is assigned to the Default-First-Site-Name site. These are the names assigned automatically when you create the first site. You should change the default names to something more descriptive.

Site Planning

You should plan thoroughly before creating and deploying an Active Directory. Site planning enables you to optimize the efficiency of the network and reduce administrative over-

head. High-performance sites are developed based on the proper planning of the physical design of your network. Site planning enables you to determine exactly which sites you should create and how they can be linked using *site links* and *site link bridges*. Site information is stored in the *configuration partition*, which enables you to create sites and related information at any point in your deployment of Active Directory.

Site planning enables you to publish site information in the directory for use by applications and services. Generally, the Active Directory consumes the site information. You'll see how replication impacts site planning later in the chapter.

Criteria for Establishing Separate Sites

When you initially create a domain, a single default Active Directory site called *Default-First-Name* is created. This site represents your entire network. A domain or forest consisting of a separate site can be highly efficient for a LAN connected by high-speed bandwidth.



NOTE

A forest is defined as multiple Active Directory domains that share the same class, site, attribute definitions, and replication information (but not necessarily the same namespace). The domains present in the same forest are linked with two-way transitive trust relationships.

When a network consists of a single subnet or multiple subnets joined by reliable, high-speed links, a single site topology offers the following advantages:

- Simplified replication management
- Regular directory updates between all DCs

Establishing a single site topology enables all replication to occur as intrasite replication, which requires no manual replication configuration. A single site topology design enables DCs to receive updates with respect to directory changes.



NOTE

Intrasite replication refers to replication among DCs within the same site. *Intersite* replication refers to replication among DCs located at different sites.

Creating a Site

Sites are created using the Active Directory Sites and Services tool of Windows Server 2003. Exercise 6.01 walks you through the steps involved in creating a site.

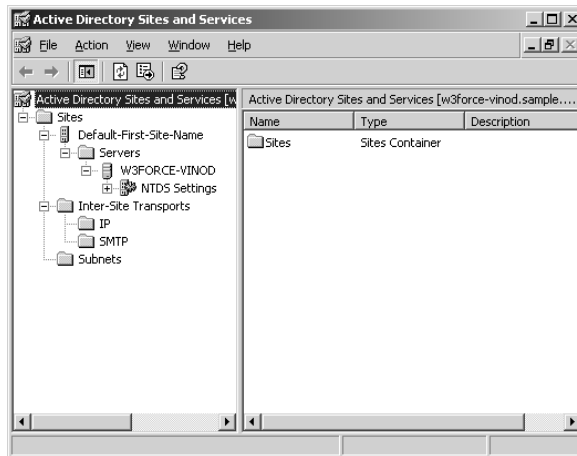
Active Directory Sites and Services tool is a Microsoft Management Console (MMC) that can be used to administer the replication of directory data. This tool can also be used to create new sites, site links, subnets, and so forth.

EXERCISE 6.01

CREATING A NEW SITE

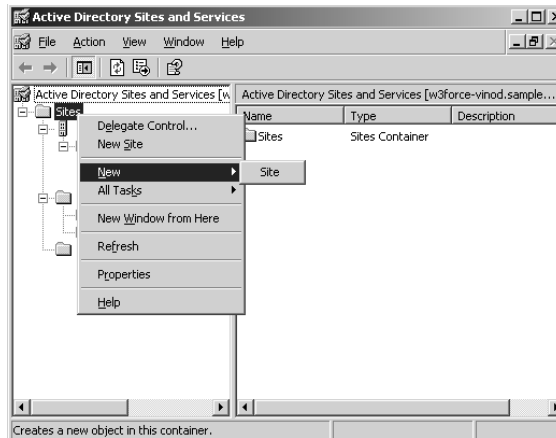
1. To open the **Active Directory Sites and Services** tool, click **Start | Control Panel | Administrative Tools | Active Directory Sites and Services**. The Active Directory Sites and Services console appears as shown in Figure 6.4.

Figure 6.4 The Active Directory Sites and Services Tool




2. Highlight the **Sites** folder in the left-hand tree pane of the **Active Directory Sites and Services** console. Right-click and select **Sites folder New | Site** option from the context menu as shown in Figure 6.5.

Figure 6.5 The New Site Option



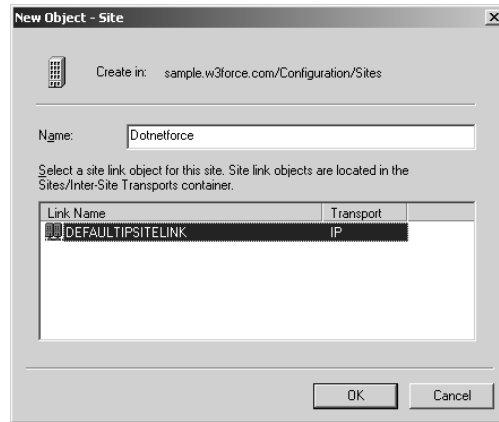
3. Selecting the **New Site** option opens a **New Object – Site** dialog box as shown in Figure 6.6.

Figure 6.6 The New Object – Site Dialog Box



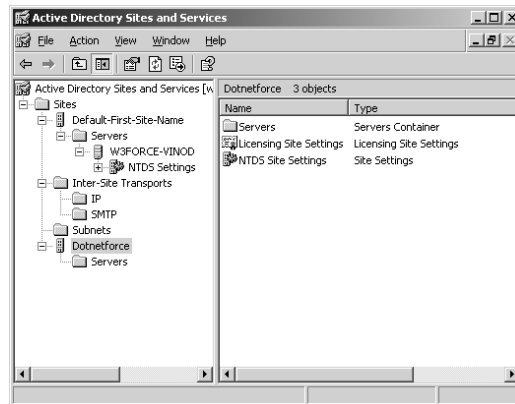
4. Type the name of the site in the **Name** box present in the **New Object – Site** dialog box as shown in Figure 6.7.
5. Select an initial site link object for the site from the **New Object – Site** dialog box as shown in Figure 6.7.

Figure 6.7 The Name of the Site



6. Click **OK**. This completes the process of creating a site using the **Active Directory Sites and Services** tool. Figure 6.8 shows the initial site link object of the site.

Figure 6.8 The Initial Site Link Object for the Site



Renaming a Site

Renaming a site is one of the first tasks you should perform when administering a site structure. When you create a site initially, it is created with the default name Default-First-Site-Name. This name can be changed based on the purpose of the site, such as the name of the physical location.

A site is also renamed when a network of an organization is expanded by one or more sites. Even if an organization is located in a single location, it makes sense to rename the Default-First-Site-Name, because you never know when the network will expand. Renaming a site enables administrators to differentiate sites present in a network easily and perform administration tasks efficiently.

When a DC becomes aware that its site has been renamed, it will update its DNS records appropriately. Because of issues with cached DNS lookups and client caching of site names that will lead to temporary delays in connectivity directly after a rename, it's best to name and rename sites as early as possible in the deployment. After renaming a site, it's advisable to manually force replication with other DCs in the same site.

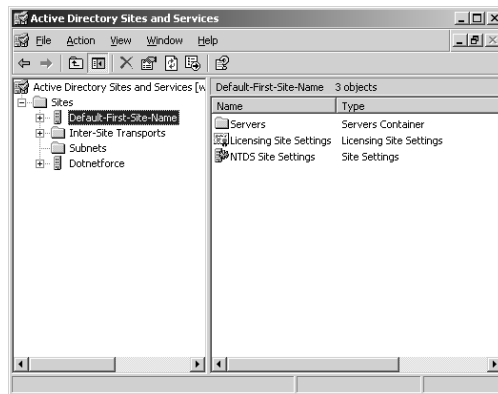
Sites are renamed using the Active Directory Sites and Services tool of Windows Server 2003. Exercise 6.02 walks you through the steps involved in renaming a site.

EXERCISE 6.02

RENAMING A NEW SITE

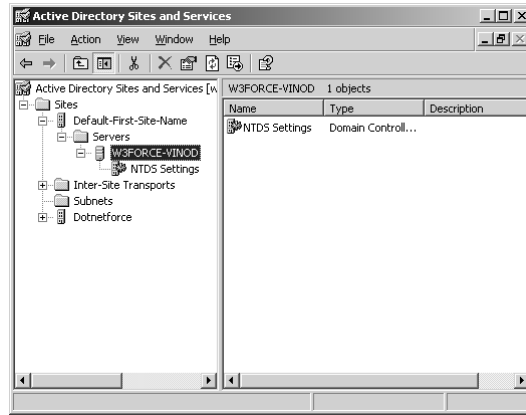
1. To open the **Active Directory Sites and Services** tool, click **Start | Control Panel | Administrative Tools**. Double-click **Active Directory Sites and Services**. The Active Directory Sites and Services dialog box appears as shown in Figure 6.9.

Figure 6.9 The Active Directory Sites and Services Tool



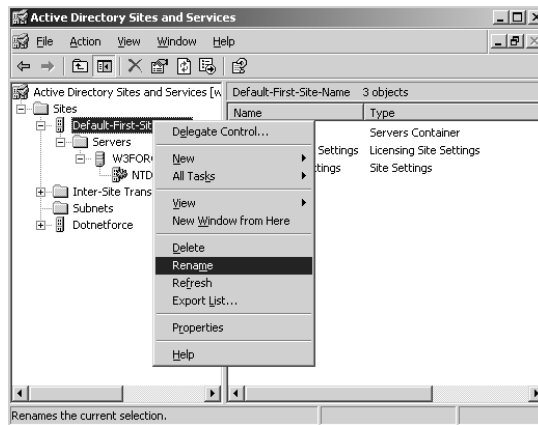
2. Highlight the **Sites** folder in the left-hand tree pane of the **Active Directory Sites and Services** console. Expand the **Sites** folder, and you'll see the sites shown with icons of small, yellow office buildings as shown in Figure 6.10.

Figure 6.10 The Sites Folder

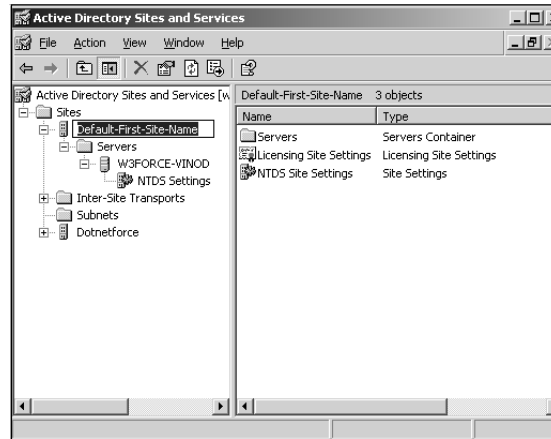


3. Right-click the site you want to rename and select the **Rename** option from the context menu as shown in Figure 6.11.

Figure 6.11 The Rename Option



4. Type the new name of the site in the **Name** box in the left console pane as shown in Figure 6.12.

Figure 6.12 The Rename Dialog Box

5. Click **OK**. This completes the process of renaming a site using the Active Directory Sites and Services tool.

Creating Subnets

Subnets are associated with the Active Directory sites to match client computers. The subnets are denoted by a range of IP addresses. The Active Directory Sites and Services user interface prevents you from having to provide the subnet names manually; instead, you are prompted for a network address. An example of a subnet name for an IP version 4 network is 10.14.208.0/20. This IP address consists of two portions: the network address appears before the slash, and after the slash is a representation of the subnet mask. Some common subnet masks and the corresponding slash notations are shown in Table 6.1. The number following the slash indicates the number of binary digits (bits) that make up the network partition of the IP address. 255 in decimal translates to 11111111 in binary (8 bits), thus you can see how the subnet masks in Table 6.1 translate to the corresponding slash notations.

Table 6.1 Subnet Masks and Slash Notation

Subnet Mask	Slash Notation
255.0.0.0	/8
255.255.0.0	/16
255.255.255.0	/24

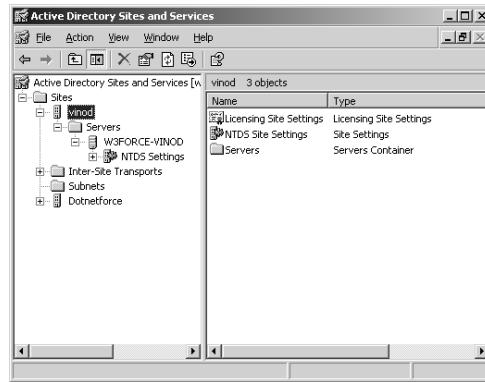
Subnets are created using the Active Directory Sites and Services tool of Windows Server 2003. Exercise 6.03 shows the steps involved in creating subnets.

EXERCISE 6.03

CREATING SUBNETS

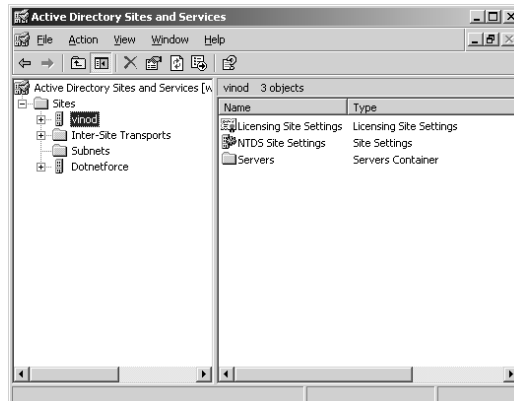
1. To open the **Active Directory Sites and Services** tool, click **Start | Control Panel | Administrative Tools**, and then double-click **Active Directory Sites and Services**. The Active Directory Sites and Services console appears as shown in Figure 6.13.

Figure 6.13 The Active Directory Sites and Services Tool



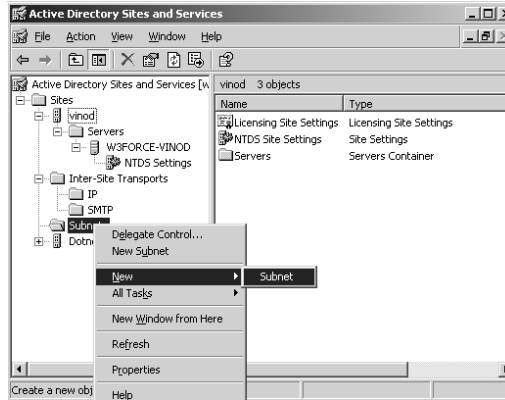
2. Highlight the **Sites** folder in the left tree pane of the **Active Directory Sites and Services** console. Expand the **Sites** folder as shown in Figure 6.14.

Figure 6.14 The Sites Folder



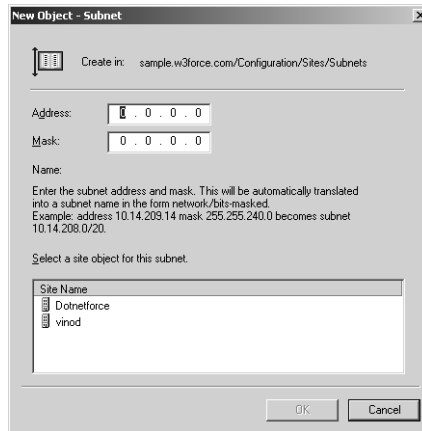
3. Right-click **Subnets** and select **New Subnet** from the context menu as shown in Figure 6.15.

Figure 6.15 The New Subnet Option

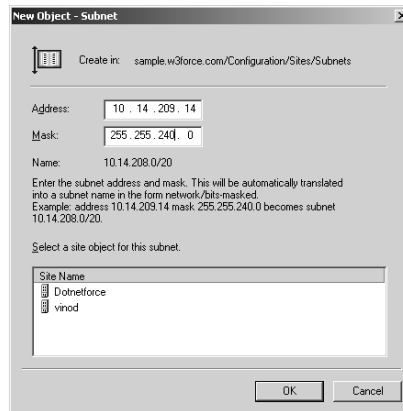


4. Selecting the New Subnet option opens a **New Object – Subnet** dialog box as shown in Figure 6.16.

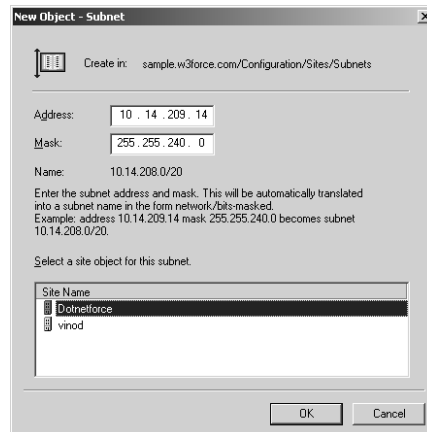
Figure 6.16 The New Object – Subnet Dialog Box



5. Type the network address and subnet mask in the form of dotted decimal notation in the text boxes present in the **New Object – Subnet** dialog box as shown in Figure 6.17.

Figure 6.17 The Network Address and Subnet Mask

6. Select a site object for this subnet from the list provided in the **New Object – Subnet** dialog box as shown in Figure 6.18.

Figure 6.18 The Site Object

7. Click **OK**. This completes the process of creating a subnet using the Active Directory Sites and Services tool.

Associating Subnets with Sites

After creating sites and subnets, the next step is to associate your subnets with sites. Computers on Active Directory networks communicate with each other using the Transmission Control Protocol/Internet Protocol (TCP/IP) assigned to sites based on their

locations in a subnet. Remember that a site consists of one or more IP subnets. You specify the subnets associated with each site on your network by creating subnet objects in the Active Directory Sites and Services console. The association of subnets with sites enables the computers on the Active Directory network to use the subnet information to find a DC in the same site, so that authentication traffic will not cross over WAN links. Active Directory also uses subnets during the replication process to determine the best routes between DCs.

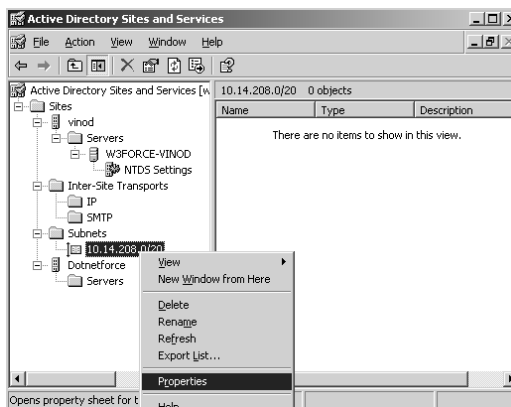
Subnets are associated with sites using the Active Directory Sites and Services tool of Windows Server 2003. Exercise 6.04 walks you through the steps involved in associating subnets with sites.

EXERCISE 6.04


ASSOCIATING SUBNETS WITH SITES

1. To open the **Active Directory Sites and Services** tool, click **Start | Control Panel | Administrative Tools**, and then double-click **Active Directory Sites and Services**.
2. Highlight the **Subnet** folder present in the left tree pane of the **Active Directory Sites and Services** console (see Figure 6.19).

Figure 6.19 The Subnet Folder



3. Right-click the newly created subnet and select the **Properties** option; this will open a Properties dialog box as shown in Figure 6.20.
4. Associate any site with this subnet by selecting the available site from the site drop-down menu, and click **OK**, as shown in Figure 6.20.

Figure 6.20 Subnet Dialog Box for Associating/Changing the Site

5. Click **OK**. This completes the process of associating a subnet with a site using the Active Directory Sites and Services tool.

EXAM
70-294OBJECTIVE
1.4.1
2.2.2

Creating Site Links

After creating and defining the scope of each site, the next step in the site configuration process is establishing connections between the sites. The physical connectivity between the sites is established between the Active Directory databases by site link objects. A *site link object* is an Active Directory object that embodies a set of sites that can communicate at uniform cost. A *site link* can consist of two or more sites. Because a site link joins two or more sites with a uniform cost and replication schedule, they are used to determine the efficiency and direction of replication traffic throughout the Active Directory topology. Each site link is based on the following four components:

- **Transport** The networking technology to move the replication traffic.
- **Sites** The sites that the site link connects.
- **Cost** The value to calculate the site links by comparing to others, in terms of speed and reliability charges.
- **Schedule** The times and frequency at which the replication will occur.

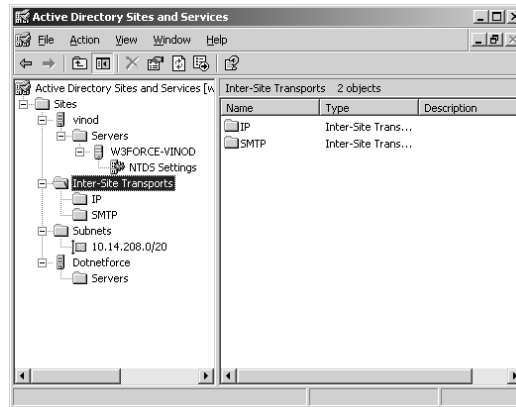
Site links are created using **the Active Directory Sites and Services** tool of Windows Server 2003. Exercise 6.05 walks you through the steps involved in creating site links.

EXERCISE 6.05

CREATING SITE LINKS

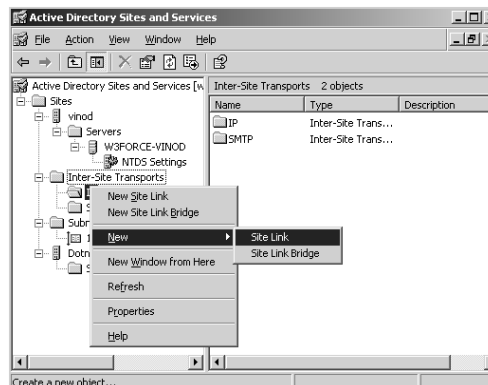
1. To open the **Active Directory Sites and Services** tool, click **Start | Control Panel | Administrative Tools**, and then double-click **Active Directory Sites and Services**.
2. Highlight the **Inter-Site Transports** folder in the left tree pane of the **Active Directory Sites and Services** console. Expand the **Inter-Site Transports** folder as shown in Figure 6.21.

Figure 6.21 The Inter-Site Transports Folder



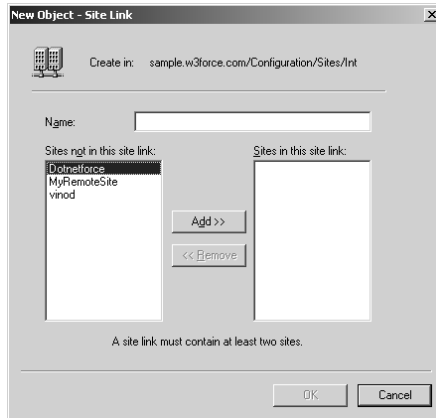
3. Right-click either the **IP** or **SMTP** folder (depending on what protocol the network is based on) in the left tree pane of the **Active Directory Sites and Services** console. Select **New Site Link** from the context menu as shown in Figure 6.22.

Figure 6.22 The New Site Link Option



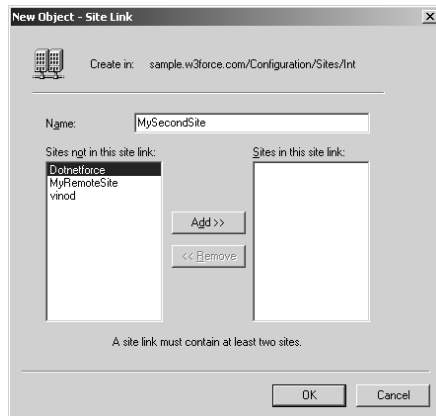
4. Selecting **New Site Link** option opens a **New Object – Site Link** dialog box as shown in Figure 6.23.

Figure 6.23 The New Object – Site Link Dialog Box

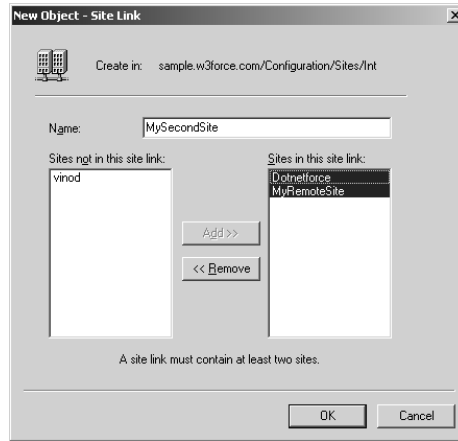


5. Type the name of the new site link object in the **Name** box in the **New Object – Site Link** dialog box as shown in Figure 6.24.

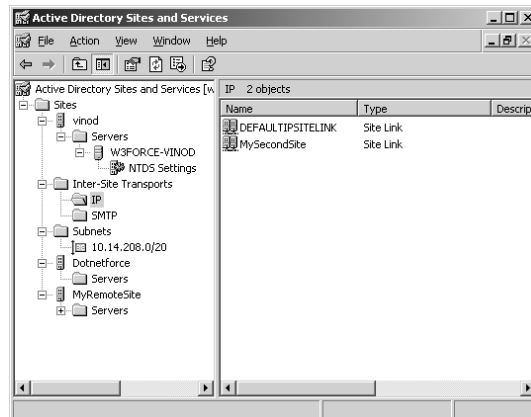
Figure 6.24 The Name of the New Site Link Object



6. Select two or more sites for establishing connection from the **Sites not in this site link** box, and click **Add** as shown in Figure 6.25.

Figure 6.25 Selecting Sites to Establish Connection

7. Click **OK**. This completes the process of creating a new site link object using the Active Directory Sites and Services tool. Figure 6.26 shows the final screen shot of the process.

Figure 6.26 ADSS Tool After Creating the New Site Link

EXAM
70-294
OBJECTIVE
1.4.1
2.2.2

Configuring Site Link Cost

Site link costs are calculated to determine how expensive an organization considers the network connection between two sites that the site link is connecting.

Higher costs represent more expensive connections. If there are two site links available between two sites, the lowest cost site link will be chosen. Each site link is assigned an IP or

SMTP transport protocol, a cost, a replication frequency, and an availability schedule. All these parameters reflect the characteristics of the physical network connection.

The cost assigned to a site link is a number on an arbitrary scale that should reflect, in some sense, the expense of transmitting traffic using that link. Cost can be in the range of 1 to 32,767, and lower costs are preferred. The cost of a link should be inversely proportional to the effective bandwidth of a network connection between sites. For example, if you assign a cost of 32,000 to a 64 kbps line, then you should assign 16,000 to a 128 kbps line and 1000 to a 2 Mbps line. It makes sense to use a high number for the slowest link in your organization. As technology improves and communication becomes cheaper, it's likely that future WAN lines will be faster than today's, so there's little sense in assigning a cost of two for your current 128 kbps line and a cost of 1 for your 256 kbps line, because these low link costs leave no room for future flexibility should you upgrade to quicker links.

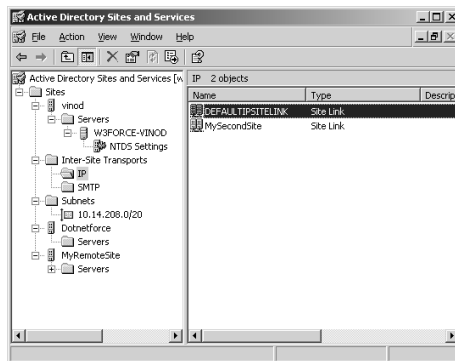
Site link costs are configured using the Active Directory Sites and Services tool of Windows Server 2003. Exercise 6.06 illustrates the steps involved in creating site link costs.

EXERCISE 6.06

CONFIGURING SITE LINK COSTS

1. To open the **Active Directory Sites and Services** tool, click **Start | Control Panel | Administrative Tools**, and then double-click **Active Directory Sites and Services**.
2. Highlight the **Sites** folder in the left tree pane of the **Active Directory Sites and Services** console and expand the **Sites** folder.
3. Highlight the **Inter-Site Transports** folder in the left tree pane of the **Active Directory Sites and Services** console and expand the **Inter-Site Transports** folder as shown in Figure 6.27.

Figure 6.27 The Inter-Site Transports Folder



- Right-click the site link whose cost you want to configure in the left tree pane of the **Active Directory Sites and Services** console, and select **Properties** as shown in Figure 6.28. Selecting **Properties** opens a dialog box as shown in Figure 6.29.

Figure 6.28 The Properties Option

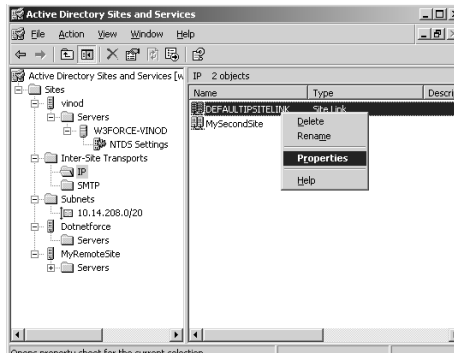
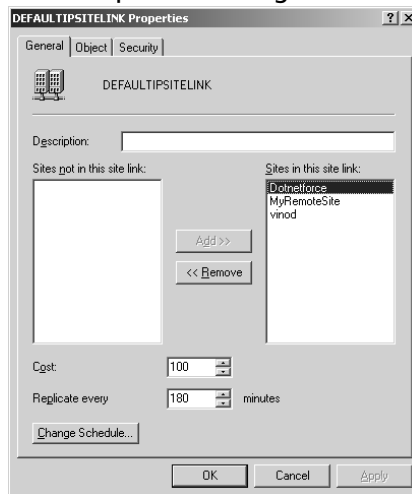
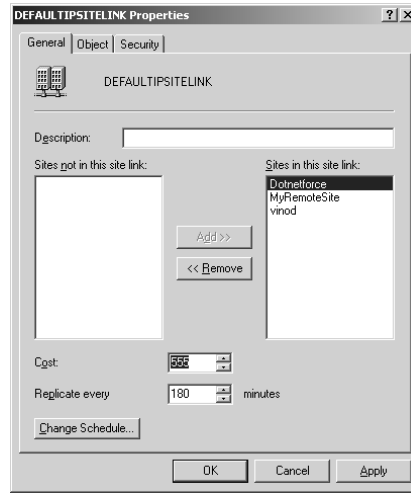


Figure 6.29 The Properties Dialog Box



- Type the value for the cost of replication of the site link object in the **Cost** box in the dialog box as shown in Figure 6.30.

Figure 6.30 The Cost of the Site Link Object



6. Click **OK**. This completes the process of configuring the site link costs using the Active Directory Sites and Services tool.

EXAM
70-294
OBJECTIVE
2.2
2.2.1
2.5.1

Understanding Site Replication

An essential process for any domain that has multiple DCs is *replication*. Replication ensures that each copy of the domain data is up to date, and is done by sending information about changes from one DC to another.

Every Windows domain typically has multiple servers that act as DCs. In Windows Server 2003, every DC is capable of making changes to the database that has domain user and computer accounts.

Purpose of Replication

Earlier versions of NT were configured in a single-master environment where the PDC was used to maintain and manage the master copy of the domain database, and was also in charge of replicating changes to the BDCs. In a *single-master environment*, if for some reason the PDC is unavailable, no changes can be made to the database.

In Windows Server 2003 domains, every DC has a complete copy of the domain partition for its own domain. This is similar to the NT model, but the difference is that each Windows Server 2003 DC first accepts and makes changes to the database and then replicates those changes to other DCs. An environment in which multiple computers are used for managing changes is known as a *multi-master environment*.

A multi-master environment has many advantages over the single-master configuration, including the following:

- There are no single points of failure, as every DC can accept changes to the database.
- DCs that accept changes to the database are distributed throughout the network. This allows administrators to make changes on local DCs and let the replication ensure that these changes are updated to all other DCs in an efficient manner.

Replication depends on a harmonized environment. Every DC involved must be a Windows Server 2003 server and should have identical schemas, and there must be a high level of trust between the servers involved.

Types of Replication

Replication in a Windows Server 2003 environment is one of two types:

- **Intrasite replication** Replication that occurs between DCs within a site.
- **Intersite replication** Replication that occurs between DCs in different sites.


It is important to understand the differences between these methods when planning the site structure and replication.

Intrasite Replication

Intrasite replication occurs between DCs within a site. The system implementing such replication uses high-speed, synchronous Remote Procedure Calls (RPCs).

Within a site, a ring topology is created by the KCC between the DCs for replication (see Figure 6.31). The **KCC** is a built-in process that runs on all DCs and helps in creating replication topology. It runs every 15 minute by default and delegates the replication path between DCs based on the connection available. The KCC automatically creates replication connections between DCs within the site. The ring topology created by the KCC defines the path through which changes flow within the site. All the changes follow the ring until every DC receives them.

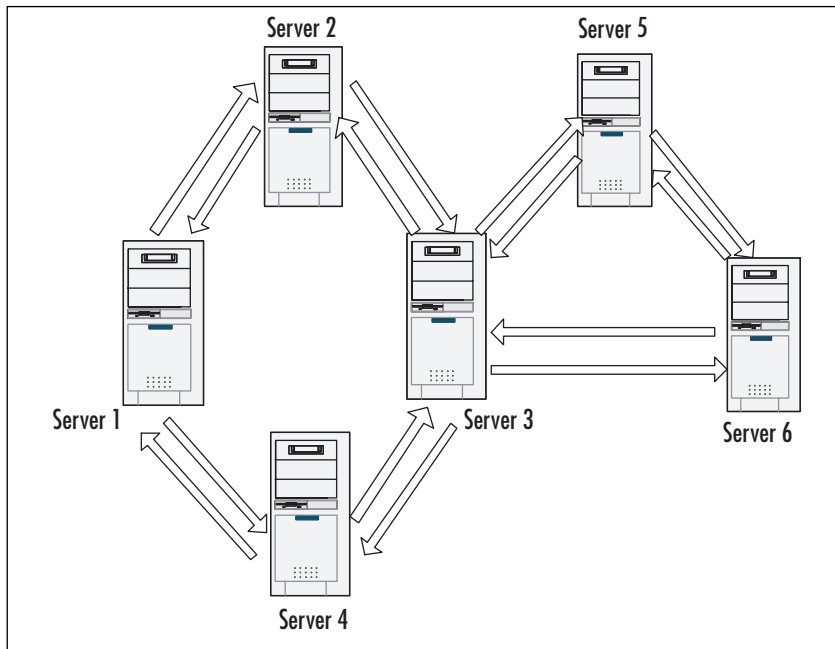
The KCC analyzes the replication topology within a site to ensure efficiency. If a DC is added or removed, it reconfigures the ring for maximum efficiency. It also configures the ring so that there will be not more than three hops between any two DCs within the site, which sometimes results in the creation of multiple rings (see Figure 6.32).

Figure 6.31 Ring Topology for Replication

Intersite Replication

Intersite *replication* takes place between DCs in different sites. The drawback of intersite communication is that it has to be configured manually. Active Directory builds an efficient intersite replication topology with the information provided by the user. The directory saves this information as site link objects. A DC running a service called the Inter-site Topology Generator is used to build the topology. An Inter-site Topology Generator is an Active Directory process that runs on one DC in a site and considers the cost of intersite connections. It ensures that the previous DCs are no longer available, and checks to determine if new DCs have been added. The KCC process updates the intersite replication topology. A least-cost spanning-tree algorithm is used to eliminate superfluous replication paths between sites.

An intersite replication topology is updated regularly to respond to any changes that occur in the network. It would be useful if the traffic needs to cross a slower Internet link.

Figure 6.32 The Three-Hop Rule of Intrasite Replication

An intersite replication across site links occurs every 180 minutes; this can be changed if necessary. In addition, you can schedule the availability of the site links for use. By default, a site link is accessible to carry replication 24 hours a day, 7 days a week, and this can also be changed if necessary. A site link can also be configured to use low-speed synchronous RPCs over TCP/IP or asynchronous SMTP transport. That is, replication within a site always uses RPC over IP, while replication between sites can use either RPC over IP or SMTP over IP. Replication between sites over SMTP is supported for only DCs of different domains. DCs of the same domain must replicate by using the RPC over IP transport. Hence, a site link can be configured to point-to-point, low-speed synchronous RPC over IP between sites, and low-speed asynchronous SMTP between sites

EXAM
70-294
OBJECTIVE
1.4

Planning, Creating, and Managing the Replication Topology

An important job when implementing replication topology is planning, creating, and managing the replication topology, as discussed next.

Planning Replication Topology

Let's now discuss how to plan a replication topology:

- Before starting a replication planning process, we need to first finish the forest, domain, and DNS.

- It is essential to have an understanding of Active Directory replication and the File Replication Service (FRS), which is used to replicate the contents of the system volume (SYSVOL) folder that contains GPO objects, logon and logoff scripts, and startup and shutdown scripts. FRS is also used for Distributed File System (DFS) replication.
- For Active Directory replication, a rule of thumb is that a given DC that acts as a bridgehead server should not have more than 50 active simultaneous replication connections at any given time.

Creating a Replication Topology

The next step is to create the replication topology. Let's discuss how to create a replication topology:

- Active Directory replication is a one-way *pull* replication whereby the DC that needs updates (target DC) gets in touch with the replication partner (source DC). Then, the source DC selects the updates that the target DC needs, and copies them to the target DC. Because Active Directory uses a multi-master replication model, each DC functions as both source and target for its replication partners. From the view of a DC, it has both inbound and outbound replication traffic, depending on whether it is the source or the destination of a replication sequence.
- Inbound replication is the incoming data transfer from a replication partner to a DC, while outbound replication is the data transfer from a DC to its replication partner.
- System policies and logon scripts that are stored in SYSVOL use FRS to replicate. Each DC keeps a copy of SYSVOL for network clients to access. FRS is also used for the Distributed File System (DFS).
- Components of the replication topology such as the KCC, connection objects, site links, and site link bridges are to be checked by the administrator.
- There are two methods for creating a replication topology:
 - Use the KCC to create connection objects. This method is recommended if there are 100 or fewer sites.
 - Use a scripted or third-party tool for the creation of connection objects. This method is recommended if there are more than 100 sites.

Managing Replication Topology

Data is usually replicated based on a change notification within sites. It's up to the administrator to force immediate replication. To do so for all data on a given connection in a single direction, perform the following steps:

1. Choose **Start | Programs | Administrative Tools | Active Directory Sites and Services**. Expand **Sites** in the left tree pane.

2. Expand the name of the site that has to replicate to.
3. Expand the name of the server for replicating.
4. Select the server's **NTDS Settings** object. The right console pane will be populated with the server's inbound connection objects.
5. In the right pane, right-click the name of the server from which you want to replicate, and select **Replicate Now**.

Replication can also be forced from the command line by using the *repadmin.exe* utility from the Support Tools.

Configuring Replication between Sites

To ensure that users can log on within a given span of time, it is necessary to locate DCs near them, which sometimes involves moving the DCs between sites.

The purpose of a site is to help manage the replication between DCs and across slow network links. In addition to creating the site and adding subnets to that site, we also need to move DCs into the site, as replication happens between DCs. The DC has to be added to a site to which it belongs so that clients within a site can look for the DCs in the site and can log on to it.

To move DCs, follow these steps:

1. Select **Click Active Directory Sites and Services**.
2. Choose the **Sites** folder and then select the site where the server is located.
3. In the site, expand the **Servers** folder.
4. Right-click on the DC you want to move, and choose **Move**.
5. Select the destination subnet from the dialog box and click **OK**.

Configuring Replication Frequency

Replication frequency can be configured by providing an integer value that informs the Active Directory as to how many minutes it should wait before it can use a connection to check replication updates. The interval of time must be not less than 15 minutes and not more than 10,080 minutes. For any replication to happen, a site link is essential. Follow these steps to configure site link replication frequency:

1. Choose **Start | Programs | Administrative Tools | Active Directory Sites and Services**.
2. Expand the **Inter-Site Transports** folder, select either the **IP** or **SMTP** folder, and then right-click the site link for which the site replication frequency is to be set.
3. Click **Properties**, and in the Properties dialog box for the site link, enter in the **Replicate Every** box the number of minutes between replications. The default value is 180.
4. Click **OK**.

Configuring Site Link Availability

After the DCs are moved, a site link has to be created between sites, as it provides a path through which replication takes place. The creation of site links gives the KCC information about which connection object should be created in order to replicate directory data. Site links also imply where the connection object should be created. Follow these steps to configure a site link:

1. Choose **Start | Programs | Administrative Tools | Active Directory Sites and Services**.
2. Open the **Sites** folder and then the **Inter-Site Transports**.
3. Right-click on the **IP** or **SMTP** folder depending on the protocol needed and then choose **New Site Link**.
4. Enter the name for the site link in the **Name** text box. From the **Sites not in this site link** list, choose the site to connect and click **Add**.
5. Click **OK**.

When creating site links, there is the option of using either IP or SMTP as the transport protocol:

- **SMTP replication** SMTP can be used only for replication over site links. It is asynchronous; that is, the destination DC does not wait for the reply, so the reply is not received in a short amount of time. SMTP replication also neglects Replication Available and Replication Not Available settings on the site link schedule, and uses the replication interval to indicate how often the server requests changes. When choosing SMTP, you must install and configure an Enterprise Certification Authority (ECA) because Public Key encryption and certificates are used to verify identity of domain controllers and provide digital signatures.
- **IP replication** All replication within a site occurs over synchronous RPC over IP transport. The replication within a site is fast and has uncompressed delivery of updates. Replication events occur more frequently within a site than between sites, and the overhead of compression would be inefficient over fast connections.

Configuring Site Link Bridges


Often, there is no need to deal with site link bridges separately, as all the links are automatically bridged by a property known as a *transitive site link*. Sometimes when you need to control through which sites the data can flow, you need to create site link bridges. By default, all the site links created are bridged together.

The bridging enables the sites to communicate with each other. If this is not enabled by the automatic bridging due to the network structure, disable the same and create an appropriate site link bridge. In some cases, it is necessary to control the data flow through the sites. In these cases, it is necessary to create site link bridges. To disable transitive site links (automatic bridging), follow these steps:

1. Choose **Start | Programs | Administrative Tools | Active Directory Sites and Services**.
2. Expand the **Sites** folder and then expand the **Inter-Site Transports** folder.
3. Right-click on the transport for which the automatic bridging should be turned off, and choose **Properties**.
4. On the **General** tab, clear the **Bridge all site links** check box and click **OK**.

To create a site link bridge, follow these steps:

1. Choose **Start | Programs | Administrative Tools | Active Directory Sites and Services**.
2. Expand the **Sites** folder and then the **Inter-Site Transports** folder.
3. Right-click on the transport that needs to be used, and choose **New Site Link Bridge**.
4. In the **Name** box, enter a name for the site link bridge.
5. From the list of **Site links not in this bridge**, select the site link to be added.
6. Remove any extra site links in the **Site links in this bridge** box and click **OK**.




Configuring Bridgehead Servers

A *bridgehead server* is a server that is mainly used for intersite replication. The bridgehead server can be configured for every site that is created for each of the intersite replication protocols. This helps to control the server that is used to replicate information to other servers.

To configure a server as a bridgehead server, follow these steps:

1. Choose **Start | Programs | Administrative Tools | Active Directory Sites and Services**.
2. Expand the **Sites** folder.
3. Expand the site in which a bridgehead server has to be created, and then expand the **Servers** folder.
4. Right-click on the server and choose **Properties**.
5. In the **Transports available for inter-site transfer** area, select the protocol for which this server should be a bridgehead and click **Add**.
6. Click **OK** to set the properties, and then close **Active Directory Sites and Services**.

The ability to configure a server as a bridgehead server gives you greater control over the resources used for replication between intersites.



Troubleshooting Replication Failure


DCs usually handle the process involved with replication automatically. Unsuccessful network links and wrong configurations prevent the synchronization of information between

DCs. There are many ways to monitor the behavior of Active Directory replication and correct problems if they occur.

Troubleshooting Replication

A common symptom of replication problems is that the information is not updated on some or all DCs. There are several steps that you can take to troubleshoot Active Directory replication, including:

- **Check the network connectivity** The basic requirement for any type of replication to work properly in a distributed environment is network connectivity. The ideal situation is that all the DCs are connected by high-speed LAN links. In the real world, either a dial-up connection or a slow connection is common. Check to see if the replication topology is set up properly. In addition, confirm if the servers are communicating. Failed dial-up connection attempts can prevent important Active Directory information from being replicated.
- **Examine the replication topology** The Active Directory Sites and Services tool helps to verify whether a replication topology is logically consistent. This is done by right-clicking the **NTDS Settings** within a Server object and selecting **All Tasks | Check Replication Topology**. If there are any errors, a dialog box will alert you to the problem.
- **Validate the event logs** Whenever an error in the replication configuration occurs, events are written to the Directory Service event log. The Event Viewer administrative tool can provide the details associated with any problems in replication.
- **Verify whether the information is synchronized** Many administrators forget to execute manual checks regarding the replication of Active Directory information. One of the reasons for this is that Active Directory DCs have their own read/write copies of the Active Directory database. Therefore, no failures are encountered while creating new objects if connectivity does not exist. It is important to regularly check whether the objects have been synchronized between DCs. The manual check, although tedious, can prevent inconsistencies in the information stored on DCs.
- **Check router and firewall configurations** Firewalls are used to restrict the types of traffic transferred between networks. They increase security by preventing unauthorized users from transferring information. In some cases, company firewalls might block the types of network access that should be available for Active Directory replication to occur.
- **Verify site links** Before any DCs in different sites can communicate, the sites must be connected by site links. If replication between sites doesn't occur properly, verify whether the site links are in proper positions.



Using Replication Monitor

The Replication Monitor tool helps you to determine whether the DCs replicate the Active Directory information correctly. This tool is available as part of the Windows Server 2003 Support Tools, which have to be installed separately. To open the Replication Monitor, we need to install the Support Tools. After installing, go to **Startup menu | Windows Support Tools | Command Prompt** and enter **replmon.exe**, which will open the Replication Monitor console.

When you open the tool, you see a blank screen that is divided into two panels: *Monitored Server* and *Log*. Follow these steps to perform replication monitoring:

1. Select the **Add Monitored Server** option from the **Edit** menu.
2. Enter the server name (if known) to be monitored, or search for a specific domain for a server to monitor. After this is done, the **Monitored Server** panel displays the Active Directory information, and the log panel shows the information stored in the log file.
3. To save the log information, select the **Save Monitored List As** and **Open Log** options from the **File** menu.
4. The **Active Directory Replication Monitor** tool can also be used for synchronizing the directory partition. DCs listed for a directory partition are treated as source servers, while the direct replication partners are represented by an icon that indicates the network-connected servers. Right-clicking a server and selecting **Properties** can also identify it. The **Properties** box displays the source server as a Direct Replication Partner, a Transitive Replication Partner, or a Bridge Head Connection.
5. Right-click the direct replication partner, and select **Synchronize Replica**. **replmon.exe** initiates replication and reports the success or failure of the request.
6. Apart from these functionalities, the **Replication Monitor** tool has various options under different menus, such as **Action**, **View**, and so forth. Under the **Action** menu you have different options. For example, under the **Domain** option you can select **Search Domain Controller**, which is used for replication errors. There is a **Server** option that is basically used for replication-related work and helps to check the replication topology.
7. Apart from these submenus, there are options such as **Site**, **Naming Context**, and **Replication Partners** that are enabled when the appropriate function for a server is selected.

The Active Directory Replication Monitor is simple and easy to use. It provides a great deal of information and is useful for fixing Active Directory replication problems.

Using Event Viewer

The Event Viewer is used for configuring Active Directory event logging. To configure Active Directory event logging, follow these steps:

1. Select **Start | Run**. In the **Open** box, type **regedit**, and click **OK**.

2. Locate and click the following Registry key: **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics**
3. Each entry in the right pane of the Registry Editor window represents a type of event that Active Directory can log. All entries are set to the default value of 0 (None).

To configure event logging for the appropriate component, follow these steps:

1. In the right pane of the Registry Editor, double-click the entry that represents the type of event that is to be logged; for example, **Security Events**.
2. Type the logging level that's needed in the **Value data** box, and click **OK**.
3. Repeat step 2 for each component that you want to be logged. Then, on the **Registry** menu, click **Exit** to quit the Registry Editor.

Some of the events that can be written to the event log include:

- KCC
- MAPI events
- Security events
- Replication events
- Directory access
- Internal configuration
- Internal processing
- Intersite messaging
- Service control setup

Each entry is assigned a value of 0 through 5, which determines the level of details of the events that are logged:

- **0 (None)**: Only critical events and error events are logged at this level. This is the default setting for all entries.
- **1 (Minimal)**: Very high-level events are recorded in the event log at this setting. Events can include one message for each major task that is performed by the service. This can be used when the location to start an investigation is not known.
- **2 (Basic)**
- **3 (Extensive)**: This level records more detailed information than the lower levels, such as steps that are performed to complete a task.
- **4 (Verbose)**
- **5 (Internal)**: This level logs all events, including debug strings and configuration changes. A complete log of the service is recorded.

Using Support Tools

The Support Tools must be installed separately from the Windows Server 2003 operating system. Table 6.2 lists some of the support tools that are used most frequently.

Table 6.2 Categorizing Support Tools

Tool	Description
Repadmin.exe: Replication Diagnostics Tool	A command-line interface that is used for Active Directory replication. This tool provides a powerful interface into the inner workings of Active Directory replication, and is useful in troubleshooting Active Directory replication problems.
Active Directory Replication Monitor (Replmon.exe)	Used to display replication topology, status, and performance of Active Directory DCs.
ADSI Edit	MMC snap-in that acts as a low-level editor for Active Directory.
Browstat.exe: Browser Status	A network browser diagnostic tool.
Dsacls.exe	Helps management of ACLs for directory services.
Dsastat.exe: Active Directory Diagnostic Tool	Compares and detects differences between naming contexts on DCs.

EXAM
70-294
OBJECTIVE
2.3.2

Monitoring File Replication Service Replication

The File Replication Service (FRS) is used to synchronize the contents of the Sysvol folders, located by default in <systemroot>\sysvol\sysvol, across multiple Active Directory domain controllers (FRS is also used in keeping Dfs replicas synched, but in the context of this book, we will be discussing FRS synchronization of Sysvol).

FRS runs as a service. It is installed by default on domain controllers and startup mode is automatic. The executable, ntfrs.exe, is located in the <systemroot>\system32 directory. FRS replication differs from Active Directory replication of changes to user, group and computer accounts in that its intrasite replication is almost instantaneous (By default, regular AD replication occurs at five minute intervals). Intersite replication mirrors that of AD; it occurs at whatever interval the administrator has set for intersite AD replication. Because FRS uses a multi-master replication model, changes can be made to the SysVol on any domain controller and then propagated to the SysVols on other DCs in the domain. These other DCs are called replication partners, and can be either inbound or outbound. Inbound partners are those that provide data to the DC; outbound partners are those to which data is provided by the DC. Tables in the FRS database, called the inbound and outbound logs, hold the change orders that come from inbound replica partners and those to be sent to outbound replica partners, respectively. There are several ways to monitor FRS: log files, FRS node in Event Viewer, and Tools and Scripts.

Examining the Log Files

FRS writes logs to the <systemroot>\Debug folder. These are text files that can be used in troubleshooting problems with FRS. The default files are named Ntfrs_0001 through Ntfs_0005.

By default there are five files saved; however, you can increase the number of log files. You can also save the files to a different location to archive them. Changing the configuration for the log files requires that you edit the registry entries in the following key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NtFrs\
Parameters.
```

Edit the data value for the **Debug Log Files** entry to change the number of log files that will be created. (If you use the `QA_check.cmd` script, you'll also need to edit its file and change the value from 5 to the new number, to correspond to the change in the registry). Edit the data value for the **Debug Log Severity** entry to specify how much detail you want recorded in the logs (0 is least amount of detail; 5 is most). Edit the value for the **Debug Maximum Log Messages** entry to control the number of lines (and thus the file size) of each log. The default is 10,000 lines, which results in a file of approximately 1MB in size.

The log files record errors, warning messages and other events related to the service.

For more information on how to analyze the contents of the FRS logs, see <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/ad/windows2000/deploy/adguide/addeploy/addch09.asp>. Although this TechNet article was written for Windows 2000, it is applicable to Server 2003 FRS log files as well.

Using the FRS Node in Event Viewer

A much more user-friendly way to monitor FRS events is through the File Replication Service node of the event viewer. This node appears only on domain controllers, and contains the familiar Event Viewer Information, Warning and Error messages.

Information events record occurrences that are not problems, such as notification that the service is starting or stopping. Warnings record events that could potentially cause problems, such as the FRS service being unable to bind to a domain controller when polling for FRS replica set configuration information. Errors record events that present existing problems.

Using Tools and Scripts

There are some tools and scripts provided by Microsoft that can be used to monitor FRS replication. We've already mentioned the `QA_check.cmd` script. You can also use the following utilities and scripts to monitor FRS:

- **The Ntfrsutl.exe tool** This utility can be used to show the ID table and inbound and outbound logs for FRS computers, examine how much memory FRS uses, view the FRS configuration, get a list of active replica sets, find out the API and version number of FRS running on the computer, and poll for changes to FRS configuration. You must be an administrator to access all of the data provided by this utility.
- **Connstat.cmd** This is a script that creates a summary of the domain controller's FRS connections and shows the states of both the inbound and outbound connections.
- **Frscheck.cmd** This is a script that uses `Ntfrsutl.exe` to collect information about the FRS version, sets, directory services, configurable, inlog and outlog, as well as registry parameters related to FRS. It also looks for error information in the Debug log files. Again, you need to be an administrator to access all of its information.

Summary of Exam Objectives

This chapter explained the role of sites, and discussed the relationship of sites to other Active Directory components. We showed you how to create sites and site links, and explained site replication. This chapter enables you to become familiar with exam objectives covering such topics as the various roles and services offered by Active Directory Sites.

The chapter also offered the basics of how each role and services of Active Directory Sites works and how Active Directory Sites can be used efficiently in terms of data transmission as part of a large network. You should be familiar with the difference between interactive and network authentication and be aware that using Active Directory sites can help to ensure that users and computers will communicate with the appropriate DCs for authentication and authorization.

Sites and subnets are essential components that make up the physical structure of your Active Directory network. Understanding how to use them for enhanced performance and best bandwidth usage is necessary, not only in mastering the objectives for Exam 70-294, but in performing your daily duties as a Windows Server 2003 network administrator.

Exam Objectives Fast Track

Understanding the Role of Sites

- ☑ Sites are used for optimizing the authentication process, by reducing authentication traffic across slow, high-cost WAN links.
- ☑ Subnets provide rapid and reliable communication between locations.
- ☑ The primary role of sites is to increase the performance of a network, which is achieved by economic and rapid transmission of data.
- ☑ Replication enables transferring data from a data store present on a source computer to an identical data store present on a destination computer.
- ☑ The Knowledge Consistency Checker (KCC) is a process that runs on a DC.
- ☑ Authentication is a process by which a system validates users, using the logon information provided.
- ☑ Network authentication verifies the user's identification to a network service.

Relationship of Sites to Other Active Directory Components

- ☑ The distribution of services information in Active Directory enables the client and applications to get information from the directory.

- ☑ Domains are defined as units of replication that receive changes and replicate those changes among all DCs present in the domain of a network.
- ☑ A single site consists of one or more Internet Protocol (IP) subnets.
- ☑ Sites and subnets are represented in Active Directory by site and subnet objects, which we create through Active Directory Sites and Services.

Creating Sites and Site Links

- ☑ Site planning enables you to optimize the efficiency of the network and reduce administrative overhead.
- ☑ Establishing separate site topology enables all replication to occur as intrasite replication, which requires no manual replication configuration.
- ☑ The renaming of a site is a primary task that should be performed when administering a site structure.
- ☑ When creating a site initially, it is created with a default name Default-First-Site-Name.
- ☑ The process of associating a subnet with a site notifies Active Directory sites about the physical networks that are represented by the site.
- ☑ Cost is the value used to calculate site links by comparing one to others, in terms of speed and reliability charges.

Understanding Site Replication

- ☑ Replication is a process that ensures that each copy of the domain data is kept contemporary, which is done by sending information about changes from one DC to another.
- ☑ In Windows Server 2003 domains, every DC has a complete copy of the Active Directory of its own domain. It is similar to the NT model, except that each Windows Server 2003 DC first accepts and makes changes to the database and then replicates those changes to other DCs.
- ☑ *Intrasite replication* occurs between DCs within a site. The system implementing such replication uses high-speed, synchronous Remote Procedure Calls (RPCs).
- ☑ The Replication Monitor Tool helps you to determine whether the DCs replicate the Active Directory information correctly.

Exam Objectives Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

- Q:** How can you restrict Active Directory replication traffic to a specific port?
- A:** By default, Active Directory replication through RPCs takes place dynamically over an available port, the RPC Endpoint Mapper, using port 135. An administrator can override this functionality and specify the port through which all replication traffic passes, thereby locking down the port.
- Q:** How do you change the time the KCC runs?
- A:** The KCC, which manages connection objects for inter- and intrasite replication, runs every 15 minutes by default. To change this, start **regedit** and go to the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Parameters** Registry entry. Then, from the **Edit** menu, select **New, DWORD Value**.
- Q:** How do I move a server to a different site?
- A:** If the sites and subnets are configured, new servers are automatically added to the site that owns the subnet. However, a server can be manually moved to a different site. To perform this task, start the **Active Directory Sites and Services**. Expand the site that currently contains the server, and expand the **Servers** container. Right-click the server and select **Move** from the context menu. There will be a list of all the sites. Select the new target site, and click **OK**.
- Q:** How can a server belong to more than one site?
- A:** By default, a server belongs to only one site. However, you can configure a server to belong to multiple sites. Because sites are necessary for replication, for clients to find resources, and to decrease traffic on intersite connections, simply modifying a site's membership might cause performance problems. To configure a server for multiple site membership, log on to the server you want to join multiple sites. Start **regedit** or **regedt32**. Go to the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters** Registry entry, select **Add Value** from the **Edit** menu, enter the name **Site Coverage** and as a **REG_MULTI_SZ** value, and click **OK**. Next, enter the names of the sites to join, each

on a new line. (Press **Shift + Enter** to move to the next line.) Click **OK**. Close the Registry Editor.

Q: How do I disable site link transitivity?

A: Site links are bridged together to make them transitive so that the KCC can create connection objects between DCs. We can disable site link transitivity manually by bridging specific site links. Start the **Active Directory Sites and Services** snap-in. (Select **Programs | Administrative Tools | Active Directory Sites and Services** from the **Start** menu.) Expand the **Sites** folder and expand the **Inter-Site Transports** folder. Right-click the protocol for which you want to disable transitivity (IP or SMTP), and select **Properties**. Clear the **Bridge all site links** check box, and click **Apply**.

Q: How do you rename a site?

A: When you install your first DC, the DC creates the default site Default-First-Site-Name. This name isn't very descriptive, so you might want to rename it. Start the **Active Directory Sites and Services** snap-in. (Select **Programs | Administrative Tools | Active Directory Sites and Services** from the **Start** menu.) Expand the **Sites** folder. Right-click the site that is to be renamed (for example, Default-First-Site-Name), and select **Rename**. Enter the new name, and press **Enter**.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Understanding the Role of Sites

1. An Active Directory environment has been configured with multiple sites and has appropriate resources in each site. The administrator of the Active Directory environment tries to choose a protocol for the transfer of replication information between two sites. The connection between the two sites has the following characteristics:
 - The link is unavailable during certain times of the day due to an unreliable network provider.
 - The replication transmission must be carried out whether the link is available or not.
 - Replication traffic must be able to travel over a standard Internet connection.

Which of the following protocols meets these requirements?

- A. Internet Protocol (IP)
 - B. Simple Mail Transfer Protocol (SMTP)
 - C. Remote Procedure Calls (RPCs)
 - D. Dynamic Host Configuration Protocol (DHCP)
2. Julie installs a Windows Server 2003 server that will be used during the installation of the Active Directory structure for her organization. She installs the DNS server, creates the domain, and configures it for dynamic updates. When she attempts to install the first DC, she gets a message that the DC for the domain is not available. She decides to continue the installation and fix the problem later. What problem will she need to fix later?
- A. The DNS server needs to be restarted.
 - B. The server she is installing needs to point to the DNS server.
 - C. The Active Directory-integrated DNS must be used while installing Active Directory.
 - D. The DNS server needs to be configured for dynamic updates and not to the zones.
3. Robin is managing an Active Directory environment of a medium-sized company. He is troubleshooting a problem with the Active Directory. One of the administrators made an update to a user object and another reported that he had not seen the changes appear on another DC. It was more than a week since the change was made. Robin checks the problem by making a change to another Active Directory object. Within a few hours, the change appears on a few DCs, but not on all of them. Which of the following are possible causes for this problem?
- A. Connection objects are not properly configured.
 - B. Robin has configured one of the DCs for manual updates.
 - C. There might be different DCs for different domains.
 - D. Creation of multiple site links between the sites.

Relationship of Sites to Other Active Directory Components

4. James is a systems administrator for an Active Directory environment that consists of three sites. He wants to set up site links to be transitive. Which of the following Active Directory objects is responsible for representing a transitive relationship between sites?

- A. Additional sites
 - B. Additional site links
 - C. Bridgehead servers
 - D. Site link bridges
5. Michael, a systems administrator of a medium-sized company, suspects that Active Directory replication traffic is consuming a high amount of network bandwidth. He wants to determine the amount of network traffic that is generated through replication. He plans to carry out the following procedures:
- Find out replication data transfer statistics.
 - Find out details on multiple Active Directory DCs at the same time.
 - Find out other performance statistics, such as server CPU utilization.

Which of the following administrative tools is most useful for meeting these requirements?

- A. Active Directory Users and Computers
 - B. Active Directory Domains and Trusts
 - C. Event Viewer
 - D. Performance
6. Steffi is an administrator of a medium-sized organization responsible for managing Active Directory replication traffic. She finds an error in the replication configuration. How can she look for specific error messages related to replication?
- A. Use the Active Directory Sites and Services administrative tool
 - B. Use the Computer Management tool
 - C. View the System log option in Event Viewer
 - D. View the Directory Service log option in Event Viewer

Creating Sites and Site Links

7. George is in charge of managing Active Directory replication traffic for a medium-sized organization that has installed a single Active Directory domain. The current setup is configured with two sites and consists of default settings that are ideal for replication. Each site consists of 20 DCs. Recently, the administrators have found that the Active Directory traffic is using a large amount of available network bandwidth between the two sites. George now has the task of meeting the following requirements:

- Decrease the network traffic between DCs in the two sites.
- Decrease the amount of change to the current site topology.
- Make no changes to the current physical network infrastructure.

George decides that it would be highly efficient to set up specific DCs in each site that will receive the majority of replication traffic from the other site. Which of the following solutions will meet the requirements?

- A. Form additional sites that are intended only for replication traffic, and move the current DCs to these sites.
 - B. Establish multiple site links between the two sites.
 - C. Establish a site link bridge between the two sites.
 - D. Configure one server at each site to act as an ideal bridgehead server.
8. James is in charge of managing the Active Directory environment for a medium-sized organization. He has to write down the procedures for creating a site for a new administrator who is starting up a new office for his organization. Which of the following is the best method for creating a site?
- A. Create the site, select the site link, add the subnets, and then move in the DCs.
 - B. Move the DCs, create the site, add the subnets, and then select the site links.
 - C. Create a temporary site link bridge, add the DCs, rename the site that's created, and then add subnets.
 - D. Create the subnets and then create a site by grouping them. Next, create the links and then move in the DCs.
9. Sofia, an administrator of a medium-sized organization, has created the site links and site link bridges for the Active Directory network. The replication between the sites is working fine, and all the sites are receiving the updates to the Active Directory. She describes the network she is working on to a colleague, and he tells her that she didn't have to configure site link bridges. Why didn't Sofia have to create site link bridges?
- A. The KCC will create the site link bridges for you.
 - B. The sites will be automatically bridged.
 - C. The Domain Naming Master will handle this for you.
 - D. The GC will handle this for you.

Understanding Site Replication

10. Peter, an administrator of an organization, has formed a Windows 2003 Active Directory structure. He has installed a single domain containing 700 users and computers. The organization is split into two offices with a 56 Kbps link between them. Peter creates two sites, one for each office, and a site link between them using SMTP. The replication between the sites doesn't seem to be working. What should Peter do?
 - A. He has to configure an enterprise CA.
 - B. He has to configure Microsoft Exchange.
 - C. He has to configure an SMTP-based mail system.
 - D. He must have a connection faster than 56 Kbps.

11. A company uses a single-master domain model, with resource domains for each of its divisions. It has registered two domains under the names `www.dotnetforce.com` and `www.w3force.com`. In this situation, which Active Directory information will be replicated between DCs in the `dotnetforce.com` and the `w3force.com` domains?
 - A. Domain-naming context
 - B. Schema-naming context
 - C. Configuration-naming context
 - D. GC
 - E. SYSVOL

12. Steffie, an system administrator, has implemented two sites that are connected by a site link. The Cost property is set to 100, and the Replicate Every property is set to 50 minutes. How often will the replication occur?
 - A. Every 5 minutes
 - B. Every 50 minutes
 - C. Every 180 minutes
 - D. The replication frequency cannot be determined.

13. A financial company with branches throughout the United States has hired a consultant to set up the Active Directory sites for their organization. Which of the following structures will he recommend?
 - A. Domain structure
 - B. Political concerns
 - C. Geographic distribution
 - D. Physical network infrastructure

14. James, a network administrator, has configured Active Directory sites. He wants to implement intersite and intrasite replication. Which of the following replication protocols uses RPCs for replication?
- A. DHCP
 - B. RPCs
 - C. IP
 - D. SMTP
15. Your Active Directory structure consists of five domains running in a single forest with 40,000 users. One domain is the Sales domain. Your organization has opened a branch office with 100 employees who are members of the Sales domain. The branch office is connected to the corporate office by a high-speed WAN link. The link is reliable, and you expect the utilization rate of the link to be low. What should you do to minimize Active Directory-related authentication traffic on the WAN link? (Choose all that apply.)
- A. Add the subnet of the branch office to the corporate site.
 - B. Add a DC from the Sales domain to the branch office and configure it as a GC server.
 - C. Add a DC from all five domains to the branch office and configure one DC as a GC server.
 - D. Add a DC for the Sales domain at the branch office.
 - E. Define the branch office as a site.

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|-------------|--------------------|
| 1. B | 9. B |
| 2. B | 10. A |
| 3. A | 11. B, C, D |
| 4. D | 12. B |
| 5. D | 13. D |
| 6. D | 14. C |
| 7. D | 15. B, E |
| 8. A | |

MCSA/MCSE 70-294

Working with Domain Controllers

Exam Objectives in this Chapter:

- 1.3.4 Install and configure an Active Directory domain controller.
- 1.2 Plan flexible operations master role placement.
- 2.5.2 Diagnose and resolve issues related to operations master role failure.
- 1.2.1 Plan for business continuity of operations master roles.
- 1.2.2 Identify operations master role dependencies.

- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

Domain controllers (DCs) in the Windows Server 2003 Active Directory network manage user logon and authentication, store directory data, and are accessed for directory searches. A copy of the Active Directory database resides on each DC, and when you create the first DC for your network by installing Active Directory, this process creates your first forest, domain, and site.

Implementing and managing DCs is an important part of the network administrator's job, because the DCs play such a vital role in the operation of the network. The focus of this chapter is the Active Directory DC, and how to plan and deploy DCs on your network. You'll learn about server roles, where DCs fit in, and how to create and upgrade DCs. We discuss placement of DCs within sites, and how to back up your DCs.

Next, we move to the subject of operations master (OM) roles, and you learn about the functions of all five: the Schema Master, Domain Naming Master, RID Master, PDC Emulator, and Infrastructure Master. We talk about transferring and seizing master roles and role dependencies, and you'll learn to plan for the placement of OMs and how to respond to OM failures.

EXAM
70-294
OBJECTIVE
1.3.4

Planning and Deploying Domain Controllers

Remember that a DC does not equal a domain. A domain is a logical entity containing potentially millions of objects, while a DC, in the context of this chapter, is simply a computer running Windows Server 2003 with a copy of the Active Directory database (of course, an NT Server or Windows 2000 Server computer can also be a DC). This server takes on a management role in granting or denying access to resources throughout the entire domain, not just those resources located on this physical machine. In order to provide acceptable connectivity performance, it is imperative that all users have adequate access to a DC close to their physical locations.



NOTE

"Domain" is a term used by Microsoft to define administrative and security boundaries, and the server role that Microsoft calls a "domain controller" is known as an "authentication server" in generic networking terminology.

Understanding Server Roles

To help you understand DCs, we will first discuss two other types of servers: standalone and member servers. A *standalone server* is a machine with Windows Server 2003 installed that is not a member of a domain. The server might be on the same physical network as your

Windows domain, but for some reason, you decide not to include it in the domain. Why would you do this? Sometimes you have a need for a machine that you want isolated from the rest of your resources by more than just domain restrictions. After all, if you can't authenticate to a standalone server, you can't gain access. Standalone servers require that all access be granted via the local Security Accounts Manager (SAM), which is a security accounts database stored on the local machine. A separate DNS server or test server is a common example of a standalone server installation.

Member servers are machines with which you want to share resources in your domain, so you add them to the domain. They have computer accounts in the domain and can be managed centrally, but do not themselves perform any domain management tasks (although an administrator can install domain administrative tools on a member server and perform those tasks from the member server). Authentication of any domain user or computer can be established centrally rather than locally. This gives you the ability to manage users and computers centrally from the domain. You have the control and it is not distributed in an unmanageable way. Member servers are often called file and print servers or application servers because they are dedicated to those types of functions, and this is the key: member servers are free from any domain management responsibilities. Some common roles played by member servers include:

- E-mail servers
- Web servers
- Name resolution servers (DNS and WINS)
- Name allocation servers (DHCP and BOOTP)
- Terminal servers
- Firewall and proxy servers (such as ISA 2000)
- Remote Access servers (dial-up or VPN)
- Network Address Translation servers (NAT or ICS)

See the *It Depends On...* sidebar for more information about the roles of a Windows Server 2003 server.

You might recall that, before Windows 2000, in order to install a server operating system, you had to decide the server's role during installation. Switching between the roles of primary domain controller (PDC), backup domain controller (BDC) and member or standalone server was difficult and required new installations with each change. Starting with Windows 2000 and continuing with Server 2003, *all* servers begin as standalone or member servers, and then you promote the server to be a DC as needed. In addition, you can now *demote* the DC back to a member server.



NOTE

Where did the BDC go? In Windows NT, there was one and only one PDC. All other DCs were BDCs—second-class citizens in the NT domain. Although they did the same amount of work, they could not be “trusted” to gather changes to the Active Directory and pass them on to other DCs. Now, all DCs are created equal, and we no longer refer to them as backups. Every DC contains a copy of the Active Directory and is trusted to accept the changes given to it by an administrator, and then pass those changes on to every other DC by replication. Of course, this is not quite the whole story. The DCs in Active Directory do have differences, as you will see when we discuss Flexible Single Master Operations (FSMOs) later.

As already mentioned, a server that contains a copy of the Active Directory database is a DC. Why don't you install Active Directory on every server machine? Wouldn't that make authentication very fast? Yes, it would, except when the DC is so busy replicating its changes to every other DC that it clogs the network. A DC has domain responsibilities, and they can interfere with other tasks. Because the Active Directory is the most important part of the domain, your DC will delay your print job or file access until it is finished with its DC duties. Your users, however, don't care about the domain and its needs. Their own needs are more important—to them. Therefore, you should separate file and print access, e-mail and Internet access, and other application-based duties from the DC. Plan your servers according to the needs of the users in your area of stewardship, and balance that with the needs of your domain.

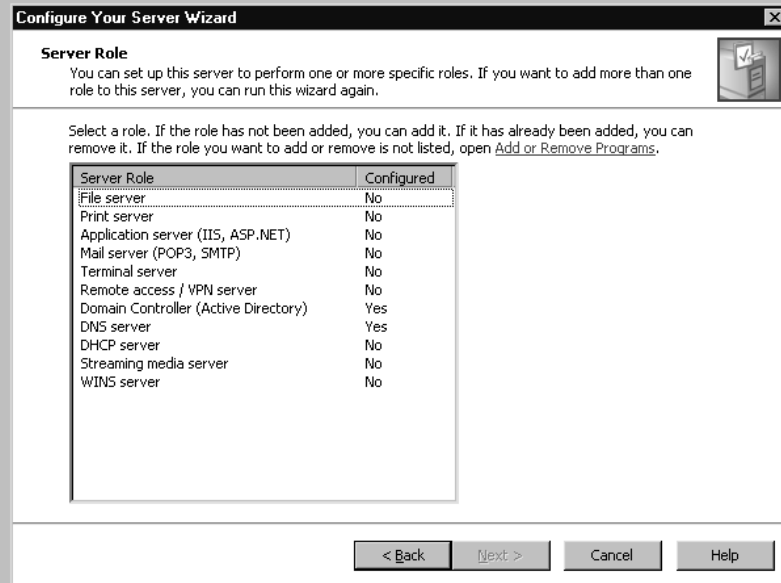
New & Noteworthy...

Configure Your Server Wizard

You are now familiar with the three *standard* server roles: standalone, member, and DC. However, have you seen the other possibilities? Network services such as DNS, DHCP, WINS, and so on can all be part of one server (or separate servers). Using the **Configure Your Server Wizard** facilitates these service installations. Access the wizard by clicking **Start | Programs | Administrative Tools | Configure Your Server Wizard**. The wizard automatically detects your network settings and then provides you with a list of services that can be installed (see Figure 7.1). Although you can assign multiple roles to one server, it is important not to overload your server.

Continued

Figure 7.1 Configure Your Server Wizard: Server Roles



It Depends On...

In economics, most answers begin with the phrase *"It depends on...."* Unfortunately, this applies all too well in the world of technology, too. We can describe individual roles and we can discuss concepts, but as you will see, applying those rules to your domain *depends on* your individual situation. For a domain with only a few users (up to 100), you can combine your network services on one machine. As the number of users and computers grows, however, you will soon run into performance issues. Then, you must decide whether to beef up the hardware (mostly memory) or begin to separate individual services onto dedicated member servers. Knowing when to do this is more art than knowledge. It is true that you must have the knowledge of what each service demands: memory, network, hard drive, or CPU, but balancing and tuning these services is a talent hewn by experience and time. Don't get discouraged. Allow yourself the time to develop these skills. Know that with each overloaded server, you are increasing your ability to handle and prevent such problems.

Function of Domain Controllers

We have alluded to the responsibilities of DCs, and now we will iterate those responsibilities or functions:

- Track all user and computer accounts
- Authenticate access to resources
- Verify passwords
- Establish secure connections
- Replicate all changes to all other DCs

DCs track much of the account information by way of Security Identifiers (SIDs). When a user logs on or a computer boots up, a DC must be found. Typically, the assigned DNS server is contacted with the question, “Where is a domain controller?” The system receives an answer containing one or more DC locations. Requests for authentication are sent to all the DCs. After a DC is found and a secure connection is established using Kerberos security, the user account is authenticated, assuming the correct credentials of a username and password that match one in the domain database are given. A session is established and information regarding this account, including rights and group membership, is stored in memory. This session information is called an *access token*. With each request by an account to access a resource, the resource’s assigned permissions (by groups and user accounts) are compared with the access token to determine the level of access to the resource.



NOTE

Windows NT 4, Windows 2000 and Windows XP workstations authenticate to its corresponding computer account in order to establish a secure session between the workstation and the DC. This is used to increase the security of the information exchanged between the machines.

An important issue to consider in a larger environment that uses wide area network (WAN) connections is where domain changes are made. Suppose that Kim, an administrator in Denver, fills the request to put Jett (a new employee) into the Managers group at the Philadelphia office. The company’s centralized support resides in Denver, so Kim is connected to a DC in Denver. The changes are made on the Denver DC. How long will it be before Jett can use his account to access Manager-level resources? The answer, of course, *depends on* what type of replication policies have been implemented. We will describe what must take place so that you will be aware of the ramifications involved in a scenario like this.

A DC receives changes to its copy of the Active Directory database. By default, all DCs within a site replicate everything to each other within about 15 minutes, or faster if Windows Server 2003 DCs are used in the environment. Between sites, the replication is managed, which is the main reason to create separate sites. If replication is immediate over fast WAN connections, then replication will be as well. If the replication is based on time or activity, the change will have to wait in Denver until the site policy decides to talk to Philadelphia and exchange data. What this means is that it could take 12 to 24 hours before Jett can use his Manager-level access. Kim knows that Jett can't wait this long, so to avoid this, she changes her DC connection to go directly to a DC in Philadelphia, over the WAN. Now the change is accomplished within 15 minutes. Of course, Kim's time to make the change is slowed by that WAN connection, but Jett is much happier!



NOTE

There is always an exception to any rule. In the Kim and Jett scenario, if the change made to Jett's account is a password reset, there is no need for Kim to connect directly to the Philadelphia DC. When Jett logs on with his new password, and the Philadelphia DC does not recognize it, a process is initiated that forces a request across the WAN to determine what the password is. This involves the FSMO roles discussed later in the chapter.

Although the previous list of DC responsibilities is by no means exhaustive, it represents most of the functions you should be concerned about first. Additionally, your DC integrates with other services for ease of administration and security. The following is a short list of some of these services.

- DNS
- DHCP
- Kerberos security
- Remote access
- Virtual private networking

Important to note here is that Active Directory, which is on the DCs, provides these services to give you centralization and control of resources. Making them work efficiently is accomplished by understanding the various services and knowing when to use them and more importantly when *not* to use them. It is easy to turn on the services and let them run, but each service has an effect on the hardware resources involved, which are limited.

Determining the Number of Domain Controllers

Since you just learned about sites in Chapter 6, "Working with Active Directory Sites," you know that each site requires at least one DC. Your site topology is very important, because

of the speed factors (actually the lack of speed) involved in the WAN connections between these sites. You must keep firm control of the replications crossing the WAN. Without sites, you break the age-old rule originally established by Novell: *Don't span the WAN*. The information that follows applies to DCs in *each* site.

Table 7.1 lists the factors you must consider in determining the number of DCs to install.

Table 7.1 Domain Controller Functions Affecting Performance

DC Functions	Description of Effect
PDC Emulator (see FSMOs later in this chapter)	This FSMO is assigned to the first DC installed, and is designed to respond to Windows NT 4 BDCs. Additionally, this FSMO receives all new password and lockout information changes immediately for the entire domain.
Active Directory replication	The process of synchronizing the Active Directory database between DCs.
Workstation logon	Computer accounts authenticating to domain.
Global Catalog (GC) operations	Required in a multidomain Active Directory forest to facilitate logons.
File and print services	A DC can store files and be a print server too.
Network services	A DC can host other important network services such as: DNS DHCP WINS
User logon	User authentication on startup and resource access.
LDAP searches	If you use LDAP applications or services, be aware of this need.
Other FSMOs	(See FSMOs later in this chapter)

Depending on the number of users, computers, and application needs in your domain, you most likely need more than one DC. At the very least, you should have two DCs for fault tolerance in case one goes down. As your network size increases, so will the number of DCs. This facilitates both load balancing and redundancy of the Active Directory. The number of nondomain functions, such as file and print services, will have to go to dedicated member servers as described earlier. Then, how many DCs do you need?

Microsoft has outlined a way to determine this and even created a Job Aid to help you. The first issue is, how much can your server physically handle? Microsoft has issued *minimum* guidelines in Table 7.2 for processors and memory, based on the number of users in the domain and the number of DCs handling the load.

Table 7.2 Minimum DC, GC, RAM, and CPU per Site

Number of Domain Users in Site	Number of DCs	Global Catalog	RAM	CPU
1 to 499	1	DC is a GC server	512MB	Uniprocessor PIII 500+
500 to 999	1	DC is a GC server	1GB	Dual PIII 500+
1,000 to 10,000	2	Both DCs are GC servers	2GB	Quad PIII Xeon+
10,000+	1 for every 5000 users	Half of all DCs are GC servers with a minimum of two GCs	2GB	Quad PIII Xeon+

When the term “bare minimum” is used, remember that it means just that: the *bare* minimum! Your servers should have much more than the minimum if you want more than minimal performance. According to the experts, if you have to choose between CPU and RAM, get more RAM. It’s always easier to get more RAM up front. Although we always say that we can add more RAM later, we often don’t, or when we want to, it is not available for that particular server because it has become obsolete.

How Many Global Catalog Servers Do I Need?

GCs were explained earlier in the book, and you need to know about Microsoft’s minimum requirements. However, according to some Windows experts who are not in the employ of Microsoft, there seems to be no reason to spread these GCs all over the domain, *if* you contain everything in one domain. We recommend doing everything possible to keep your network within one domain. The only reasons to separate it into multiple domains are the following:

- **Political** Sometimes you just have to do it because the check-signers say so.
- **Linguistic** Your company is multilanguage and requires administration to follow these linguistic lines.
- **Security** If different account policies need to be applied.
- **That is how you found it when you got here** This is often the result of merging companies.

Back to GCs. If you do have multidomain forests, GCs are critical; otherwise, it would take too long, if ever, to find a DC or other objects for authentication of

Continued

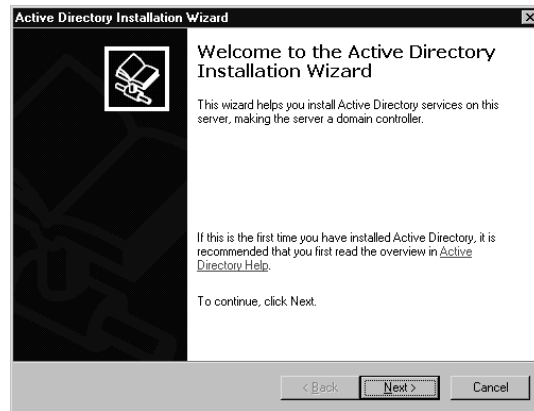
access. As a brief reminder, GCs are an index of the objects in Active Directory. Every domain must have at least one, and it is always installed on the first DC created in a domain. The purpose is to provide a quick search tool to locate objects anywhere in the entire forest. GCs also replicate amongst themselves; and it is a separate process from the Active Directory replication. Therefore, having too many GCs increases this replication, but having too few can increase the queries made by the many services in your domain.

Using the Active Directory Installation Wizard

You know what a DC is, what hardware to buy, how many to buy, and where to put them, and now we will show you how to create one. Microsoft's Active Directory Installation Wizard (ADIW) is used to create DCs, domains, trees, and forests, so you need to understand how to start it and which options to choose.

To start the ADIW, click **Start | Run**. Type **dcpromo** and press **Enter**. Figure 7.2 shows the initial Welcome window with a link to Windows' Help files. Use the Help files if you have to—they are very good. Just click **Next**.

Figure 7.2 Active Directory Installation Wizard Initial Dialog

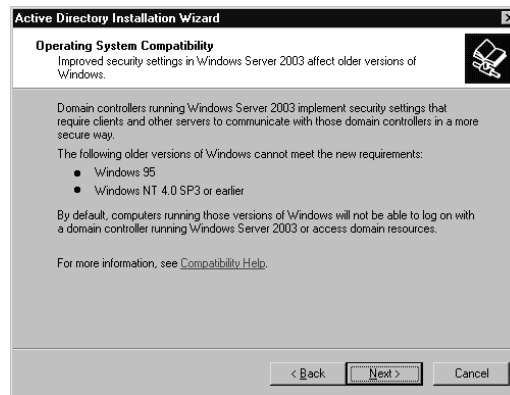


NOTE

You can also start the ADIW from the **Configure your Server Wizard**, by selecting the domain controller role.

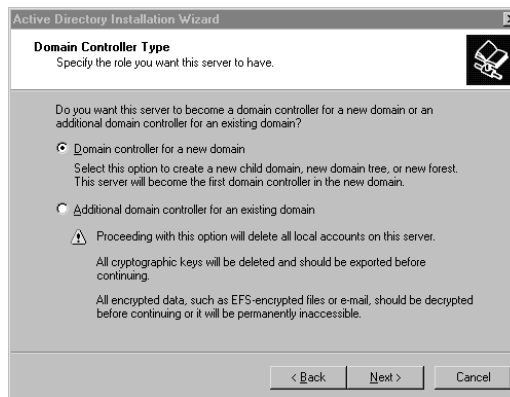
Operating system compatibility is described in the window in Figure 7.3. See the sidebar *Compatibility with Previous Operating Systems* for more detail. Since there is nothing you can do about this issue right now, address it later, if needed. Just click **Next**.

Figure 7.3 Operating System Compatibility



In Figure 7.4, you are given your first window with more than one option. **Additional domain controller for an existing domain** is used to create all other DCs within that same domain. Use this to set up DCs for each site. Selecting this option takes you to a window that requires administrator-level credentials in order to create the DC. The server you are promoting must be able to find another DC via DNS, so make sure you are connected to the network and you have set up your TCP/IP settings to find both the DNS and the DC.

Figure 7.4 Domain Controller Type



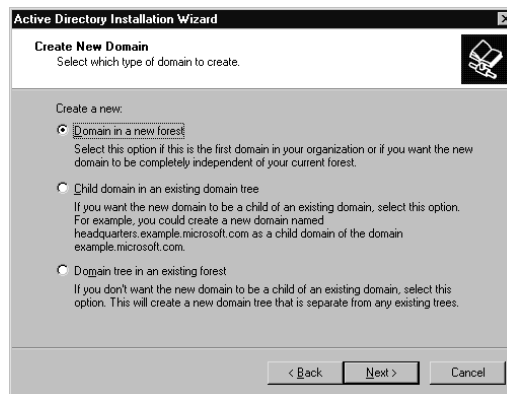
Choosing **Domain controller for a new domain** will make this server a DC, and it will be the first DC in a new domain. Use this for each *new* domain. Following most experts' recommendations, you will only do this once, because a single domain network is the best way to go. Of course, reality dictates that you might have to create additional domains, and this is where you do it.

Three choices are presented to you in Figure 7.5: Create a new:

- **Domain in a new forest** This choice is for the very first DC in your first tree in your first forest.
- **Child domain in an existing domain tree** This choice is used when you already have a domain tree (for example, yourfirm.biz) and you need a second domain or child to this domain (for example, MyPlace.YourFirm.biz).
- **Domain tree in an existing forest** With this option you are sharing the forest and allowing some communication, but you have different tree names. For example, you could have a forest like YourFirm.biz and then add another domain tree that uses a different DNS name, like MyFirm.biz.

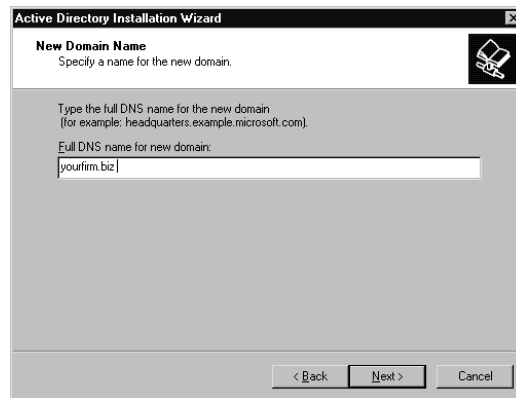
The options of **Child domain in an existing forest** and **Domain tree in an existing forest** require an existing entity to which you are adding. The next window requests administrator credentials at the tree and forest levels. Again, the TCP/IP settings must already be in place in order to find the corresponding DCs to authenticate your credentials and allow you to add on to the tree or forest.

Figure 7.5 Create New Domain



For your first DC, you create the domain, tree, and forest all in the process of creating the one DC, so for our scenario, select **Domain in a new forest** and click **Next**.

Your new domain requires a name. Name your domain using the fully qualified domain name (FQDN). If your domain name is registered with an authorized registrar such as Network Solutions that controls the top-level domains, then use the registered name; otherwise, select a name that suits you and your business, but don't plan on using an unregistered name on the Internet. Usually, a business has an Internet public name that they either have registered or plan to register. Figure 7.6 shows the window and an example of an FQDN.

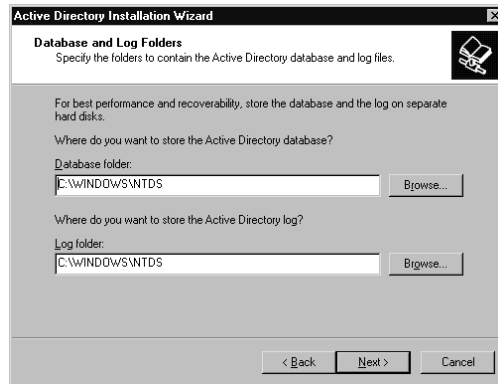
Figure 7.6 New Domain Name

The ADIW next asks you to confirm the NetBIOS domain name to use. Although NetBIOS (and thus WINS) is dying out, it is not dead yet, so you must have this name for your legacy clients. If you are in a pure Windows 2000+ environment, you can disable NetBIOS over TCP/IP (NetBT), but you still have to fill in this information. The default is your registered domain name without the hierarchical identifier of .biz, .com, .us, and so forth. Figure 7.7 depicts the window you see at this stage with the business name, YOUR-FIRM, emblazoned in the field. For our example, accept the default.

Figure 7.7 NetBIOS Domain Name

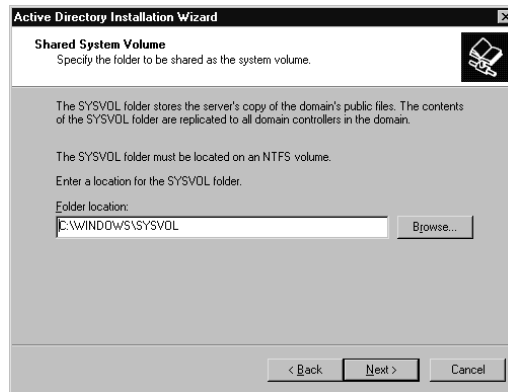
Continuing with the data collection, the ADIW asks where to store the database and log files. It is recommended that you store these on an NTFS partition, and that you store them on separate disk drives, running SCSI, FireWire, or EIDE on separate channels. It does no good to put these files on separate drives if the drives cannot be accessed simultaneously. Figure 7.8 shows both files being stored on the same drive because the author only had one drive available for this exercise. Clicking **Next>** brings up **Shared System Volume** settings.

Figure 7.8 Database and Log Folders



While it is recommended that the Active Directory database and log files be stored on NTFS partitions, the SYSVOL folder *requires* NTFS. SYSVOL exists on every DC and is used to share data between the servers. One such file that should be distributed among all DCs is a logon script. This is much like the NETLOGON share in Windows NT 4. Files placed in this folder *automatically* replicate and synchronize to all other DCs in the domain. It's a nice feature that is especially appreciated by former Windows NT administrators. As shown in Figure 7.9, the default location is C:\WINDOWS\SYSVOL folder. As long as this is on an NTFS volume, go ahead and accept the default.

Figure 7.9 Shared System Volume



Pressing **Next** initiates the DNS Registration Diagnostics; in other words, your server looks for a compliant DNS server on which to create the SRV records and set up the domain structure for dynamic updates. You have three options to choose from if your server is unable to locate said DNS (see Figure 7.10).

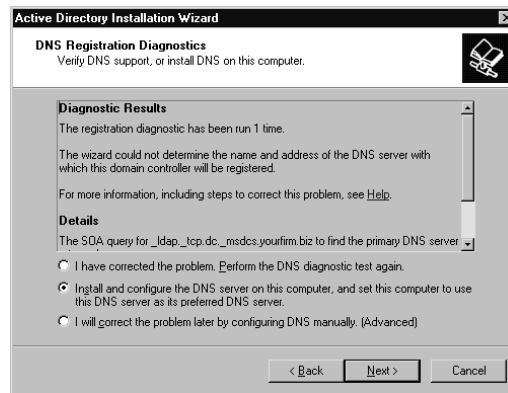
- Correct the problem and try again. A Successful DNS Registration window appears if you successfully register the new domain.
- Make this server a DNS server.
- Correct the problem later.



NOTE

We *highly* recommend that you set up DNS before installing your DC; this can save you many hours of grief and troubleshooting.

Figure 7.10 DNS Registration Diagnostics

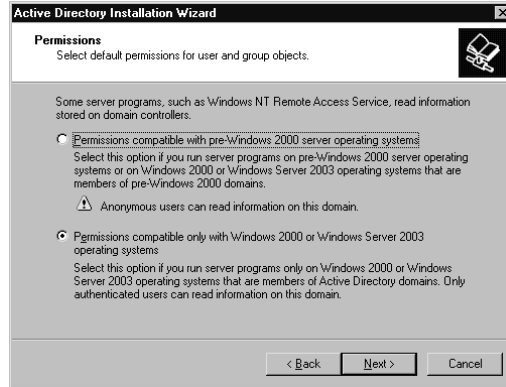


As far as this exercise is concerned, select the default option number **Install and configure the DNS server on this computer to use this DNS server as its preferred DNS server** and let the ADIW create and configure your DNS server. Experts recommend that you go with the option **I have corrected the problem. Perform the DNS diagnostics test again**. This really means exit ADIW, fix the DNS server problem, and run **dcpromo** again. Having your DNS set up correctly before creating the DC prevents extra irritation later when connecting additional DCs and joining workstations and member servers. Option three, **I will correct the problem later by configuring DNS manually**, should be selected only as a last resort.

Now that you have resolved the DNS crisis, you are faced with a permissions dilemma (see Figure 7.11). Some legacy (a.k.a. downlevel), operating systems (pre-Windows 2000) require an anonymous user account to make inquiries to DCs. If you need that “hole” in your security to keep those pre-Windows 2000 server programs running, choose the first option, **Permissions compatible with pre-Windows 2000 server operating systems**. The best example of a service that needs this is the Windows NT 4 Remote Access Service (RAS). If you are not sure if you need this, select the next option, **Permissions Compatible only**

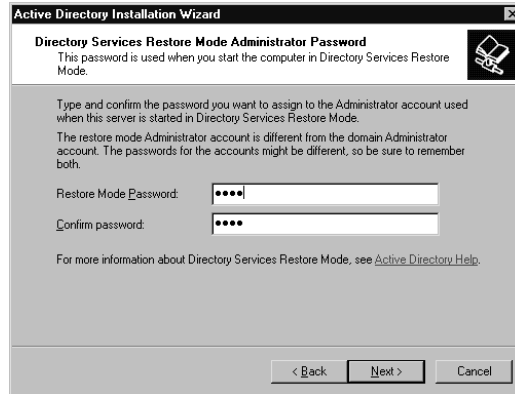
with **Windows 2000 or Windows Server 2003**. You can modify it later by adding the **Everyone** group to the Pre-Windows 2000 Compatible Access group, which resides in the **Builtin** folder in the **Active Directory Users and Computers MMC**.

Figure 7.11 Permissions Compatibility



Select **Permissions compatible only with Windows 2000 or Windows Server 2003** and click **Next** to get the screen in Figure 7.12.

Figure 7.12 Directory Services Restore Mode Administrator Password



You have one last major decision before you let the wizard do its stuff. Windows 2000 and Windows Server 2003 have some restore features that can be selected at boot time. To protect the Active Directory database, you must assign a password to the *restore administrator*. Some administrators use the same domain administrator password they set up originally (see Figure 7.12), but in a production environment, never use a short password as illustrated in the figure, especially for an administrative account.

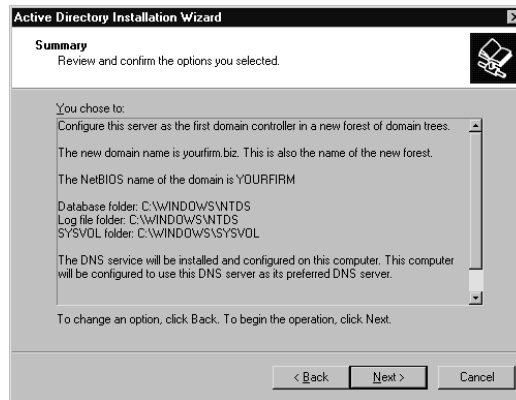


NOTE

It is recommended that you always follow best password practices, creating passwords with a minimum of eight characters that include both upper- and lowercase alphabetical characters, some numeric and some symbol characters, avoiding words found in dictionaries or phrases that can be easily guessed.

Figure 7.13 shows you the Summary screen, where you find all the settings you chose during the wizard, including, if you scroll down, that your Restore Administrator password is the same as your domain administrator password. If you are satisfied with these settings, choose **Next**. The next step is shown in Figure 7.14, and the only option shown is a Cancel button. This is an illusion, because it does not cancel the process as promised.

Figure 7.13 Summary or Last ADIW Dialog Before You Can Bail Out



The wizard will configure the Active Directory next. You can click the **Cancel** button (shown in Figure 7.14), but if you change your mind now you still sit through the writing process until it is completed. You still reboot the server, and run **dcpromo** again to “erase” the Active Directory information. Then, you must reboot the server a second time, and run **dcpromo** with the “intended” settings. Wait. Reboot—and finally, it’s finished as you intend it. Each iteration of **dcpromo** requires 20 to 30 minutes, so be careful you don’t make mistakes that require repeating this process. When the writing process completes successfully, you see the final window of completion, Figure 7.15.

Figure 7.14 Active Directory Installation Wizard—Writing**Figure 7.15** Completing the Active Directory Installation Wizard

After clicking **Finish** you will have to reboot, but take note of the information in this window. “This domain controller is assigned to the site **Default-First-Site-Name**.” Your first DC is automatically placed in this default site, so you will know where to find it the next time you use the Active Directory Sites and Services tool.

Why Did Your DNS Registration Fail?

At the DNS Diagnostic dialog, the ADIW has attempted to locate a DNS and create the appropriate SRV records for your domain—and failed. Three situations can cause a DNS failure:

1. There is no DNS server. Let the ADIW install DNS for you.
2. No response came back from **dcpromo**'s request, because the IP settings on the server or somewhere in between were misconfigured. Find that error and fix it.
3. A response from a DNS server was received, but the DNS server would not accept dynamic updates. Remember that your DNS server must support both SRV records and dynamic updates. Also remember that Microsoft is not the only vendor that makes a compatible DNS server; if you already have a non-Microsoft DNS server on your network, determine whether it meets the qualifications before automatically installing a Microsoft version.

Non-Microsoft DNS servers that are RFC 2782 (SRV records) and RFC 2136 (dynamic updates) compliant include UNIX or UNIX-derivative machines running BIND or Lucent's QIP DNS/DHCP.

Compatibility with Previous Operating Systems

According to the Windows Server 2003 Help File, the default security setting is to require Server Message Block (SMB) signing and encryption or signing of secure channel traffic. Some downlevel Windows-based operating systems do not have built-in support for SMB signing or secure channel encryption and signing. This means that only computers running Windows 2000 or later can communicate with your Windows Server 2003 DC.

You have the choice of upgrading the operating system or installing an additional helper service to allow communication with your new 2003 DC.

- **Windows for Workgroups (Windows 3.1x)** Your only choice is to upgrade.
- **Windows 95** Upgrade or install Active Directory Client.
- **Windows NT 4.0** Upgrade or install service pack 4 or greater.

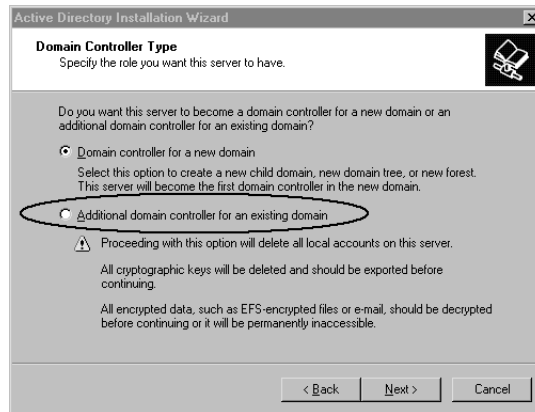
One final option is to disable the Windows Server 2003 default security. These settings are in the Default Domain Controllers Policy. See *Compatibility in Windows Server 2003 Help* for more specifics.

Creating Additional Domain Controllers

To add more DCs to your new domain, yourfirm.biz, you must install Windows Server 2003 on another machine. Remember the initial server installation is either a standalone or member server and then it can be promoted to a DC. Dcpromo, otherwise known as the ADIW, accomplishes this feat. You just created your first DC, so the steps are still fresh in your mind, right?

Since the domain and DNS servers already exist, when you see the window shown in Figure 7.16, select the second option, **Additional domain controller for an existing domain**. Next, enter the credentials for your parent domain administrator, and the ADIW creates the new DC with replication, dynamic updates, and DNS SRV records all in place.

Figure 7.16 Domain Controller Type



EXERCISE 7.01

PROMOTE A SERVER TO AN ADDITIONAL DOMAIN CONTROLLER

This exercise walks you through promoting a member server to be an additional DC in your existing domain. There are a few prerequisites you will need to meet before you begin:

- One DC
- One standalone or member server
- Both of these servers connected on the same network
- Both of these servers set up with TCP/IP pointing to the same DNS server

Once these have been met, follow these steps:

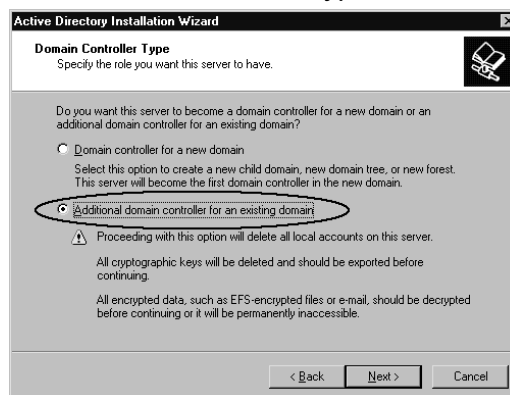
1. On your member or standalone server, make sure you are logged on with administrator permissions and that the prerequisites are met.
2. Begin the promotion process by clicking **Start | Run** and typing **dcpromo**. Click **OK**.
3. The ADIW is launched (see Figure 7.17). Click **Next**.

Figure 7.17 Welcome to ADIW



4. Click **Next** on the next dialog labeled **Operating System Compatibility**.
5. Select **Additional domain controller for an existing domain**, and click **Next** (Figure 7.18).

Figure 7.18 Domain Controller Type



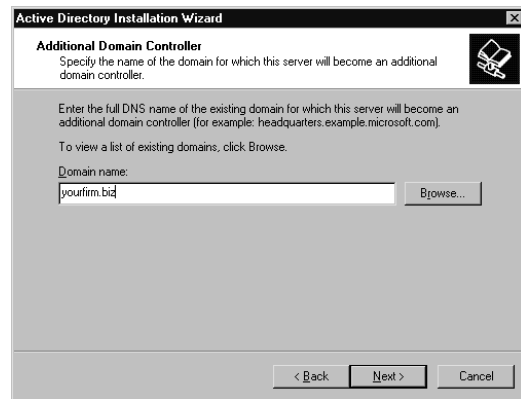
6. Type in the Administrator account and password. Type in the domain name if it is not already there (see Figure 7.19).

Figure 7.19 Network Credentials



7. This dialog requires the FQDN that matches the A record in your DNS server. By default, it puts whatever was in the last dialog. Make sure it is correct and click **Next** (see Figure 7.20).

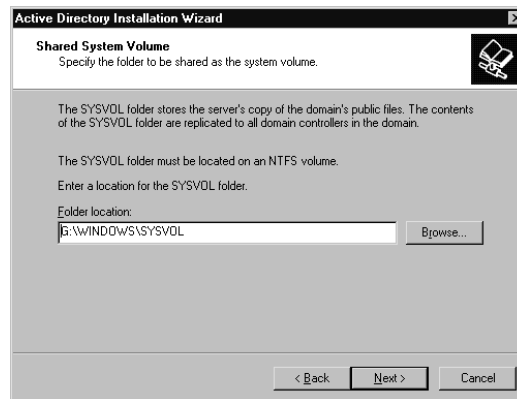
Figure 7.20 Full DNS Name of Existing Domain



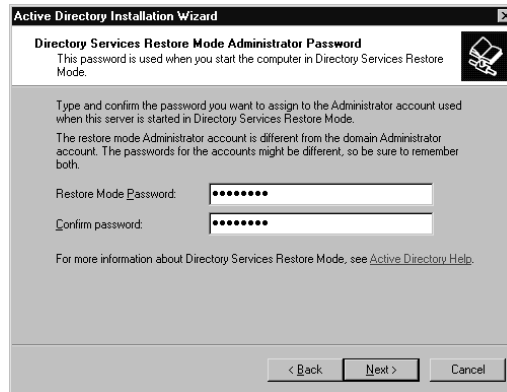
8. The next two dialogs should seem familiar. You must specify the location of the Active Directory database and log. Keep the defaults of \WINDOWS\NTDS. Click **Next** (see Figure 7.21).

Figure 7.21 Location of Active Directory Database and Log Files


9. Specify the location of your system volume (SYSVOL) folder, which *must* be on an NTFS partition. The default is fine as long as you formatted your disk with NTFS. Click **Next** (Figure 7.22).

Figure 7.22 Shared System Volume Location

10. Type in the same password you used for your domain password. Figure 7.23 requests a password for the directory services restore administrator which is different from the domain administrator; however, unless there are compelling reasons otherwise, put in the same password. It can be difficult enough dealing with the crash requiring a restore, without adding to the stress of remembering a password not often used. Click **Next**.

Figure 7.23 Directory Services Restore Mode Administrator Password

11. The next step is easy enough. Review the Summary (see Figure 7.24), and click **Next**. Wait! Think about what `dcpromo` is about to do and what must be in place for it to work. After clicking Next, you will test the settings outlined in the prerequisites. The ADIW must find a DNS server, ask for the location of a DC in the existing domain, locate and authenticate to that DC, and transfer a copy of the Active Directory database, making a new DC. If this fails, you have a great opportunity to determine why. Use the hints just mentioned and verify that all is in place. We requested a Retry more than once before noticing the we misspelled *administrator*—sometimes, it's the simplest answer.

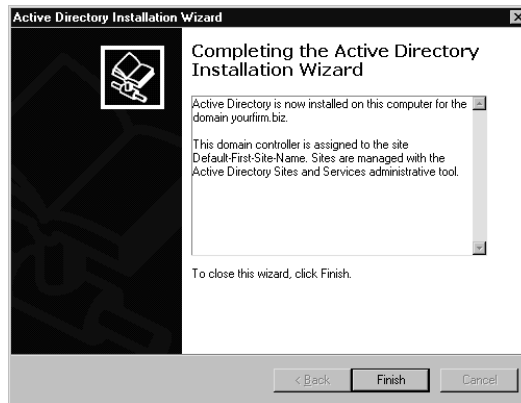
Figure 7.24 ADIW Summary

12. After the transfer is complete, which should take several minutes (see Figure 7.25), you will see the dialog in Figure 7.26—a sign that all went well. Again, note the location of the DC in the Default-First-Site-Name site, and click **Finish**.

Figure 7.25 ADIW Transferring a Copy of the Active Directory Database



Figure 7.26 Completing the ADIW



13. The last step is to restart the machine. As shown in Figure 7.27, click **Restart Now**. This obviously reboots the server, which then comes back online as a DC in your existing domain.

Figure 7.27 Restart the Server to Complete the Promotion



Use this new DC to practice the concepts in this book and better prepare yourself for the exam. In particular, the next exercise in this chapter requires two DCs.

With the advent of Windows Server 2003, you can now create DCs another way. Using a backup of the Active Directory (see the section *Backing Up Domain Controllers*) you can just do a restore to the server you want to promote. Why is this so great? Think about the sites that you might have in your domain. Remember that these are geographical divisions of the domain, and each site should have at least one DC for users to authenticate to at their local site. How did that DC come into existence? You have to promote a server using **dcpromo**, and that required direct and immediate access to another DC in order to extract and replicate all that the new DC needs. Again, this is at another site, so your options are to run the ADIW across the WAN, which could be excruciatingly slow and possibly late at night; or create the DC at the original site that already has a DC and then ship the whole computer to the remote site. The ability to create a DC from a backup gives you the option of backing up to a CD, sending the CD to the remote site, and performing a restore on a server, forcing it to become a DC. Now that *is* great news!

A few notes worth mentioning about this great feature:

- You will still need to have access to another DC to get the latest updates.
- Always use the most recent backup. The older the backup, the longer it will take the new DC to get the latest changes from the online DC.
- Never use a backup that is older than the tombstone lifetime of the domain, which by default is 60 days.
- If the DC that was backed up contained an application directory partition, it will not be restored to the new DC.
- This only works on a backup taken from a DC running Windows Server 2003.

Upgrading Domain Controllers

If you administer an existing domain and are looking to upgrade to Windows Server 2003, your best bet is to upgrade the DCs one by one until they are all at the same level. Server 2003 can co-exist with Server 2000 and NT, as long as you are aware of the caveats associated with such an environment. Even if you plan to upgrade all of the DCs, you must still temporarily run the network as a mixed environment.



NOTE

In Windows 2000, you had two choices of modes: *mixed* and *native*. Mixed mode exists as long as any of your DCs are running a downlevel operating system. Windows Server 2003 provides four “modes” to consider, but they are no longer called modes. Now, Microsoft refers to them as *functional levels*. These levels are discussed in detail in Chapter 4, “Working with Forests and Domains,” but be aware that as you migrate your existing DCs, you can begin to increase the functional level of your domain and consequentially your forest.

Upgrading your existing domain can be done in one of two ways: *in-place upgrade* or *migration*. An in-place upgrade means that you take your existing DC and install Windows Server 2003 right over the top of it. Your existing domain structure with all of its user, group, and computer accounts will be migrated into the new Windows Server 2003 Active Directory. The advantages are clear:

- It's simple and quick.
- You don't need a new computer.
- No new SIDs or trusts have to be created, which keeps all your existing member servers and resource domains happy.
- Everyone gets to keep his or her password.
- Migrating from Windows 2000 to Windows Server 2003 in this manner works well.

Why would you upgrade any other way? Experts suggest that you avoid the in-place upgrade in the following situations:

- If you are trying to get an NT 4 PDC to become an Active Directory DC in an existing Active Directory domain, you can't do it this way. Upgrading an NT 4 PDC will always create a new Active Directory domain.
- Upgrading an NT 4 DC allows you to create a new Active Directory domain name, but you are forced to keep the NetBIOS name.
- You cannot merge your NT 4 domain into your Active Directory domain.
- All accounts are upgraded and there is no way to roll this back. We suggest that you take one of your BDCs in NT or one of your Active Directory DCs in 2000 and move it offline in case there are any problems. You can then bring the DC back online, and your original domain will still exist.
- All the extra "stuff," such as unused groups and users, in the NT SAM is there in your new Active Directory domain.

A migration is accomplished by creating a new pristine Active Directory on a new server. Then, you use a migration tool to copy the domain information from your old domain to your new one. Here are some of the advantages of this method:

- Migration is gradual. You can migrate one department at a time.
- Accounts are copied rather than moved, so you can return to the old domain if necessary.
- You avoid the complexity of taking existing database bugs and moving them into your new Active Directory.
- You can re-evaluate your existing domain structure and consolidate or expand your domains, as you deem necessary.

There are also disadvantages to migration:

- You need new computers to install your new domain.
- Generally, users have to create new passwords.
- A migration tool might have to be purchased. Microsoft has a free migration tool, Active Directory Migration Tool (ADMT), but it is designed for the small to medium domains. You can purchase other tools for enterprise-level migrations. At around \$10 per user, this can get expensive.
- You cannot use the same NetBIOS name that exists in your old domain.
- Migration is more work. You might have to go to every member server and re-do all of the groups (some migration tools provide ways to avoid this by using Security ID histories (SIDHistory)).

Whether to upgrade or to migrate is an important decision. To master the exam objectives, you must understand the differences and know the pros and cons involved. Either choice will produce issues to consider and plan for.

Placing Domain Controllers within Sites

Sites were discussed earlier in this book, in Chapter 6, “Working with Active Directory Sites.” Remember that you don’t need more than the default site unless you have a network with subnets that are connected by slower WAN links. If you have multiple sites, you need to put your DCs in the right places; otherwise, you will “span the WAN”; in other words, your DCs will replicate continuously over your WAN, eating up the bandwidth needed by your users.

Here is a brief review of what you need to manage your sites. The tool of choice is the Active Directory Sites and Services console.

- Create a site name
- Create subnets to match your actual IP subnets
- Move the servers listed in ADSS to their respective assigned sites

First, you define the site itself. We suggest at least renaming the Default-First-Site-Name to something befitting your company location, such as CorpHQ. Next, you need to define subnets. Your physical and logical IP subnets should already exist, but you need to define them in the site tool, by specifying the subnet address and assigning it to your site. To place your DC in a site, open the **Servers** folder and move the DC into the appropriate site. Remember that all servers within a site will automatically determine the replication process, but you must configure the replication between sites.



NOTE

If a site does not replicate at least once every 60 days, the default tombstone lifetime, your DC will begin to delete inactive objects. Thus, if a site was unable to communicate with another site for a couple of months, it would start throwing out objects that were still being used by other DCs in other sites. For this reason, be careful to *never* reconnect a site that has not replicated in the last 60 days.

Backing Up Domain Controllers

Every Windows server has a *system state* that includes the Registry of that server (among other things). On a DC, the system state also includes Active Directory. Since replication of Active Directory occurs automatically, you only need to back up the system state of one of your DCs to back up Active Directory. However, your other DCs might run other applications or have files that only exist on that machine, so be sure that those are included in your routine backup. If you have multiple sites, consider backing up the system state of one DC per site to facilitate easier access to the Active Directory backup data should you need to restore it.

To back up the system state of any computer, you must be connected locally. In other words, the computer that you are logged on to, and are running the backup application from, is the only system state you can back up. If you are using a tape drive and you want to back up the Active Directory, you will have to connect the tape drive to a DC directly. The local computer rule applies to a restore as well: you must be directly connected to the computer on which you want to restore the system state. Backup media options have been increased from the limited Windows 2000 Backup to include removable media (CD, DVD) or a shared resource.

What Is in the System State?

The system state, in general, contains the machine-specific data of the Registry, COM+ class registration database, files under Windows file protection, and system boot files. Additional data are included depending on the computer configuration. As you know, each server might perform a different role in your domain, and as the role changes, so will the data included in the system state. For example, a Web server running IIS includes its metadirectory in the system state, and a DC, which we are most concerned with here, includes Active Directory and the SYSVOL directory.

Backing up and restoring the system state is an “all or nothing” proposition. You *cannot* back up or restore individual components of the system state. However, you can restore a system state to an alternate location. This is new in Windows Server 2003, and as mentioned in the text, can be used to create additional DCs.

Restoring Domain Controllers

To restore a DC from backup, you must determine which part needs to be restored. The first question to ask is, does the Active Directory need to be restored authoritatively or non-authoritatively?

- **Non-authoritative restore** means that you just restore the Active Directory to whatever point it was at when you backed it up, and then let the new changes from the other DCs automatically replicate to this DC to bring it up to the most current state.
- **Authoritative restore** means that the Active Directory that you restore is the master, and even though the data on it is “old” compared to the other DCs, its data is to be taken as the *authority* or final word on the Active Directory.

Use the non-authoritative restore when you have lost the DC but the data on the other DCs is accurate; in other words, there is nothing the “downed” DC knows that no other DC knows. Authoritative restore is used when the “downed” server *does* know something the other DCs don’t. For example, suppose you delete the user account Hannah on Monday. On Friday, you learn that Hannah was not supposed to be deleted. You can’t just create a new user called Hannah because the new account takes on a new SID, and all the permissions, rights, and privileges that were associated with the first Hannah are lost. You must restore the original account, which by now is removed from all the DCs. Fortunately, you can perform an authoritative restore from Sunday night’s backup to get Hannah’s account back. This forces all the other DCs to re-accept Hannah’s original account.

That is a simplified version of what the backup and restore capabilities can do. When you restore authoritatively, you can restore the entire Active Directory or select different levels of the domain hierarchy, even down to the single object restore, as was needed in the previous scenario. As long as you know the exact FQDN for the object to be restored, you can recover it.

The steps to restore Active Directory start with a good recent backup. Remember that your restore is only as good as your backup. Spend the time, effort, and money to ensure that you have good valid backups. With the backup in hand, you are ready. On which DC should you run the restore? See Chapter 11, “Ensuring Active Directory Availability,” for more information about backing up and restoring the Active Directory. Also remember that you can now use a backup from a DC running Windows Server 2003 to create additional DCs.



EXAM WARNING

Watch for questions that ask you about how to restore Active Directory. Remember that you can only back up the Active Directory if your tape backup hardware is directly connected to the DC, and you can only restore the Active Directory back to a directly connected server. This is true for any system state data, not just for DCs.

EXAM
70-294
OBJECTIVE
1.2
2.5.2

Managing Operations Masters

Flexible Single Master Operations (FSMO, pronounced *fizz-moe*) are certain roles assigned to DCs that need only exist on one DC and not all DCs. They are also called *operations masters*. These operations are critical in managing such objects as the schema and determining uniqueness among a forest, tree, and domain. Earlier in the chapter we declared all Windows Server 2003 DCs equal—that was not entirely accurate. FSMOs make some DCs more important than others, at least in regard to certain domain tasks, and it is your job to know which DCs perform these roles and what to do if a role needs to be switched to another DC. You must also know how to *seize* a role should you lose one.



NOTE

Which is it? Technically, Microsoft coined the term *FSMO* during the Windows 2000 beta, and it was used in most of the books that came out at that time. Officially, Microsoft uses the term *operation masters* (OMs), so if you want to be politically correct, use OM, but FSMO (fizz-moe) sounds cooler to say.

Understanding the Operation Masters Roles

What are these OMs, and why do you have to understand them? There are five OMs, and they are all automatically installed on the first DC. If you never create another DC or domain, the roles are self-sufficient. For the most part, things run very smoothly without any manual configuration. However, what happens if that one DC crashes or is replaced without much thought? This is plausible, because every DC contains the same Active Directory information, and in theory, there would be no problem if you lost one DC. However, if you lose the DC that is running one of these roles, you might find that you can't add anything new to your domain, and suddenly the importance of the roles becomes paramount. They exist, as you will see, to facilitate new additions to your domain structures by preventing duplications and confusion. To help you understand the FSMO roles, we discuss them individually next. Note that some roles are forestwide, and others are domainwide.

Forestwide Roles

Forestwide FSMOs, as the name suggests, reside on only one DC within the entire forest and perform the role for all domains in the forest. Only one of each of these FSMOs exists in a forest. They are automatically installed on the very first DC you used to create your domain, as are all the other roles. The forestwide roles are:

- Schema Master
- Domain Naming Master

In the following sections, we will discuss each in more detail.

Schema Master

The Schema Master controls what is in the schema. The schema controls what is allowed in the Directory. The fact that you can create a user account and assign attributes like a last name and password is due to the schema. If the schema changed to not allow last names, then the user account attribute of last name would go away. This is not to be trifled with, and so not even Enterprise Admins can modify the schema. Only Schema Admins have the right to modify the schema, and the most common way for the schema to be modified is by an Active Directory-aware application such as Exchange 2000.

Currently, not many applications can modify the schema, but that will change in time. If you have an application that modifies the schema, or if you do want to modify the schema directly (which can be done using the Schema MMC snap-in as described in Chapter 8, “Working with Global Catalog Servers and Schema”), you will have to be logged on as a Schema Admin *and* have access to the DC with the Schema Master role. One other note: the schema is forestwide, so modifying the schema affects every DC in the entire *forest*, not just the domain where the application is running.

Domain Naming Master

The Domain Naming Master FSMO keeps track of all the domain names in the entire forest, ensuring that no one creates duplicate domain names, and resolving issues in the unlikely event of “simultaneous” domain creation. This FSMO must be available when you create a new domain; otherwise, the domain creation will fail. This, too, is a forestwide FSMO and there is only one within the forest. There might be some delay in your domain creation if, for example, you are in California and your Domain Naming Master is in Australia. Fortunately, domain creation is a rare occurrence, so the problem rarely comes up.

Domainwide Roles

The remaining three FSMO roles are domainwide, meaning there is one of each in every domain. The first DC you create in a domain is relegated the following three roles:

- Relative ID (RID) Master
- PDC Emulator
- Infrastructure Master

We discuss each of these in more detail in the following sections.

RID Master

As you know, every object in Active Directory is assigned a SID, but did you know that all but the last 32 bits of a SID in each domain is the same? A SID looks like this:

S-1-5-21-X1-X2-X3-RID

S-1-5-21 is the same in every SID on every domain ever created. X1, X2, and X3 are randomly generated 32-bit values created at the time the domain is created. These are the

same values on every SID in that domain. Only the last 32-bit number represented by RID is unique to each SID. To ensure uniqueness, there must be a master DC that keeps track of all RIDs; hence, the RID Master was conceived. A pool of RIDs exists on the RID Master to be doled out to DCs as needed whenever a new object is created. This is because it would be inefficient to have to generate a new RID with every new object creation. Instead, the RID Master generates 500 RIDs and then, as the pool drops by 100 or so, it generates 100 to 250 more to keep the pool relatively full.

Each DC is allocated a group of RIDs, and since the RID Master keeps track of these allocations, you won't have the problem of two DCs handing out the same SID, which would be a real possibility if there were no RID Master.

PDC Emulator

If you understand NT 4 domains and their reliance on a PDC, it is easy to see why, if your domain contains NT BDCs, you must have a PDC Emulator. NT 4 has no clue what Active Directory is, nor does it know that it could ask any DC for the latest updates. It is programmed to look for a PDC and accept changes *only* from that source.

Other downlevel operating systems are programmed to look for a PDC as well:

- Windows 95, when used to log on to a domain, must go to the PDC to look for its policy settings with every logon. (Windows 9x computers cannot be members of domains, but they can be used by users with domain user accounts to log on.)
- Trust relationships in NT 4 require that the PDC be consulted before creating and establishing a trust between domains, so the PDC Emulator handles this.
- Windows populates the Network Neighborhood and My Network Places by consulting the *master browser*. You guessed it; the PDC Emulator fills that role.

It is clear why this FSMO is important. The PDC Emulator makes it possible for your downlevel systems to cohabitate with the Windows 2000/2003 Active Directory world. Does this mean that after you rid your domains of these legacy nuisances, you can be rid of the PDC Emulator? Actually, the answer is no. It turns out that this FSMO plays an important role in two other Active Directory functions essential to your 2000+ environments.

A common administration request is to change a user's password. For example, it's Friday, your system tells you to change your password, and you do. Monday morning, you absent-mindedly put in the old password, which fails. "Oh yeah," you say, "I changed that on Friday. Hmm, I think I changed it to this..." You type it in and it fails again. "I'm sure that was it!" You check your Caps Lock key and try again. It fails! Your account has now been locked out. Now you have to call the help desk, which is in Chicago (you work in Dallas), and request a new password. The help desk operator, says, "Sure thing, that's about all I do on Mondays."

However, here's a question to consider: Which DC is the help technician logged on to? It's not very likely that it is your DC in Dallas. Therefore, do you have to wait a few hours or more for the change to replicate to the Dallas DC before you can log on again? Fortunately,

the PDC FSMO is there to rescue you. With password changes and account unlocks, a higher authority is consulted if a DC doesn't find the information it needs in its local database. You log on again and your Dallas DC does not recognize the new password and sees that your account has been locked out. However, rather than reject you automatically, the DC looks for the PDC Emulator and passes the logon request to it. The PDC Emulator is in charge of these two settings and is the final authority. Whenever a password change or unlock request is made the change is reported to the PDC Emulator immediately.

That's one Active Directory function that requires the PDC Emulator in a Windows Server 2003-only environment, but what's the other one? Domains need to track changes constantly, and time is of the essence. If there is a change to Becca's account at 2:30 P.M. and another change at 2:35 P.M. that conflicts with the original change, the domain must recognize in what order to perform these changes, based on the timestamp. In a multidomain forest, it is difficult to keep everyone's time synchronized. Microsoft deals with this by setting up a forestwide time synchronization system, using the PDC Emulator FSMO. Each PDC Emulator in each domain acts as the timekeeper for its domain. All DCs check in with their respective PDC Emulators to synchronize their time. So, how do the PDC Emulators know what time to use? They go to the master timekeeper: the first PDC Emulator installed in your forest. Remember that the first DC of the forest that you create contains all five FSMOs, so it is the first PDC Emulator you created with which all other PDC Emulators synchronize their time.

Authoritative Time Server

As mentioned, Microsoft provides a time-keeping process to try to establish a common time among all the computers in the domain. It uses a hierarchical scheme as follows:

- Each client workstation nominates its authenticating DC as its timekeeper.
- Each member server does the same.
- Each DC in a domain nominates the PDC Emulator as its timekeeper.
- Each PDC Emulator follows the same hierarchy as the domains in the forest, selecting the PDC Emulator above as its timekeeper until the "root" PDC Emulator is reached, which is nominated as the master timekeeper for the forest.

The time service that Microsoft uses and is required by Kerberos authentication is *W32Time*. This service is automatically in place and working between the DCs in your forest. It is a good idea to find your master timekeeper PDC Emulator and set up a time server for it to synchronize

You can use the Simple Network Time Protocol (SNTP) to establish a time standard for your domain. Use the following command to access an SNTP server:

Continued

NET TIME /SETSNTP:SERVER_LIST. A list of servers can be found on the Internet. Here are two provided by the United States Naval Observatory:

- Ntp2.usno.navy.mil (192.5.41.209)
- Tock.usno.navy.mil (192.5.41.41)

Other time servers are managed by the National Institute of Standards and Technology (NIST) found at www.nist.gov. For in-depth instruction and reference on this topic, refer to Microsoft's white paper, *wintimeserv.doc* on their Web site.

Infrastructure Master

According to Microsoft, without the Infrastructure Master, changes between your DCs would be slow. The Infrastructure Master speeds this process up. There is one infrastructure FSMO per domain, and it is on the first DC you installed for that domain, unless you have transferred or seized the role (discussed in the next section).

The Infrastructure FSMO appears to be somewhat enigmatic in how it does its role. We have seen references to this FSMO indicating as just stated that the speed of services managing the domain is increased—but how? First, it updates the group-to-user references whenever a change is made; and second, this FSMO is in charge of seeking and destroying those *stale* objects floating around your Ethernet. Actually, this refers to the references that are no longer valid. This can occur when an object is moved, renamed, or deleted. The infrastructure FSMO uses the GC to check for these stale references and then removes them. Because the GC and the Infrastructure FSMO have to work so closely together, Microsoft recommends that these two roles run on separate DCs. Of course, by default, they are on the same DC, so it is up to you to move one of these roles to your second DC as soon as you have one.

EXAM
70-294
OBJECTIVE
1.2.1

Transferring and Seizing Operations Master Roles

With your newfound understanding of FSMOs, you can see that they are essential for domain consistency and integrity. It has been said more than once that these roles are created automatically, but the defaults assigned by that automatic creation might not suit your environment, and you might consequently need to either transfer these roles to a better machine or move them before retiring a server. It is also possible for you to lose a DC containing one or more of these roles and be unable to recover it. This section describes how to transfer and seize these operations master roles.

Transferring FSMOs

You've decided to transfer your FSMOs from the original location, on the first DC, to another server that will be your super-server. When the transfer is planned, you can manually move these roles by following the steps outlined in the next sections.

Transferring the Schema FSMO

First, you must be a member of the Schema Admins group. Next, you need to access the Active Directory Schema snap-in, which is not in the Administrative Tools menu but must be added to an MMC.

To install the Active Directory Schema snap-in, follow these steps:

1. Open a command prompt [**Start** | **Run** | **cmd**, and click **OK**.
2. At the command prompt, type **regsvr32 schmmgmt.dll**. This command will register schmmgmt.dll on your computer. Successful registration produces the dialog box shown in Figure 7.28.

Figure 7.28 Register Service



3. Click **Start** | **Run...** and type **mmc /a**. Then, click **OK**. This opens a blank MMC in author's mode.
4. On the **File** menu, click **Add/Remove Snap-in**, and then click **Add**.
5. Under **Snap-in**, double-click **Active Directory Schema**, click **Close**, and then click **OK**.

Looking at the schema attributes, you can identify a few. Figure 7.29 shows the *cn* or *Common-Name* attribute, which is mandatory in a user account. Right-clicking on the object named **Active Directory Schema** affords you several options (see Figure 7.30). From this tool, you can see which DC is currently assigned the Schema Master by selecting **Operations Master...**, and you can transfer the FSMO to another DC by selecting **Change Domain Controller**.

Figure 7.29 Active Directory Schema Tool

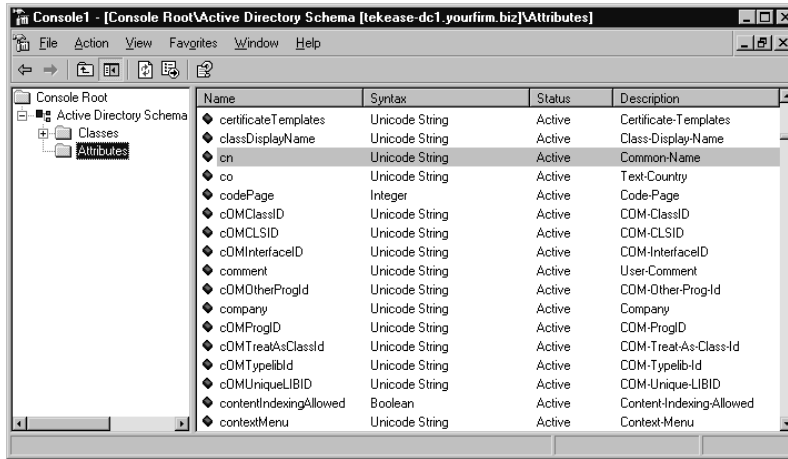


Figure 7.30 Management Options




Figure 7.31 depicts the next dialog in our quest. You are then given the choice to transfer the FSMO to **Any DC** or **Specify a Name**. Specify the new location and click **OK**. The new location is the FQDN of the DC to which you are transferring the FSMO. The system will refresh the screen and you will see that the focus has changed to the other DC you just specified (see Figure 7.32). To complete the task, you still need to right-click the **Active Directory Schema** object again, and this time choose **Operations Master**, which brings up the dialog box shown in Figure 7.33. In our example, we are moving the schema FSMO to the DC, skyline.yourfirm.biz. Click **Change** and the system will ask you to verify that you *really* want to make this change. Click **OK**. After a short pause, the confirmation dialog in Figure 7.34 appears. Click **OK**. The Schema FSMO is now on the skyline.yourfirm.biz DC.

Figure 7.31 Change Domain Controller




Figure 7.32 Change in Focus Prior to FSMO Transfer

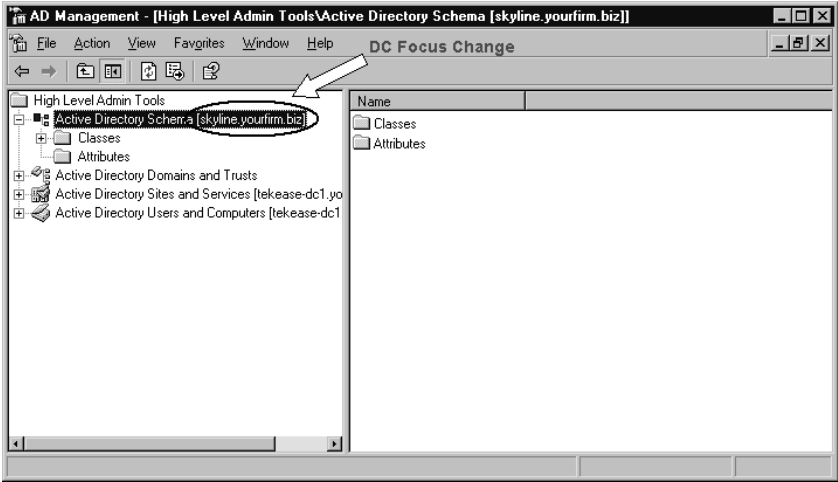


Figure 7.33 Change Schema Master

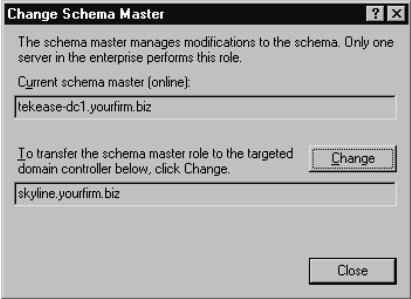



Figure 7.34 Confirmation of FSMO Transfer



Finding FSMO

If all you ever do is go with the defaults, you probably know where all the FSMOs are. However, there is a good chance of inheriting someone else's undocumented domain or walking into a foreign network as the perceived network guru. In these cases, you need to know how to find FSMOs. Microsoft has a tool to do just that: www.microsoft.com/windows2000/techinfo/reskit/tools/existing/dumpfsmos-o.asp. Okay, so the tool is not that impressive when you edit the *.cmd file, but the information is. Here are the steps to get a list of the FSMO roles and who has them:

1. Make sure you are logged on as either the local BuiltIn\Administrator for local access or Domain\Administrator or Enterprise\Administrator for remote access.
2. Open a command prompt: **Start | Run | "cmd" | OK**.
3. If you have access to the **dumpfsmos.cmd**, go ahead and run it and you are finished; however, you can do the same thing manually by reading on...
4. At the command prompt, type the following (bold text indicates what you should type, the rest depicts the DC's responses. Note that indents and bolding have been added for emphasis and easier reading):

```
C:\>ntdsutil roles connections

ntdsutil: roles
fsmo maintenance: connections
server connections: connect to server skyline
Binding to skyline ...
Connected to skyline using credentials of locally logged on user.
server connections: quit
fsmo maintenance: select operation target
select operation target: list roles for connected server
Server "skyline" knows about 5 roles
Schema - CN=NTDS Settings,CN=SKYLINE,CN=Servers,CN=Default-First-Site-
Name,CN=Sites,CN=Configuration,DC=yourfirm,DC=biz
Domain - CN=NTDS Settings,CN=TEKEASE-DC1,CN=Servers,CN=Default-First-
Site-Name,CN=Sites,CN=Configuration,DC=yourfirm,DC=biz
PDC - CN=NTDS Settings,CN=TEKEASE-DC1,CN=Servers,CN=Default-First-Site-
Name,CN=Sites,CN=Configuration,DC=yourfirm,DC=biz
RID - CN=NTDS Settings,CN=TEKEASE-DC1,CN=Servers,CN=Default-First-Site-
```

Continued

```

Name, CN=Sites, CN=Configuration, DC=yourfirm, DC=biz
Infrastructure - CN=NTDS Settings, CN=TEKEASE-DC1, CN=Servers, CN=Default-
First-Site-Name, CN=Sites, CN=Configuration, DC=yourfirm, DC=biz
select operation target: quit
fsmo maintenance: quit
ntdsutil: quit
Disconnecting from skyline...
C:\>

```

From the output of our request, **list roles for connected server**, we see that the Schema FSMO is on the Skyline DC, which is where we transferred it, and the other four FSMOs remain on the original DC, tekease-dc1. Ntdsutil is a great tool, so learn how to use it.

Another tool you can use uses VBScripting as a GUI approach to the same goal: finding FSMO. This tool is user friendly by generating a pop-up dialog for your input of a server, and then displaying five pop-up dialogs, each with the location of a FSMO. Try searching the Internet for “finding fsmo vbs,” or go to www.server-watch.com/tutorials/article.php/10825_1472341_5.

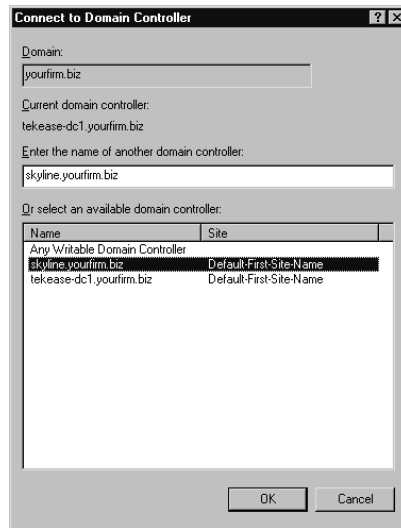
Transferring Domain Naming FSMO

Transferring this FSMO requires you to have Enterprise Admin level permissions and uses the Active Directory Domains and Trusts (ADDT) tool. Your first step requires you to change the focus of the tool to the DC to which you want to transfer the domain-naming FSMO. In the **Active Directory Domains and Trusts** tool, click **Action | Connect to Domain Controller...** That brings up the dialog shown in Figure 7.35. Fill in the name of another DC and click **OK**. You are returned to ADDT and nothing appears to have changed; however, your focus is now on the other DC. As with the Schema FSMO change, right-click the **Active Directory Domains and Trusts | Operations Master... | Change** and the transfer is complete.



NOTE

The Domain Naming Master can only reside on a DC that contains a GC. It appears that this FSMO requires access to the GC to function.

Figure 7.35 Connect to Domain Controller

Transferring RID, PDC, or Infrastructure FSMOs

To transfer the RID, PDC Emulator, or Infrastructure FSMOs, you use the Active Directory Users and Computers (ADUC) tool. You must be a Domain administrator to perform this function. First, change your focus to the DC that will receive the transfer by right-clicking the domain object. Select **Connect to Domain Controller... | Enter the name of another domain controller OR Select an available domain controller**, and click **OK**. Right-click the domain object again and select **Operations Masters...**. Notice in Figure 7.36 that there are three tabs: one each for the RID, PDC, and Infrastructure operations masters. These three FSMOs are domain specific, not forest specific, and they are all transferred using this same dialog box. As with the forest-specific FSMO transfers, click **Change...**, confirm that you want to transfer the FSMO, and the ADUC completes the function.

Figure 7.36 Operation Masters: RID, PDC, and Infrastructure



NOTE

Creating an Infrastructure Master FSMO on a DC that contains a GC is undesirable unless every DC in your domain is a GC. In a single DC domain, that's easy; all five FSMOs and the GC are on the sole DC. However, GCs are not automatically placed on each new DC, so you should move the Infrastructure FSMO over to a different DC when you begin creating additional DCs.

EXAM
70-294

OBJECTIVE
1.2.1

Responding to OM Failures

As long as you know where the FSMOs in your domain reside and ensure that they are transferred before decommissioning a DC, you can avoid most problems. A good rule of thumb to follow is to always demote a DC before taking it offline or replacing the computer on which a DC exists. By demoting a DC, you ensure that all Active Directory information is synchronized and any FSMO is automatically transferred. What happens if you lose a DC that had a FSMO on it?

If a FSMO is lost in your domain, there is no automatic response within the domain to elect a replacement; you just don't have a DC performing that role. Depending on which FSMO you lost, this can cause some interesting and sometimes fatal disasters in your domain. Forcing a FSMO into existence is called *seizing the master*. This process is not generally as user friendly as the transfer process, except when the role being seized is that of the PDC Emulator or the Infrastructure Master.

Seizing the PDC Emulator or Infrastructure FSMO

Seizing the PDC Emulator or Infrastructure FSMOs is still accomplished through the same GUI tool used previously: Active Directory Users and Computers. Since the DC with the lost FSMO is unavailable, the DC you are focused on should suffice. However, you can switch the focus by right-clicking on the domain object, selecting **Connect to Domain Controller...** | **Enter the name of another domain controller OR Select an available domain controller**, and clicking **OK** (see Figure 7.35). To seize or force a transfer of the PDC or Infrastructure, right-click the domain object and select **Operations Masters...** | **[PDC or Infrastructure]** |. Notice that the service has attempted to contact the FSMO in question, and the dialog displays a message that it is offline (see Figure 7.37). Click **Change...** anyway. Confirm your request. This time, a warning dialog box will appear asking you again if you are sure you want to transfer the operations master role. Click **OK**. A third dialog then appears with an explanation and question:

The current operations master cannot be contacted to perform the transfer. Under some circumstances, a forced transfer can be performed.

Do you want to attempt a forced transfer?

Figure 7.37 Failed to Connect to PDC FSMO

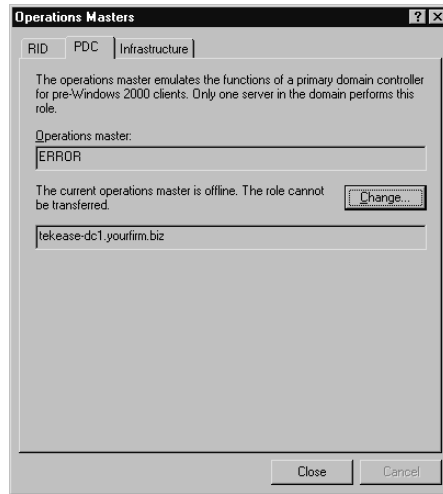



Figure 7.38 Forcing a FSMO Transfer



Click **Yes** to complete the seizure or forced FSMO role transfer. To summarize, the process requires *three* confirmations to perform the process, so be patient. Remember, this only applies to two of the domainwide FSMOs: PDC Emulator and Infrastructure. The RID FSMO *cannot* be seized from the GUI tool.

Seizing the RID Master, Domain Naming Master, and Schema Master FSMOs

Seizing the roles of RID, Domain Naming, and Schema Master requires the command-line utility **NTDSUTIL**. Follow these steps to perform this type of seizure:

1. Click **Start | Run** and type **cmd**. At the command prompt, type **ntdsutil** and press **Enter**.
2. Type **Roles | Enter**. The prompt will change to **fsmo maintenance:**.
3. Type **Connections | Enter**. The prompt changes to **server connections:**. As in the GUI ADUC, you have to change your DC focus to the DC that is receiving the transferred role.

4. Type **Connect to server** <servername> and press **Enter**, where <servername> is the name of the DC receiving the transferred role.
5. Type **Quit** and press **Enter**. This completes the focus change and returns you to the **fsmo maintenance:** prompt.
6. Type **Seize** <fsmo> **master** and press **Enter**, where <fsmo> is the operations master role you are trying to transfer: RID, Domain Naming, or Schema.
7. Type **Quit** and press **Enter** to exit the FSMO maintenance, and type **Quit** and press **Enter** a second time to exit NTDSUTIL.

Here is an example of the messages that appear when you seize the RID FSMO from the DC named dc3.yourfirm.biz and give it to dc1.yourfirm.biz:

```
C:\>ntdsutil
Ntdsutil: roles
Fsmo maintenance: connections
Server connections: connect to server dc1.yourfirm.biz
Binding to dc1.yourfirm.biz...
Connected to dc1.yourfirm.biz using credentials of locally logged on
user.
Server connections: quit
Fsmo maintenance: seize rid master
```



NOTE

A pop-up dialog will appear, requesting confirmation that you want to proceed.

```
Attempting safe transfer of RID FSMO before seizure.
Ldap_modify_sW error 0x34(52 (Unavailable)).
Ldap extended error message is 00002DAF: SvcErr: DSID-03210300, problem
5002 (UNAVAILABLE), data 1722

Win32 error returned is 0x20af(The requested FSMO operation failed. The
current FSMO holder could not be contacted.)
Depending on the error code this may indicate a connection, ldap, or
role transfer error.
Transfer of RID FSMO failed, proceeding with seizure...
Searching for highest rid pool in domain
Server "dc1.yourfirm.biz" knows about 5 roles
```

```

Schema - CN=NTDS Settings,CN=DC1,CN=Servers,CN=CorpHQ,CN=Sites,
CN=Configuration,DC=yourfirm,DC=biz
Domain - CN=NTDS Settings,CN=DC1,CN=Servers,CN=CorpHQ,CN=Sites,
CN=Configuration,DC=yourfirm,DC=biz
PDC - CN=NTDS Settings,CN=DC1,CN=Servers,CN=CorpHQ,CN=Sites,
CN=Configuration,DC=yourfirm,DC=biz
RID - CN=NTDS Settings,CN=DC1,CN=Servers,CN=CorpHQ,CN=Sites,
CN=Configuration,DC=yourfirm,DC=biz
Infrastructure - CN=NTDS Settings,CN=DC1,CN=Servers,CN=CorpHQ,CN=Sites,
CN=Configuration,DC=yourfirm,DC=biz
Fsmo maintenance: quit
Ntdsutil: quit

```

That's it—no reassurance that it was completed. Instead, it is simply stated that server “dc1.yourfirm.biz” knows about five roles and it lists the change requested. Remember, this works for the Schema and Domain –Naming seizures as well. Just replace the word **rid** in the text with **schema** or **domain naming**. By the way, unless you enjoy extreme chaos, do *not* bring an old master back online in your domain. Format the disk and build the machine from scratch.

EXERCISE 7.02


TRANSFERRING A FSMO

In this exercise, you will transfer a FSMO in your domain from one DC to another; essentially, you will “lose” the FSMO and seize it back. To do this exercise, you must have at least two computers networked together that are both DCs. The prerequisites include:

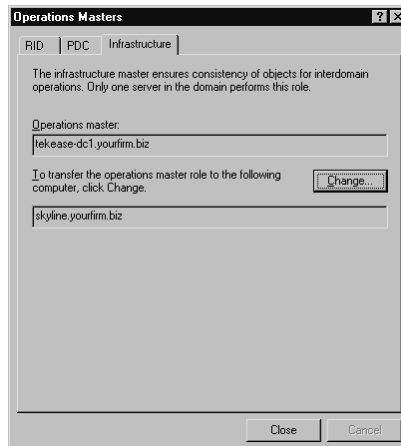
- Two DCs in the same domain and networked together.
- The first DC you install automatically assumes all five FSMO roles: Schema, Domain Naming, PDC Emulator, RID, and Infrastructure. If you are doing the examples in the book, you might have transferred the schema FSMO already. As a bonus, you can seize that as well.

First, you must transfer the Infrastructure FSMO:

1. Open Active Directory Users and Computers: **Start | Programs | Administrative Tools | Active Directory Users and Computers**.
2. Right-click your domain object and select **Connect to Domain Controller** (see Figure 7.39).
3. Select your other DC and click **OK**.

Figure 7.39 Connect to Domain Controller

4. Right-click on the domain object again.
5. This time, select **Operations Masters**.
6. Click on the **Infrastructure** tab (see Figure 7.40).

Figure 7.40 Infrastructure FSMO

7. Click **Change...**
8. Click **Yes** to confirm that you are sure. Note: You might receive a warning dialog if you are attempting to transfer the Infrastructure FSMO to a DC that is also a GC server. For our exercise, that's fine;

however, in a production environment you must avoid this because the Infrastructure FSMO cannot function on the same server as the GC.

9. Click **OK** in the Transfer Successful dialog.

Now you need to seize the Infrastructure FSMO.

10. Using the ADUC that should still be open, right-click the domain object and select **Connect to Domain Controller**.
11. Select the original DC to make it your focus.
12. To simulate the loss of this FSMO and DC, just unplug it from the network.
13. Right-click the domain object again and select **Operations Masters**.
14. Click on the **Infrastructure** tab. With the other DC offline, you get the hourglass for a minute or so and then you get the **ERROR** dialog shown in Figure 7.41. The message states that the FSMO cannot be transferred, which is correct—a transfer implies smooth transition and fully connectivity to all parties. We are simulating the loss of this FSMO, so an error is what we want.


Figure 7.41 ERROR Infrastructure FSMO Offline



15. To force the Infrastructure FSMO back to your original DC, click the **Change...** button.
16. If your original DC is also a GC, the warning dialog appears. In either case, click **Yes** to indicate that you are sure you want to transfer this FSMO.

17. After some mulling over by the DC and of course the failure to find the existing FSMO, you are presented with a nice message specifying that this FSMO can be forced, if we so choose (see Figure 7.42). We so choose, so click **Yes**.

Figure 7.42 FSMO Holder Could Not Be Found



18. A nice dialog appears indicating success. Click **OK** and then **Close** the dialog.

Now, what to do about that other DC that is only offline physically? You are one plug away from domain chaos. If you re-insert the network plug, you will have two Infrastructure FSMOs in the same domain. Depending on the FSMO in question, the chaos will vary. With the Infrastructure FSMO, symptoms might not be readily apparent because its role is more interdomain than intradomain. In our example, it is a single domain forest so the Infrastructure FSMO is not used. An interesting note is that depending on which DC you ask, the Infrastructure FSMO is in a different location. According to our tekease-dc1 DC, the Infrastructure FSMO is on tekease-dc1, which we just forced in the exercise. After bringing skyline DC back online and asking it, the answer was that the skyline DC holds the infrastructure FSMO. What happens in 15 minutes when these two DCs compare notes?

Summary of Exam Objectives

This objective of this chapter is to teach you how to manage domain controllers (DCs) and operation masters (also called Flexible Single Master Operations, or FSMOs). Planning your DCs requires an understanding of what Active Directory is. Active Directory is the NTDS database containing the domain hierarchical structure, users, groups, and permissions. DCs are special servers that have been promoted to that role after having Windows Server 2003 installed. The DC(s) on your network play an administrative role in maintaining the domain. Such functions as user logon authentication and Active Directory replication make up the bulk of the DCs' responsibilities.

DC planning requires knowledge of your geographical and site structure, the number of users, groups, and computers on your network (both current and planned), and the amount of network traffic. These items, combined with the physical limitations of the computer running as a DC, control how many DCs you'll need and where you'll place them.

DCs are created from one of two types of servers: standalone or member (the server type is selected when you install Windows Server 2003). To create a DC, run the Active Directory Installation Wizard (**dcpromo**). Beyond the minimum hardware requirements, a DC requires a DNS server that supports both SRV records and dynamic updates. Creating additional DCs requires the same ADIW, but you must have access to a parent DC and know the credentials for an administrative account on that DC.

Backing up your DC consists of backing up two parts: the data that is stored on it, and the system state. The system state includes the Registry, and on a DC includes the Active Directory database. To back up the system state to a tape device, the device must be directly (locally) connected. The system state can only be restored back to the same DC from which it was backed up.

FSMOs are also referred to as operation masters. FSMOs handle required system operations that control the domain naming and internal infrastructure of your domains. By default, these FSMOs are automatically installed and managed by the domain, but if you need to move them or if you lose them, there are ways to manually manage them.

The Schema Master is managed through the Active Directory Schema tool. The Active Directory Domains and Trusts tool is used to transfer the Domain Naming FSMO. For the remaining domainwide operation masters, PDC Emulator, RID, and Infrastructure, use the Active Directory Users and Computers tool. As long as the existing FSMO is available, these FSMOs move to the designated new DC smoothly. Should a FSMO be lost, your domain will not function fully, and depending on which operation master is lost, your domain might not function at all. To force the creation of a new FSMO, use the *ntdsutil* from the command prompt.

Exam Objectives Fast Track

Planning and Deploying Domain Controllers

- ☑ Plan to have at least one DC per site. Sites keep the domain from flooding the WAN, so a DC at each site keeps the replication traffic off the WAN as well.
- ☑ The number of DCs to deploy is based on the number of users, geographic location, and fault tolerance. Each DC can handle 5000 objects, but you should have a second DC on the network in case the first DC crashes. A second DC also improves the network load by balancing the number of objects, changes, and authentication requests between the two DCs.
- ☑ If a DC is overloaded with account activity, you can add another DC or improve the hardware of that DC. Adding RAM is usually the first choice.
- ☑ DCs are created by promoting a server. Run `dcpromo` (Active Directory Installation Wizard) to do this.
- ☑ The first DC you create is the domain, tree, and forest all in one. Additional DCs are created using the same ADI^W tool.

Backing Up Domain Controllers

- ☑ DCs contain the Active Directory database and need to be backed up daily.
- ☑ The Active Directory is stored in the system state on a DC, so to back up the Active Directory, you must back up the system state.
- ☑ Backup of a DC's system state requires a direct connection to the DC. This means that the backup program must run on the DC, and if you are backing up to a tape drive, the drive must be locally connected.
- ☑ A backup of a Windows Server 2003 DC can be used to create additional DCs in the same domain. This is useful for creating DCs at remote sites.

Managing Operations Masters

- ☑ Operation masters are also known as Flexible Single Master Operations (FSMO).
- ☑ FSMOs fulfill services that must be assigned to one DC and cannot be shared like all other Active Directory objects.
- ☑ There are two forestwide operation masters:

- Schema Master controls all modifications to the schema. This FSMO must be accessible whenever the schema is extended. The schema is usually extended by applications that are Directory enabled. Microsoft Exchange is a good example of a Directory-enabled application.
- Domain Naming Master directs the naming of any domains within the forest. It ensures that names are unique. This FSMO must be accessible when you create a new domain anywhere in the forest.
- ☑ There is only one Schema and one Domain Naming FSMO for the entire forest.
- ☑ There are three domainwide operations masters:
 - PDC Emulator has three responsibilities. It acts as the PDC for all Windows NT 4 BDCs. It is also the central depository for all password and account lockouts. This means that any change to a password or account lockout is immediately delivered to the domain's PDC Emulator. If a DC determines that an account is locked out or the password is incorrect, the DC will consult its PDC Emulator to see if there is a change that it just has not yet received. Finally, the PDC Emulator is used by Microsoft to be the timekeeper for each domain. This ensures that all DCs use a common time when sharing information, and is used to determine the order in which changes that might conflict occur. The first PDC Emulator created in your forest is the master timekeeper for all other PDC Emulators.
 - RID Master is a pool of Relative IDs used to create unique Security IDs (SIDs). Every object in a domain requires a unique SID, and when the object is created, the DC must grab a SID from a pool of SIDs assigned by the RID. Consequently, if the RID FSMO is unavailable in your domain, you cannot create any new objects after the DC's pool dries up, because when it asks for more, it can't find the RID FSMO.
 - Infrastructure Master exists to maintain a "clean" database. It looks for invalid object references in the Global Catalog (GC) and removes them. This is primarily used in multidomain forests.
- ☑ There is only one domainwide FSMO for *each* domain.
- ☑ Due to the importance of each FSMO, it behooves you to make sure they exist and are centrally located for efficient access by the other DCs in your forest. Transfers of each FSMO are accomplished through different tools. Manual, *planned* transfers are done using the appropriate GUI tool:
 - Use the Active Directory Schema Tool to transfer the schema FSMO. This tool must be installed by registering the schmmgmt.dll.
 - Transfer the Domain Naming FSMO with the Active Directory Domains and Trusts GUI tool.

- The Active Directory Users and Computers tool is used to transfer the three domainwide FSMOs: PDC, RID, and Infrastructure.
- ☑ If a FSMO is lost when a DC goes down, you have to seize the master. This means forcing the creation of a new FSMO. Take care to not restore the FSMO that resided on another DC after seizing it.
- ☑ Seize an operation master using the **ntdsutil** from the command prompt.

Exam Objectives

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

Q: What happens if there is no DC at a site?

A: A DC outside of the site will handle all authentication requests, which is usually across a WAN connection. By virtue of the nature of sites, the WAN link will most likely be slow (since you would not create a separate site for subnets that are connected by fast WAN links). This can slow every logon and connection request made by users.

Q: After a site is created and DCs are assigned to the site, how is domain replication managed?

A: By default, DCs between sites don't replicate; therefore, it is imperative that you define the site connectivity using Active Directory Sites and Services. Right-click the DC to manage and define the type of connectivity to use. Intersite replication (replication between different sites, as opposed to intrasite replication between DCs within the same site) must be configured manually.

Q: When can I switch domain functional levels?

A: With the advent of the upgraded domain to Windows Server 2003, there are new functional levels added. To move from one functional level to another requires that your DCs be upgraded—not your workstations, but your DCs (you can still have NT 4.0 workstations in a domain that is operating at Windows Server 2003 functional level). As you move your DCs from Windows NT 4 to Windows 2000 and finally Windows Server 2003, you can increase your domain functionality by changing the domain functional level. There are four functional levels: Windows 2000 mixed, Windows 2000 native, Windows 2003 interim, and Windows 2003.

Q: What happens if I lose a DC and choose not to replace it?

A: Depending on your domain infrastructure and how often that DC gets updates, you might have no worries at all. If that DC contained information that had not been replicated to any other DC, then you must re-do that change. A solid domain infrastructure usually has enough redundancy and sound replication strategy that losing that DC should be of little consequence. However, if the lost DC performed any of the OM roles, you will need to seize the role for another DC.

Q: How many DCs should I back up?

A: Since every DC contains the same data within 15 minutes to a day, backing up every DC in your domain might be overkill. Single-site domains can generally get away with backing up just one DC. Multisite domains should back up one DC for each site. However, if the DC is also functioning in another role (as a file server, for example), there might be data on it that needs to be backed up, in addition to the system state/Active Directory partition.

Q: What happens if I lose the Schema Master?

A: Any task that requires the schema to be extended will fail. You will not be able to manually extend the schema using the Schema snap-in, nor will applications be able to extend the schema, which might cause their installation to fail, or cause the application to function incorrectly. To add software that extends the schema, permission must be granted. That is the job of the Schema Master.

Q: What happens if I lose the Domain Naming operations master?

A: You will not be able to create any new domains. As long as you do not need to create new domains, the network will function with no problem.

Q: What happens if I lose the PDC Emulator?

A: Any Windows NT 4 BDCs on the network will not function, because there is no PDC from which it retrieves DC replications. In addition, time synchronization and the up-to-date passwords and lockout account information is done on the PDC, so you might experience problems related to these issues.

Q: What is a symptom of losing my RID Master?

A: You can't create any new objects in your domain. The RID Master doles out SIDs, which are required to make any new domain object. You might not notice the absence of the RID Master immediately, because it will have already handed out SIDs to other DCs, and they can use these when new objects are created. However, eventually, the other DCs will run out of SIDs, and will not be able to create new objects.


Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Planning and Deploying Domain Controllers

- As a domain administrator you have seen the success of other departments using a RAS server to allow remote access to their domains. The other administrators use Windows NT 4 RAS and it has worked well for them. You want the same, so you install a Windows NT 4 RAS server in to your Windows Server 2003 domain. As you test this configuration, you continually get “Access Denied,” no matter which user you use to dial in with. What is a likely explanation for the continual failure to allow access?
 - Your domain was created using **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**.
 - Your domain mode is set to Windows Server 2003 domain functional level.
 - Your domain mode is set to Windows 2000 native domain functional level.
 - A Windows NT 4 RAS server cannot authenticate to a Windows Server 2003 domain.
- Using the services depicted in Figure 7.43, select the components that are required to create a domain and place them in order of implementation.

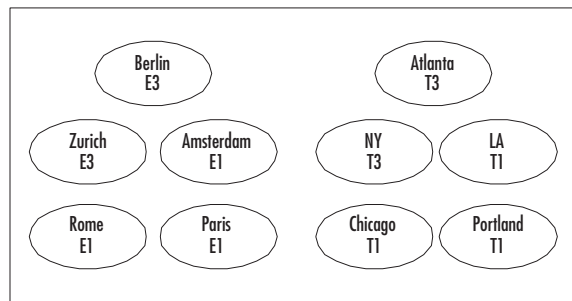
Figure 7.43 Network Services



- A. Site
 - B. DNS
 - C. WINS
 - D. DHCP
 - E. DC
 - F. BDC
 - G. Standalone server
 - H. Workstation
 - I. RIS
3. DasSchmeckt, the leading food services company outside the United States, has just merged with Yummy, Inc. in the United States. DasSchmeckt's headquarters are in Berlin, and Yummy, Inc. is in Atlanta. Most of the clients they serve are remote and have no need to connect directly to the company's LAN; they just use Internet mail and VPNs to access the intranet. With the merger, it has been decided that you will expand the forest by creating two domains: one in Europe and the other in the United States. To improve performance and accessibility, you will create sites at each major management location and link them all for Active Directory replication. Each management location only has 10 to 30 people, and most are connected with T1 Internet access. Use the information provided in the following table to determine the minimum number of sites and DCs you need.

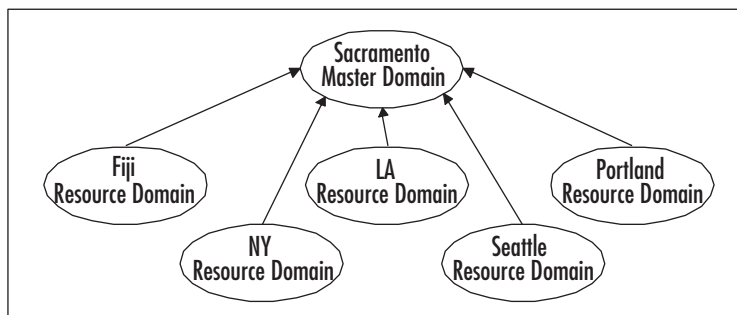
Europe	WAN Speed	United States	WAN Speed
Berlin	E3	Atlanta	T3
Amsterdam	E1	New York	T3
Paris	E1	Chicago	T1
Zurich	E3	Dallas	T1
Rome	E1	Portland	T1

Figure 7.44 Multisite Domain




- A. Ten sites, with two DCs in each site
 - B. Ten sites, with one DC in each site
 - C. Eight sites, with one DC in each site
 - D. Eight sites with two DCs in each site
 - E. Eleven sites with two DCs in each site
4. Currently, the POTC Company uses Windows NT 4. They have a single-master domain structure with five resource domains (see Figure 7.45). The IS in Oakland manages all except the offshore connection in Fiji, where most everything is in French. The POTC Company has a chance to improve on their multidomain network as they migrate to Windows Server 2003. Using the information provided, determine which domains can become OUs and which must remain a domain.

Figure 7.45 Single-Master NT Domain



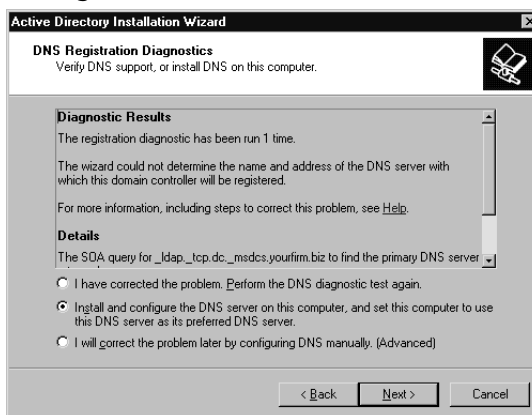
- A. Create one domain incorporating all five resource domains into OUs.
 - B. Create two domains: one for the root, Sacramento, and one subdomain for the five resource domains, creating OUs for each location.
 - C. Create three domains: Sacramento with LA, Portland, and Seattle as OUs; Fiji as a subdomain, and New York as another subdomain.
 - D. Create two domains: Sacramento with LA, Portland, Seattle, and New York as OUs; Fiji as a subdomain.
5. Referring to Figure 7.46, determine the minimum number of DCs required. Each oval represents a physical location and lists the WAN connection speed available at that site. The arrows indicate the proposed replication strategy.

Figure 7.46 Site Topology



- A. Ten, one at each site
 - B. Twelve, one at each site plus one more in Berlin and Atlanta
 - C. Twelve, two at each site plus one more in Berlin and New York
 - D. Twenty, two at each site
6. You are installing Windows Server 2003 and promoting it to the first DC of your new domain, BusyBees.biz. During the Active Directory Installation Wizard process you get the dialog box in Figure 7.47. What is the solution to the problem? (Choose all that apply.)

Figure 7.47 DNS Diagnostics



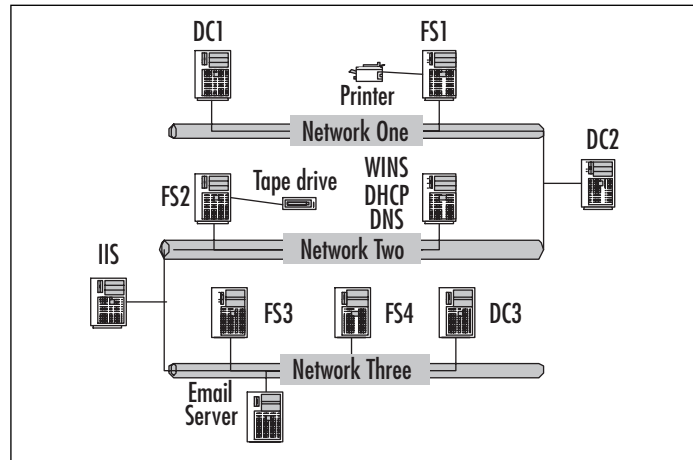
- A. Do nothing. The ADIW will create a DNS server for you.
- B. Cancel ADIW. Install a DNS server that supports RFC 2136 (dynamic updates).
- C. Cancel ADIW. Install a Windows 2000 DNS server using the defaults.
- D. Cancel ADIW. Install a Windows 2000 DNS server. Create a primary zone called BusyBee.biz and enable dynamic updates.
- E. Cancel the ADIW. Install a Windows 2000 DNS server. Create a primary zone called BusyBee.biz and don't enable dynamic updates.

Backing Up Domain Controllers

- 7. Mark is the local administrator for the site in Portland. His duties include the backups for the servers in his site, using Windows Backup. His site includes a DC that he does not back up because the DC in San Francisco is backed up and all Active Directory replications come to Portland once a night. When Mark loses his DC in Portland to a lightning strike, he replaces the server and now wants to restore the computer to a DC. What is the simplest and fastest way to do this?
 - A. Restore the system state from a DC in San Francisco.
 - B. Promote the server to a DC using ADIW.
 - C. Ship the server to San Francisco and have the dcpromo run there and then ship it back.
 - D. Set up a VPN and then run dcpromo from San Francisco.
- 8. Stephanie is the administrator for the scrapbook company, Book On Over Co. (BOOC), which was recently bought by their competitor, Buecher Sind Toll GMBH (BST, a German company). Consequently, the two Windows Server 2003 domains were brought into one tree with two domains. Manfred, the systems engineer for BST, recently performed an authoritative restore of the Active Directory in his domain successfully and informed Stephanie of it. Now the Managers group in the BOOC domain can no longer access data on the Forms server in the BST domain. Based on the information given, the authoritative restore seems to have caused the problem. What is the likely cause of this problem?
 - A. The Managers group was deleted by the authoritative restore.
 - B. The authoritative restore removed the Trust between the domains.
 - C. The authoritative restore replaced the Security ID (SID) of the Managers group to an old SID that makes it no longer valid in either domain.
 - D. The password used by the Trust between the two domains was changed to an old password by the authoritative restore.

9. Using the diagram in Figure 7.48, determine which data can be included in the daily backup routine to the tape device connected to FS2.

Figure 7.48 Backing Up a Domain



- A. Net 1: AD on DC1; FS1
 Net 2: FS2; IIS
 Net 3: Email; FS3; FS4
- B. Net 1: FS1
 Net 2: AD on DC2; FS2; IIS
 Net 3: Email; FS3; FS4
- C. Net 1: FS1
 Net 2: FS2; IIS
 Net 3: AD on DC3; Email; FS3; FS4
- D. Net 1: FS1
 Net 2: FS2; IIS
 Net 3: Email; FS3; FS4
10. Brayden is the domain administrator for a multisite Windows Server 2003 domain. The headquarters is located in South Bend, Indiana. A new branch is being opened remotely in San Jose, California. Brayden needs two DCs to place at the new San Jose site. The WAN link won't be up for two more weeks, but he wants to get the DCs online and in place this week so his San Jose technicians can begin setting up the workstations in San Jose right away. What can Brayden do to create those DCs before the WAN is installed?


- A. Create the two DCs in South Bend, and then ship the servers to San Jose.
- B. Create a backup of a DC in South Bend to a CD or DVD and ship it to San Jose.
- C. Create the DCs in San Jose, and then when the WAN link is installed, synchronize them with the DCs in South Bend.
- D. Nothing. Brayden must wait for the WAN link before creating the remote site's DCs.

Managing Operations Masters

- 11. James comes to work on Monday and opens the Active Directory for Users and Computers. His task today is to create three new users and create a new group. James attempts this and it fails repeatedly. He knows that one DC went down over the weekend, but he is not connected to that DC and can see all the objects in Active Directory. Users are logging on just fine as well. What is a possible explanation for not being able to create new objects in Active Directory?
 - A. James is not logged on as a Schema Admin.
 - B. The DC that went down had the Domain Naming FSMO on it.
 - C. The DC that went down had the RID FSMO on it.
 - D. The DC that went down had the Schema FSMO on it.
 - E. The DC that went down had the PDC Emulator on it.
- 12. Ryan is a domain administrator for Astronauts Ltd. It is a multidomain tree with five sites. Today, he must add some users to the Marketing group. He uses ADUC to open the group and adds the users, Brayden and Hannah, from the SD.CA.COM domain. The users Rebecca and McKay are already members of this group from the LA.CS.COM domain. In testing the access of these users to the Contact database used by the Marketing department, Ryan finds that the users Brayden and Hannah are still unable to access the database, while Rebecca and McKay can. Which of the following is an applicable troubleshooting step in diagnosing this problem?
 - A. Verify that the group is a distribution list.
 - B. Verify that the group is a local group.
 - C. Verify that the RID FSMO is online and available.
 - D. Verify that the Infrastructure FSMO is online and available.
- 13. As an enterprise administrator for the Sports Agents of America (SAA), you must migrate the newly acquired agency's domain into your existing forest as a child domain to SAA.us. The new agency is called Alternative Sports, Inc. The new Windows Server 2003 domain is called AS. Figure 7.49 shows the current domain and

site topology of SAA.us. To set up the migration, your first step is to create the child domain, AS.SAA.us. This fails repeatedly. What is a possible reason for this?

Figure 7.49 Sports Agency of America Domain Tree



- A. The Domain Naming FSMO located in the Montana site is offline.
 - B. The Schema FSMO in the Montana site is offline.
 - C. The FSMOs for AS.SAA.us need to be created before you can create a child domain.
 - D. The Infrastructure FSMO is unavailable.
14. Michael is an enterprise administrator for NuttyNuts, Inc. He is installing Microsoft Exchange 2000 into his domain. His domain, nuttynuts.biz, has two sites and one child domain: CA.nuttynuts.biz, a subsidiary in Sacramento, California. Michael logs on to the domain with his focus on a local DC and as a member of the Enterprise Admins group. During the Exchange installation, he runs across errors that restrict him from completing the installation. Which is a possible reason for this problem?
- A. Exchange 2000 cannot run on Windows Server 2003 domains because the schemas are incompatible.
 - B. The RID FSMO is unavailable.

- C. The Domain Naming FSMO is unavailable.
 - D. Michael must log on as a member of the Schema Admins group.
15. Heather has been hired to come into your company and install a customized Directory-enabled application. Only the users in your branch office located in Fresno, California use this application. Your headquarters is in Santa Rosa, California, and you created a site for each location and set up directory replication over the slow WAN link to occur only at night. Access between the sites occurs at that time, but occasionally you allow the sites to connect during the day when a certain threshold of requests is reached. You create a temporary account for Heather and place the new account in the Schema Admins group. Heather begins to install the application but soon realizes that the schema will not let her extend it, as the application requires? Which is a possible reason for this?
- A. She must install the application in Santa Rosa and then set up Terminal Services for the users in Fresno to access the application remotely.
 - B. She needs to wait for the schema extension requests to be processed between the two sites.
 - C. The Schema FSMO is unavailable.
 - D. The schema can only be extended on the DC that holds the Schema FSMO.

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|-------------------|--------------|
| 1. A | 9. D |
| 2. B, G, E | 10. A |
| 3. B | 11. C |
| 4. D | 12. D |
| 5. D | 13. A |
| 6. A, B, D | 14. D |
| 7. A | 15. C |
| 8. D | |

MCSA/MCSE 70-294

Working with Global Catalog Servers and Schema

Exam Objectives in this chapter:

- 2.1.3 Add or remove a UPN suffix.
- 1.1 Plan a strategy for placing global catalog servers.
 - 1.1.1 Evaluate network traffic considerations when placing global catalog servers.
 - 1.1.2 Evaluate the need to enable universal group caching.
- 2.1.2 Manage schema modifications.

- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

Active Directory uses the Global Catalog (GC), which is a copy of all the Active Directory objects in the forest, to let users search for directory information across all the domains in the forest. The GC is also used to resolve user principal names (UPNs) when the domain controller (DC) that is authenticating logon isn't aware of the account (because that account resides in a different domain). When the DC can't find the user's account in its own domain database, it then looks in the GC. The GC also stores information about membership in Universal Groups.

Because the GC performs all these functions for the multidomain network, it is important for administrators to understand how it works and how to create, manage, and place the GC servers that hold the GC. In this chapter, we look at this special type of DC: the GC server. You'll learn about the role the GC plays in the network, and how to customize the GC using the Schema Microsoft Management Console (MMC) snap-in. We show you how to create and manage GC servers, and explain how GC replication works. You'll learn about the factors to consider when placing GC servers within sites.

Next, we address the Active Directory schema itself. You'll learn about schema components: classes and attributes, and the naming of schema objects. We show you how to install and use the Schema management console, and you'll learn how to extend the schema and how to deactivate schema objects.

Working with the Global Catalog and GC Servers

The GC is a vital part of Active Directory functionality. Given the size of enterprise-level organizations, on many networks, there will be multiple domains and at times, multiple forests. The GC helps in keeping a list of every object without holding all the details of those objects; this optimizes network traffic while still providing maximum accessibility.



NOTE

The first DC in a domain becomes the GC server by default.

Whenever a user is searching for an object in the directory, the GC server is used in the querying process for multiple reasons. The GC server holds partial replicas of all the domains in a forest, other than its own (for which it holds a full replica). Thus, the GC server stores the following:

- Copies of all the objects in the domain in which it resides
- Partial copies of objects from other domains in the forest



NOTE

When we say that the GC server holds a partial copy of an object, we mean that it includes only some of the object's attributes in its database. Attributes are object properties, and each object has a number of attributes. For example, one attribute of a *User Account* object would be the username. You can customize the attributes of a particular object type by editing the *schema*, which we discuss later in this chapter.

The key point is that the GC is designed to have the details that are most commonly used for searching for information. This allows for efficient response from a GC server. There is no need to try to find one item out of millions of attributes, because the GC has the important search-related items only. This makes for quick turnaround on queries.

Functions of the GC

The GC serves various purposes, which we discuss later in the chapter. GC servers are important for the UPN functionality of Active Directory. Universal Groups are also a responsibility of the GC server.

The scope of Directory Services has changed from the days of Windows NT 4.0 Directory Services. With Active Directory, a user record holds more than just a username for an individual. The person's telephone number, e-mail address, office location, and so forth can be stored in Active Directory. With this type of information available, users will search the directory on a regular basis. This is especially true when Microsoft Exchange is in the environment.

Whether a person is looking for details on another user, looking for a printer, or simply trying to locate another resource, the GC will be involved in the final resolution of the object. As mentioned previously, the GC server holds a copy of every object in its own domain and a partial copy of objects in other domains in the forest. Therefore, users can search outside their own domains as well as within, something that could not be done with the old Windows NT Directory Services model.

EXAM
70-294

OBJECTIVE
2.1.3

UPN Authentication

The UPN is meant to make logon and e-mail usage easier, since the two (your user account and your e-mail address) are the same. An example of a UPN is Brian@syngress.com. The GC provides assistance when a user from a domain logs on and the DC doesn't know about the account. When the DC doesn't know the account, it generally means that the account exists in another domain. The GC will help in finding the user's account in Active Directory. The GC server will help resolve the user account so the authenticating DC can finalize logon for the user.



EXAM WARNING

With Windows Server 2003 and beyond, you will see more and more references to UPN use in single or multiple domain environments. Be sure to understand how the UPN works in relation to logon, and how the GC keeps this information available efficiently.

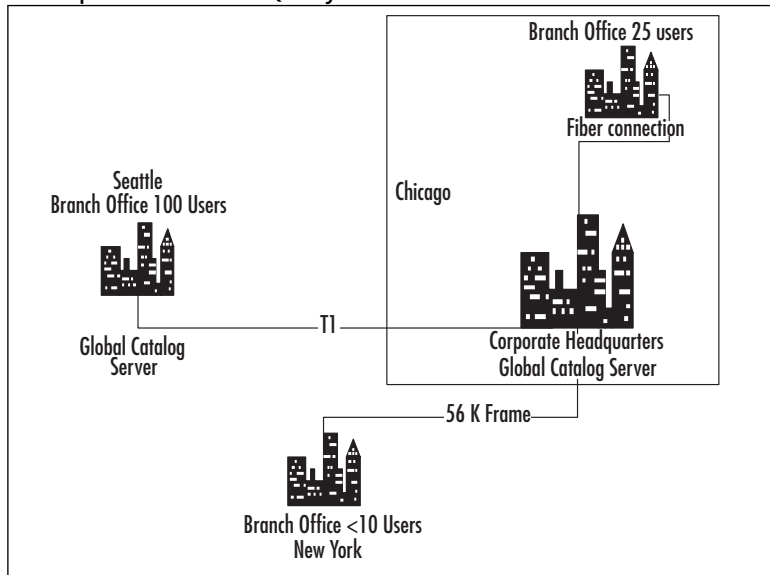
Directory Information Search

With Active Directory, users have the ability to search for objects such as other users or printers. To help a user who is searching the database for an object, the GC answers requests for the entire forest. Since the complete copy of every object available is listed in the GC, searches can be completed quickly and with little use of network bandwidth.

When you search the entire directory, the request is directed to the default GC port 3268. The GC server is also known to other computers on the network because of SRV records in DNS. That is how a node on the network can query for a GC server. There are SRV records specifically for GC services. These records are created when you create the domain.

When users search for information in Active Directory, their queries can cross WAN links, depending on the network layout. Each organization is different. Figure 8.1 shows an example layout with GC servers in the corporate office in Chicago and a branch office in Seattle. The other two sites do not have GC servers. When queries are initiated at the Chicago branch office, the queries use the corporate office GC server. With a high-speed fiber connection, bandwidth isn't an issue.

Figure 8.1 Example GC Search Query



The branch office in New York has a slow link but less than 10 users. These users will use the GC in Chicago as well. Even though the pipe between these locations is only 56K, the minimal amount of users doesn't warrant having a GC server in New York. The Seattle office has a T1, which is decent connectivity, but there are over 100 users in this location. Considering that, searches will be more efficient with a GC server locally. We will look at sites later in the chapter, but Figure 8.1 will help you get a basic understanding of how the query process works.



EXAM WARNING

Be prepared to see diagrams similar to Figure 8.1 that show network layouts and the various GC servers you have on your network. Part of being a successful network administrator is being able to determine if the design is good.

Universal Group Membership Information

When setting up your network, you will have certain features available based on the Forest Functional Level and Domain Functional Level. Universal Groups is one of these features that will or will not be available depending on your functional level. If your Domain Functional Level is set to at least Windows 2000 Native or later, you will have Universal Groups available on your network. Universal Groups can have members belonging to various domains in the forest. Without a GC server, Universal Groups could not exist. That is because Universal Group membership is stored in the GC only. This means that every DC will not have a copy of Universal Group membership; only the DCs serving as GC servers have this information. When a user logs on, his Universal Group membership is checked. The GC provides this information to the authenticating DC.

Universal Group membership information is stored in all GC servers, so you need to consider the design of your GC server layout when adding to or changing the GC server configuration. The number of users at a location will help determine when you need a GC server. A large number of queries of the GC information over slow links isn't recommended; placing a GC at each site is a better design. With sites with a small number of users, you can get away with not having a GC server at each site. We discuss this in more detail later in the section *Placing GC Servers within Sites*.

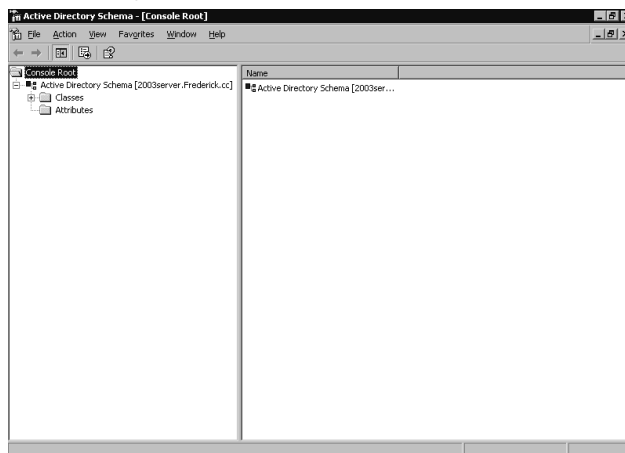
Customizing the GC Using the Schema MMC Snap-In

There might be occasions when you need to make a modification to the GC. You might want to include more attributes than were originally set up. You have to be careful, though, and consider the replication of data. The more attributes there are for the GC servers to replicate, the more network traffic is generated.

To modify the GC, use the Schema snap-in within the MMC. Before you can run the console, you must install it. You complete the installation by registering a .dll. To install the Active Directory Schema snap-in, open a command prompt and type **Regsvr32 schmmgmt.dll**.

You should then see a message that the dll was registered (“DllRegisterServer in schmmgmt.dll succeeded”). Now you can run the Active Directory Schema snap-in as shown in Figure 8.2.

Figure 8.2 Active Directory Schema Snap-In



Now that the .dll is registered, you can create a custom MMC. If you click **Start** and select **Run**, you can start a blank console by typing **MMC** in the Run window and clicking **OK**. You have to add the snap-in to the blank MMC. Exercise 8.01 walks you through the steps of registering and running the Active Directory Schema snap-in. Remember that you must be a Schema Admin member to make changes to the schema. If you are not, you will be able to run the Schema Admin snap-in and view properties of classes and attributes, but you won't be able to make changes.

EXERCISE 8.01

ACTIVE DIRECTORY SCHEMA SNAP-IN SETUP AND USE

This exercise gives you some hands-on experience with installation of the Active Directory Schema snap-in. You need to be logged on as an Enterprise Administrator in Active Directory.

1. Log on to your server with an Enterprise Administrator account.
2. Open a command prompt by Clicking **Start**, and then select **Run**.

3. In the **Run** box, type **cmd** and press **Enter**.
 4. At the command prompt, type **regsvr32 schmmgmt.dll**.
 5. You should see a box that shows registration if the dll was successful.
 6. Click **OK** in the dialog box confirming that the registration succeeded.
 7. Now, click **Start**, type **mmc /a**, and press the **Enter**.
 8. In the MMC window, click on **File** and select **Add/Remove Snap-in**.
 9. In the **Add/Remove Snap-in** window, click **Add**.
 10. Find the **Active Directory Schema** snap-in listed in the **Add Standalone Snap-in** window.
 11. Select the snap-in and then click **Add**.
 12. Now, click **Close** in the **Add Stand Alone Snap-in** dialog box.
 13. Click **OK** in the **Add/Remove Snap-in** dialog box.
 14. You should now have a console that you can use for modifying the schema or GC. You can save this as a .msc file to easily click on it next time versus adding a custom snap-in.
-

Creating and Managing GC Servers

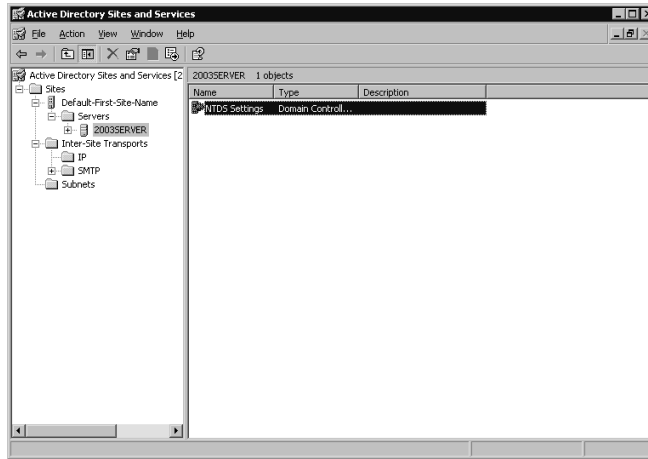
When you initially install Active Directory, the first DC created is also the first GC server. As your network changes, you might require additional GC servers to help manage network traffic. To specify whether a server is a GC server, you use the Active Directory Sites and Services console. Open the Active Directory Sites and Services console, expand **Sites**, and then expand the site with the DC you want to be a GC server. Next, expand **Servers** and find the *Domain Controller* object.

In the details pane you should see NTDS Settings. Figure 8.3 shows an example of the Active Directory Sites and Services screen.

If you right-click **NTDS Settings** and select **Properties**, you will have the option to enable or disable the GC on the DC you select, as Figure 8.4 shows.

The check box on the General tab is used to enable or disable GC functionality. To be able to change the state of the GC check box, you must be a member of the Domain Admins group or the Enterprise Admins group.

As stated previously, the planning of GC server placement is important for a successful network. Each GC server creates additional replication traffic on the network.

Figure 8.3 NTDS Settings in Active Directory Sites and Services**Figure 8.4** General Tab of NTDS Settings Properties**TEST DAY TIP**

Getting comfortable with the Schema snap-in will prepare you if you encounter point-and-click questions or actual hands-on questions on the exam.

Understanding GC Replication

You know now that GC servers hold information for all of the objects in their own domains and a partial copy of the objects from other domains in the forest. For this to be possible, some type of replication has to happen between the GC servers. The default attributes included in the GC make up the most commonly searched for items. These items are part of normal Active Directory replication.

The Knowledge Consistency Checker (KCC) generates the GC replication topology. The GC is only replicated between DCs that are GC servers; the information is not replicated to other DCs. A few things can affect replication; for example, Universal Group membership, and the number of attributes included in the GC.

Universal Group Membership

The GC holds the sole responsibility of maintaining Universal Group membership. The names of the Global Groups and Domain Local Groups are also in the GC, but their membership lists are not. This helps keep the size of the database small enough to efficiently answer queries.

For replication purposes, it is best to keep Universal Group membership relatively static. Every change made to a Universal Group is replicated to every GC server. Keeping these changes to a minimum will keep the GC replication traffic to a minimum.



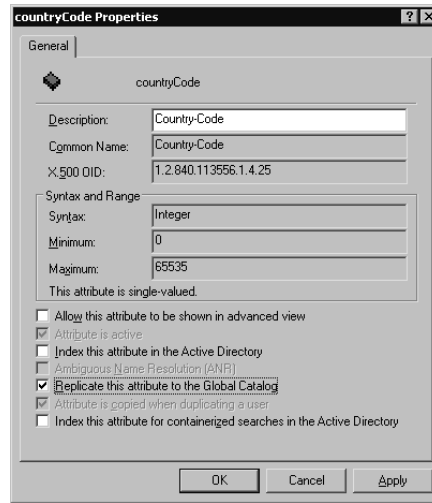
TEST DAY TIP

Universal Groups can only exist if the functional level of your network is Windows 2000 native or later. Universal Group information is replicated between GC servers. Replication traffic can consume bandwidth, which is why site topology is important; putting a GC at each site keeps replication traffic to a minimum.

Attributes in GC

When you first set up Active Directory, there is a series of default attributes from Active Directory in the GC. Sometimes, the default set of attributes is missing an item you would like to see. For example, perhaps you want to have a coworker's department number as part of his user record; you can accomplish this by adding an attribute. You can use the Active Directory Schema snap-in to include additional attributes in the GC by using the General tab as shown in Figure 8.5. To get to this option, open the **Schema** snap-in, and expand the **Attributes** section. Right-click any attribute, and select **Properties**.

Figure 8.5 Adding Attributes to the GC



Prior to Windows Server 2003, each time the attribute set is extended, a full synchronization of all attributes stored in the GC is completed. In a large network, this can cause a serious amount of network traffic. With Windows Server 2003, only the additional attribute or attributes are replicated to other GC servers. This makes more efficient use of network bandwidth.

Placing GC Servers within Sites

Another consideration when it comes to replication is placement of your GC servers. In a small network with one physical location, GC server placement is easy. Your first DC that is configured will hold the GC role. If you have one site, but more than one DC, you can move the role to another DC if you want to. Most networks today consist of multiple physical locations, whether in the same city or across the country. If you have high-speed links connecting your branch offices you might be okay, but many branch office links use limited bandwidth connections. If the connection between locations is less than a T1, you might have limited bandwidth depending on what traffic is crossing the wire. As a network administrator, you will have to work with your provider to gauge how much utilization there is across your WAN links.

Another factor is reliability. If your WAN links are unreliable, replication traffic and synchronization traffic might not successfully cross the link. The less reliable the link, the more the need for setting up sites and site links between the locations.

Without proper planning, replication traffic can cause problems in a large network. Sites help control replication traffic. Making the most of available bandwidth is an important factor in having a network that allows your users to be productive. Logon and searching Active Directory are both affected by GC server placement. If users cannot find the information they need from Active Directory, they might not be able to log on or find the information or data they need.

GC in an Exchange Server Environment

Now that Active Directory is the single directory used in Windows 2000 and Windows Server 2003 networks, there is very tight integration with Microsoft Exchange. Prior to Exchange 2000, Exchange had its own directory and the domain had its own directory service. There were links between the two, but they were still technically separate directories.

Because all user information (first name, last name, and contact information) is kept in Active Directory, users will be searching more and more throughout the directory. In previous versions of Exchange, there was a Global Address List that you could search to locate people within your organization. Information such as telephone numbers, fax numbers, and office locations can be part of your GC strategy with Windows Server 2003.

It is important for administrators to ensure that users can reach the data for which they are searching as quickly and easily as possible. Proper planning and location of your GC information is important to successful queries of your directory information.

EXAM
70-294
OBJECTIVE
1.1.1

Bandwidth and Network Traffic Considerations

Active Directory replication works differently depending on whether it is *intersite* or *intrasite* replication. DCs that are part of the same site (intrasite) replicate with one another more often than DCs in different sites (intersite). If you have sites that are geographically dispersed, you need to be careful how you handle your GC server placement. The bandwidth between geographically dispersed offices is often minimal. The rule of thumb is to have GC servers in selected sites. In most cases, you do not want to have a GC server in every site because of the vast amount of replication that would occur. The following examples describe situations in which you should have a GC server within a site:

- If you have a slow WAN link between geographic locations and a domain controller at each location, you'll want to also configure domain controllers at these locations to be global catalog servers.
- If you have an application that relies heavily on GC queries across port 3268, you'll want to have a GC server in the site that the application runs in. An example of this is Exchange 2000, which relies heavily on GC information.
- If the domain functionality level is Windows 2000 native or later, you'll want to have GCs in as many sites as possible because Universal Group membership comes into play. We look at caching of Universal Groups, which can reduce traffic related to this, in the next section.



TEST DAY TIP

Microsoft's documentation recommends that if you have 50 or more users at a given location, you should give that location a DC serving as a GC server. This will help reduce the number of queries crossing the WAN for Active Directory object searches.

Data replicated between sites is compressed, which makes better use of available bandwidth. Because the data is compressed, more can be sent over a limited amount of bandwidth. This is how site placement and design can be critical to efficient network operation.

EXAM
70-294
OBJECTIVE
1.1.2

Universal Group Caching

The Windows Server 2003 Active Directory introduces Universal Group caching as a new feature. When a user logs on to the network, his membership in Universal Groups is verified. For this to happen, the authenticating DC has to query the GC. If the GC is across a WAN link, the logon process will be slow every time. To alleviate this, the DC that queries the GC can cache this information, which cuts down on the amount of data traveling across the WAN link for Universal Group information.



EXAM WARNING

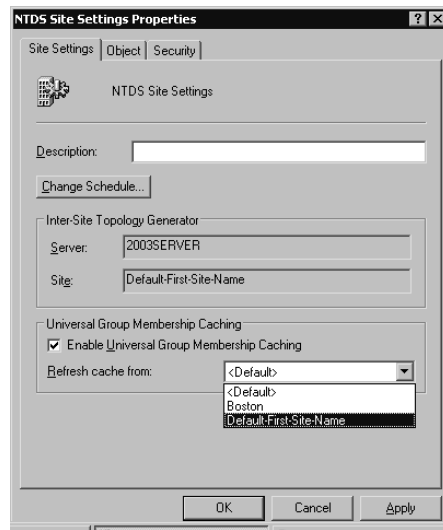
Because Universal Group caching is a new feature in Windows Server 2003, you'll probably see it mentioned in at least one question on the exam. Be sure to understand the purpose of Universal Groups and the caching of that information.

The cache is loaded at the first user logon. Every eight hours by default, the DC will refresh the cache from the nearest GC server. Caching functionality is administered in Active Directory Sites and Services as shown in Figure 8.6, and can be turned off if desired. You can also designate the GC server from which you want the cache to refresh, giving you more control over traffic distribution on the network.



NOTE

The NTDS Site Settings Properties box is not the same NTDS Settings Properties box you accessed to make a DC act as a GC. Instead of accessing the properties of NTDS settings under the DC node in the Servers container, you must access the properties of NTDS Site Settings in the right console pane when you select a site name (for example, Default-First-Site-Name). The similarity of these two settings can be confusing if you haven't worked with the console much.

Figure 8.6 Configuring Universal Group Caching

Prior to Windows Server 2003, Active Directory logon failed if a GC could not be located to check Universal Group membership. With Universal Group caching, DCs cache complete group membership information, so even if a GC server cannot be reached, logon will still happen based on cached Universal Group information.

New & Noteworthy...

Universal Group Caching

Windows Server 2003 Active Directory introduces a new feature that allows DCs to keep a copy in cache of group membership information. If a user logs on and cannot find a GC server to resolve Universal Group membership, the user still has a way to connect and work. This is facilitated by Universal Group caching at the DC. If the authenticating DC cannot find a GC server to query for Universal Group membership, the DC will look in its cache for that membership information. This is an important difference between Windows 2000 and Windows Server 2003 based networks, and eliminates the common problem in Windows 2000 networks of users not being able to log on if a GC is down.

Troubleshooting GC Issues

As with anything on your network, you will spend a certain percentage of your time troubleshooting issues. Common issues with GC include:

- Replication latency between GC servers
- Slow query response
- Overall load too high

Determining the proper course of action depends on the problem you are encountering. The basic answer to any of the preceding issues will result in either moving or adding GC servers.

If you are experiencing replication latency, you need to look at your GC servers and possibly your site configuration. Adding sites might help with replication traffic, as traffic between sites replicates differently than it does within sites. Between sites, the data is compressed and will cut down on bandwidth usage, which could help with the latency problem.

Slow query response can be a result of slow links between locations. Adding a GC server to the location experiencing this problem if one isn't already present will help. In addition, checking your site configuration can help as well. Workstations within a site will query the local GC server versus going across a WAN link.

If the overall load is too high, you need to look at adding more GC servers to balance the traffic. This also results in more replication traffic if you are not careful, so planning and consideration of the impact on your network is important.

There is not one single answer to troubleshooting your GC and the servers involved. Even with proper planning, problems sometimes arise. Working through the problem and testing will take time, and location of GC servers can make all the difference in a successful Active Directory layout.

EXAM
70-294
OBJECTIVE
2.1.2

Working with the Active Directory Schema

To have a directory, you must have a framework on which your directory structure is based. The Active Directory schema defines your directory. For an object to store data such as a username or telephone number, there has to be a field in which this information can be entered. These fields have names and belong to another component of the schema, which we will look at shortly.

It is important to remember that there is only one schema per forest. Thus, all the domains in the forest share this single schema. Schema information is the backbone of your Active Directory and the data within. If problems develop in the schema, your entire network could be out of commission. This means that you must be very careful with the schema, which is why only members of the Schema Admins groups have write permission to the schema. The only default member of that group is the Administrator account in the forest root domain. You should be very selective when choosing members of the Schema Admins group.

Understanding Schema Components


We started to look at the schema from a database perspective at the beginning of this section. Every database is built on a foundation that defines its structure, the database model. The foundation, or model, is based on various components (see Figure 8.7). The first component of the foundation for the schema in Active Directory is the *class*. Objects in the database belong to classes. A little later in the chapter, we look at how the class defines the structure of your data.

Associated with each object are fields, or *attributes*, that you fill in with data. We'll look at how this component relates to the class, and then bring it all together. For now, remember that the important components of the schema are:

- Object classes
- Object attributes

As with any database, there must be a naming standard. Within this section, we will look at schema object naming and how information can be referenced in different ways. If you decide that you need to extend the schema by adding additional classes or attributes, you need to plan exactly how you want it done. It is important to take extreme care in extending the schema. In most circumstances, software installations will extend the schema rather than an administrator manually editing it. Software such as Microsoft Exchange or Microsoft SQL can use the schema and sometimes require that changes be made. This is generally done as part of the software installation process and is the reason why you have to be logged on with a particular user account when installing that type of software.

Figure 8.7 Schema Components Diagram





EXAM WARNING

Be sure to review the relationship between a class and an attribute. They are separate entities within a Windows Server 2003 domain's schema that are related based on each class having certain attributes associated with it.

Attributes

You need to define various attributes for each object you create. Remember that the *Object* class determines which attributes are required and which are optional. All attributes associated with that class will exist, but some (the optional attributes) can be left blank. You will be required to enter information for the required attributes. An example of an attribute is *First Name* or *Telephone Number*. This attribute is associated with the *User Object* class. These attributes can be filled in when you create the user account and are defined in Active Directory as containing a certain type of data. The *AttributeSchema* object defines the characteristics of a given attribute. Configuration items such as common name, syntax rules, and other things make up the *AttributeSchema* object. For example, a *Telephone Number* attribute is generally in a specified format, such as Access code–Area or Country Code–Prefix–Number (for example, 1-512-555-1234). However, the schema is not this specific; it specifies that the syntax must be a Unicode string of characters, with a minimum of one and a maximum of 64 characters.

In addition to defining the syntax, each attribute's properties will also include an X.500 object identifier (OID) for interoperability with other directories that comply with X.500 specifications, and a statement as to whether the attribute is single or multivalued as shown in Figure 8.9.

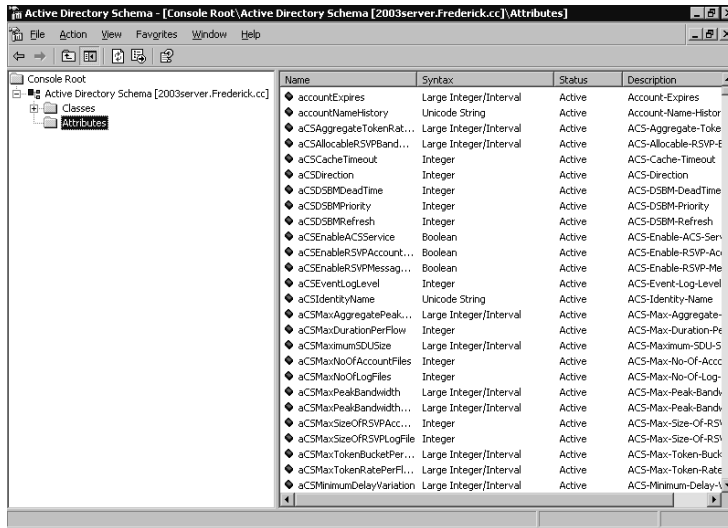
Figure 8.9 Properties of an Attribute

The screenshot shows the 'name Properties' dialog box with the following details:

- General Tab:**
 - Description:** RDN
 - Common Name:** RDN
 - X.500 OID:** 1.2.840.113556.1.4.1
- Syntax and Range:**
 - Syntax:** Unicode String
 - Minimum:** 1
 - Maximum:** 255
- Value Type:** This attribute is single-valued.
- Options:**
 - Allow this attribute to be shown in advanced view
 - Attribute is active
 - Index this attribute in the Active Directory
 - Ambiguous Name Resolution (ANR)
 - Replicate this attribute to the Global Catalog
 - Attribute is copied when duplicating a user
 - Index this attribute for containerized searches in the Active Directory
- Footer:** This is a system attribute; some edits are not allowed.
- Buttons:** OK, Cancel, Apply

You use the Active Directory Schema snap-in to maintain attributes just as you do with objects. Figure 8.10 shows some of the default attributes within Active Directory.

Figure 8.10 Default Attributes in Active Directory



Single-Value Attributes

Most of the attributes you will work with will be single-value attributes. A single-value attribute is just what its name implies; it is an attribute with one piece of data entered. An example would be *First Name*. The *First Name* attribute cannot hold multiple values of data. After the first name is entered, you have to create a completely different object if you want to have another object with a different *First Name* attribute.

Schema Administration

When working with the schema, you have to be very cautious about making changes. In most organizations, only a few people have the ability to modify the Active Directory schema. The Schema Admins group is used to control who has the authority to make modifications to the schema. In most cases, the administrator will not be directly making changes. Many times, software that is installed will make modifications to the schema. Classes and attributes will be added automatically to help support the application.

Another component that assists in the schema not becoming corrupted or duplicated in any way is the Single Master Operations Roles feature of your DCs. In every forest, there is one Operations Master role for schema modifications. When you log on, you can authenticate and do your administrative work on any DC on

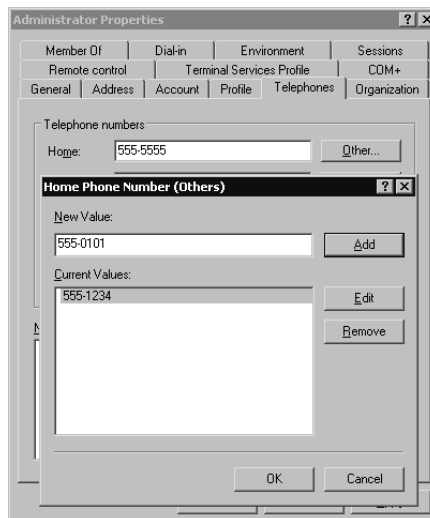
Continued

the network. However, when changes are to be made to the schema, the DC you are working on will contact the DC with the Schema Master role if that DC is not serving in the role itself. The DC that is filling the Operations Master role of Schema Master will then make the actual changes to the schema. Having one DC responsible for this prevents the possibility of two DCs attempting to write to the schema at the same time. There can only be one Schema Master for each Windows Server 2003 forest.

Multivalued Attributes

Although most attributes are of the single-value variety, there are also cases where an attribute will hold more than one piece of information. Attributes such as *Telephone Number* can hold multiple values. When you create a user account or edit its properties, you can enter a main number of 555-5555 for a user, and if that user has a secondary line, you can click the **Other** button to add additional data as shown in Figure 8.11. (You can access the properties for a user account by opening the **Active Directory Users and Computers (ADUC)** administrative tool, clicking the **Users** container in the left console tree, right-clicking the username whose properties you want to edit in the right console pane, and selecting **Properties**.)

Figure 8.11 Multivalued Attributes



Multivalued attributes do not sort or keep track of the order of the entries if there are multiple entries. They are simply there for convenience in the case of common attributes that can have more than one entry. Each value within a multivalued attribute must be unique.

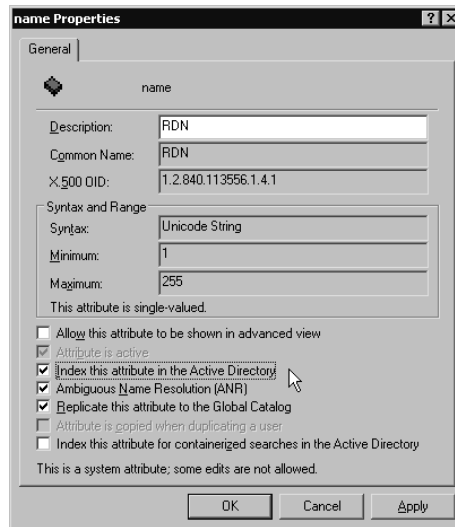
Indexing Attributes

When you index data in a database, you are organizing the information so you can have efficient responses to queries based on that data. You can set attributes as *indexed* to help users find the information they need. This means that the attribute will be indexed in the Active Directory. With indexing attributes, wildcard searches will function, allowing the user the ability to enter a partial word with an asterisk and return multiple hits.

When deciding which attributes to index, you have to be careful because you can slow your network down with extra replication traffic. When you mark an attribute as indexed, every attribute in that instance is added to the index. For example, if an attribute such as *Location* is part of a *Printer* object and a *User* object, both objects would be added to the index. With multivalued attributes, you could be using more bandwidth because you are replicating a large amount of information. The rule of thumb is to only index common attributes.

To index an attribute, use the Active Directory Schema snap-in. Expand the **Attributes** section and right-click on the attribute you want to index. Select **Properties**, and then check the option to **Index this attribute in the Active Directory** as shown in Figure 8.12.

Figure 8.12 Indexing Attributes



EXAM WARNING

Be sure you know how to set up indexing of attributes. In addition, remember that you can add attributes to the GC. Indexing is enabled on the General tab of the Properties of the attribute.

Naming of Schema Objects

If you are going to be working with *Schema* objects a lot, you need to be comfortable with the naming conventions that apply to *Schema* objects. There are different ways to reference objects in Active Directory, the most common of which are:

- **Lightweight Directory Access Protocol (LDAP)** LDAP is the primary access protocol for Active Directory. LDAP is an industry standard protocol for commonality among directories, and is based on the ISO's X.500 directory naming conventions. LDAP names identify an entire path within the directory; for example, CN=JDoe, OU=Sales, DC=Frederick, DC=cc.
- **Common name** The common name is a simplified way to identify an object. Common names are much easier to read than LDAP names. Common names must be unique within the container. An example common name would be JDoe.
- **Object Identifier (OID)** An identification number issued by another authority. The International Organization for Standardization (ISO) and American National Standards Institute (ANSI) have developed standards for OIDs as part of the X.500 directory services specifications. Every OID is unique. An example is the *Department* attribute in Active Directory. The OID for Department is 1.2.840.113556.1.2.141. This same OID will be used for the *Department* attribute in any directory that follows X.500 standards.

You should follow the LDAP or common name naming standards when setting up *Schema* objects. If you write software that modifies the schema, certain standards must be followed for the software to meet the "Certified for Windows" requirements. If you stick with the standards, any changes you make will be less likely to cause problems.

What Is x.500?

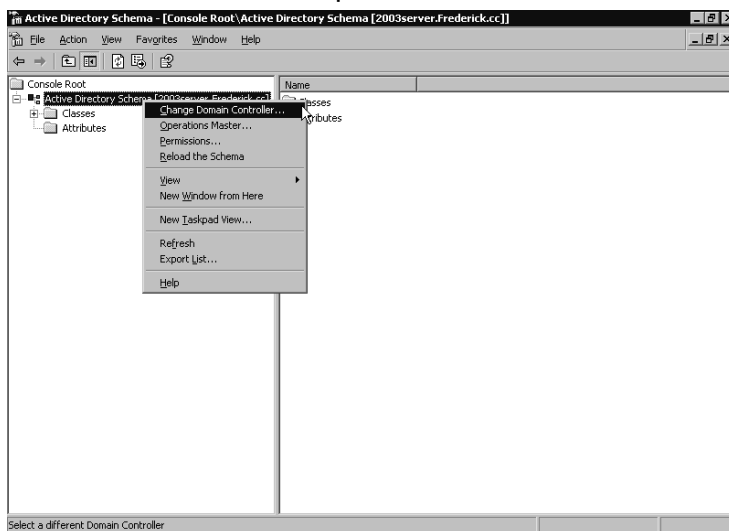
Active Directory is based on X.500 standards, which were first finalized in 1988 and then revised in 1993. X.500 is a directory access protocol that determines how a directory client and a directory server interact. LDAP is based on the original X.500 standards, but is simplified.

As with most standards, the early directory access protocols were developed specifically for applications versus a networkwide or worldwide directory. Eventually, the need for a standard across networks and the Internet forced the development of the X.500 directory access protocol.

Working with the Schema MMC Snap-In

When working with the Schema snap-in, you need to be aware of some other configuration items. If you right-click **Active Directory Schema**, you are presented with various options as shown in Figure 8.13.

Figure 8.13 Schema Administrative Options



If you select **Change Domain Controller**, you can choose what DC you want to feed the schema information. If you select **Operations Master**, you can see what server is holding the **Schema Master** role, or change the server responsible for that role. You can use the Permissions option to change permissions on the schema. In most cases, network administrators would have no reason to go into the Permissions tab. The last option in the first section of the list is Reload the Schema. This option will reload the schema from the database to make sure you don't have cached information that could be outdated.

When working with the Schema snap-in in a mixed environment, you might find yourself at a Windows 2000 server. If you are making schema modifications from a Windows 2000 server, you must ensure that Service Pack 3 for Windows 2000 has been installed.

Modifying and Extending the Schema

There will be times when the default schema layout doesn't meet your needs. If this is the case, you can modify the schema by changing existing classes or attributes. You could also extend the schema by adding classes or attributes that do not exist. Again, you must be extremely careful when making changes to the schema; modifying or extending the schema should only be done when absolutely necessary.

To modify or extend the schema, use the Schema snap-in. Begin by making your changes in a test environment and testing thoroughly before making modifications or extensions on your production network. Remember that the user using the snap-in must be a member of the Schema Admins group. Before you modify the schema by changing or adding classes or attributes, keep the following guidelines in mind:

- Double-check to be certain that the existing schema configuration does not meet your needs. It is possible that there is an existing class or attribute that will work for your requirements.
- When you add a class or attribute, that class or attribute cannot be removed. You can, however, deactivate a class or attribute. We will look at that in the next section.
- Make sure you have a valid OID; do not just pick one out of thin air.
- Default system classes cannot be modified. Windows uses these classes for basic functionality.
- Review documentation on the schema. In particular, review the Active Directory Programmer's Guide, which can be downloaded at www.microsoft.com, if you intend to make extensive modifications or extensions.
- Remember that schema changes affect the entire forest, because only one schema exists in a Windows Server 2003 forest and is shared by all domains in that forest.

When creating a new class, various attributes need to be filled out as shown in Figure 8.14. The first section is the Identification section. You will have to complete both Common Name and LDAP Display Name. You also have to enter the object ID, so you need to know how they are assigned. There is also an optional *Description* attribute that you use if you want to.

The other section is Inheritance and Type. The *Parent* class will have permissions assigned. Being a *Child* class, we would inherit the permissions from the *Parent* class objects.

Figure 8.14 Create a New Object Class

The screenshot shows a dialog box titled "Create New Schema Class". It has two main sections. The first section, "Identification", contains four text input fields: "Common Name:", "LDAP Display Name:", "Unique X500 Object ID:", and "Description:". The second section, "Inheritance and Type", contains a "Parent Class:" text input field and a "Class Type:" dropdown menu with "Structural" selected. At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

Deactivating Schema Classes and Attributes

If changes or additions are made to the schema, they cannot be deleted. Windows Server 2003 does not allow for deletion of classes or attributes after they are defined in the schema. However, you can deactivate a class or attribute if you don't want to use it anymore. This is essentially the same as deletion, because the class or attribute is no longer available for use. However, the class or attribute still exists within the schema. The deactivated class or attribute is called *defunct*. Default classes and attributes cannot be deactivated. If you decide that you need to have the attribute available, you can reactivate it later.

When you deactivate a class or attribute, you can redefine it if your forest is at the Windows Server 2003 functional level. For example, if you have an attribute that has the wrong syntax, you can deactivate the existing attribute and then create a new attribute with the proper syntax. You can reuse the LDAP display name and the OID. Note that you have to rename the original attribute after you deactivate it and before you create the new attribute to prevent conflicts.

You use the Schema snap-in to deactivate or reactivate an attribute or class. Figure 8.15 shows where you can activate or deactivate an attribute.

Figure 8.15 Activating or Deactivating



EXERCISE 8.02

DEACTIVATING CLASSES OR ATTRIBUTES

In this exercise, you will use the Schema snap-in to create an attribute, and then you will deactivate it.

1. Open the **Schema snap-in**.
 2. Expand Active Directory Schema, right-click Attributes, and select Create Attribute.
 3. Click Continue at the warning dialog box.
 4. In the Common Name dialog box, type Telephone number 2.
 5. In the LDAP Name dialog box, type Telephone number 2.
 6. For the OID, type 2.5.4.20.2.
 7. Change the syntax drop-down to Integer, and then click OK.
 8. Now, find the new attribute, right-click, and choose Properties.
 9. On the General tab, you should see a check box for Attribute is Active.
 10. Click the check box to remove the check. Click **Yes** to the question about the making the object defunct.
 11. Click **OK** and the status window in the details pane should show Defunct under the Status column.
-

Troubleshooting Schema Issues

You might run into issues when working with the schema. They could be as simple as not finding the Schema snap-in to not being able to extend the schema. We will look at some different issues you might encounter when working with the schema.

The most common problem is running or finding the snap-in. Make sure you register the snap-in, and then create a customized MMC to run the snap-in.

There might be times where you simply cannot extend the schema; for example, if you are trying to add a class and are unable to complete the operation. A few things could cause this; the most common being that the user trying to make the changes is not a member of the Schema Admins group. In addition, the Schema Operations Master role has to be up and available on the network. If the Schema Operations Master role is across a WAN link, you might be experiencing too much latency. You can move this role if needed to solve network connectivity problems.

You might also experience an issue where you cannot associate an attribute with a class. This is because the schema cache is not up to date. If this happens, you need to make sure the Schema cache is updated by reloading the schema. This could also be caused by trying to make changes on a server other than the Schema Operations Master. When modifying the schema, it is recommended that you make changes on the server running the Schema Operations Master role.

Summary of Exam Objectives

The Global Catalog (GC) server is one of the most important roles played by one or more DCs in your network. It might not appear to do much on the surface, but the GC is responsible for helping resolve names for objects throughout your forest. The GC server holds a copy of all the objects in the domain in which the server is located. That same GC server holds a partial replica of other domains in the forest. The information that the GC holds from other domains includes common search items. This limited but frequently accessed information makes queries very efficient.

GC servers are responsible for UPN authentication. When a user logs on using the UPN, the GC is queried to locate the user account and a domain controller (DC) in the appropriate domain. GC servers are also responsible for answering queries against Active Directory. If a user wants to locate another person within the organization, that user could use his workstation to search Active Directory. The queries are sent to the IP port 3268, which is used for GC communication.

Placement of GC servers has to be considered early in the design process for your network. If you don't determine where you do and do not need a GC server and plan accordingly, you could have communication problems and users could be adversely affected. A good rule of thumb is to remember that if a location has over 50 users, a DC is needed at that location. Dividing the network into *sites* makes a difference in how replication traffic is handled in regard to GC information. Replication within a site (intrasite replication) is handled differently than replication between different sites (intersite replication). Placement of GC servers within every site might not be necessary, but you should keep track of how much bandwidth computers are using. GC queries in large quantities can tie up significant bandwidth.

If the domain functional level is at least Windows 2000 Native, Universal Groups will be available. The GC is the only location in which Universal Group information exists. When users log on, their Universal Group membership is verified. The authenticating DC makes this request of the GC server. If the GC server cannot fulfill the request, logon can be denied. However, with Windows Server 2003, Universal Group membership can be cached to prevent this problem. Caching must be turned on under **NTDS Site Settings Properties** in the Active Directory Sites and Services console as explained previously in this chapter. With this setting turned on, the authenticating DC will query the nearest GC for Universal Group membership. The information received will be cached on the authenticating DC, and refreshed every eight hours by default. With caching enabled, that authenticating DC will be able to process logons in the event the GC cannot be reached because the information has been cached.

The schema defines the structure of your Active Directory. Various types of objects can be administered in Active Directory. An object in Active Directory is an instance of a *class*, such as *User* or *Printer*. A class defines the type of object. Associated with each *Object* class are attributes that can be modified. For example, an attribute can be the *Location* or *First Name*. There are two different types of attributes. The most common is the single-value attribute, which contains one piece of data. You might also work with multivalued attributes,

which can contain more than one piece of data. An example of the latter is a telephone number. The Other button allows you to add additional entries in the event that someone has more than one telephone number.

To speed queries and make searches easier, attribute indexing can be enabled. This process builds an index of every attribute in an instance. Common attributes should be indexed, but not all attributes should be indexed. Special consideration should be given to indexing multivalued attributes. You can produce a lot of extra traffic because of replication of all the multivalued attributes in an instance. When you are working with *Schema* objects, there are different ways you can reference an object. Common ways to describe objects include LDAP names, Common Names, and OIDs. LDAP is an industry standard protocol and the primary access protocol for Active Directory. The Common Name is an easier way to identify an object. The OID is assigned by a third-party authority. There are standards that must be followed in regard to OIDs. We recommend that you follow the naming standards laid out for LDAP and Common Name.

You can use the Schema MMC snap-in to do all modifications in regard to GC and schema. To install the snap-in, you must first register the schmmgmt.dll file; then you can create a custom MMC and add the Schema snap-in. The Schema snap-in is used to extend the schema if the default classes and attributes do not meet your needs. When considering extending the schema, you need to make sure you have tested the changes thoroughly before applying them to a production network. A problem with the schema can mean serious trouble for your network. You must log on as a member of the Schema Admins group to make any modifications or extensions to the schema. The only default member in this group is the Administrator of the forest root domain.

Changes made to the schema cannot be deleted, but they can be deactivated. Windows Server 2003 doesn't allow for deletion of classes or attributes within Active Directory. A deactivated class or attribute is still in the schema database, but is unavailable for use.

Exam Objectives Fast Track

Working with the Global Catalog and GC Servers

- GC servers hold Universal Membership data.
- Universal Membership information can be cached on non-GC servers in Windows Server 2003 networks.
- GC servers assist in searches for objects within the Active Directory.
- The GC handles UPN authentication.
- Dividing your network into sites helps with replication traffic over WAN links.

Working with the Active Directory Schema

- ☑ The schema is made up of *Object* classes such as *User*, *Printer*, and *Server*.
- ☑ Each *Object* class has a series of attributes associated with it.
- ☑ There can be multivalue attributes and single-value attributes.
- ☑ You must be a member of the Schema Admins group to modify the schema.
- ☑ *Schema* objects follow the LDAP or Common Name standards.
- ☑ Classes and attributes cannot be deleted, but can be deactivated.

Exam Objectives Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

Q: I want to enable GC functionality on a DC. Where do I do that?

A: In the NTDS Settings Properties window on the General tab. You simply check the box next to **Global Catalog** and click **OK**.

Q: I have an office with only 10 users. Should I put a GC server at this location?

A: Probably not; Microsoft recommends that 50 or more users at a location constitutes the necessity for a local DC at that office.

Q: I am noticing a large amount of traffic between my corporate office and branch office. I recently added a GC server/domain controller at my branch office. Why all the extra traffic?

A: More than likely, you didn't set up a site for each location. Having GC servers located in sites helps to control replication and should cut down on bandwidth usage. Data is compressed before being sent between sites, which keeps bandwidth usage down.

Q: I am trying to modify the schema but cannot make any changes. Why?

A: Make sure you are logged on as a member of the Schema Admin group. Only Schema Admin members can modify the schema.

Q: What is the difference between a class and an attribute?

A: A class defines the type of object you are working with, such as a *User* object or *Computer* object. The object is associated with various attributes, which are fields of data such as username, first name, location, and so forth.

Q: I want to delete a new attribute I added and cannot find the option. Why?

A: You cannot delete classes or attributes. You can deactivate a class or attribute, which will make the class or attribute no longer available for use although it will still be defined in the schema. It can then be reactivated if you ever want to use it again.

Q: What do you do if your GC server is overloaded?

A: Add another GC server to balance the traffic.

Q: If you cannot modify the schema and you have verified that you are a Schema Admin, what other possible cause is there that will cause schema extensions not to work?

A: The Schema Operations Master might be unreachable

Q: If I have four locations separated by WAN links that are 56K or less, how many GC servers should I have if each location has over 1000 users?

A: In this situation, you should have a GC server at each location and possibly set up sites.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Working with the Global Catalog and GC Servers

1. You are working on your DC and want to be able to run the Schema snap-in. You click on Start and select Run. You type MMC and press Enter. When you go to add the snap-in, you don't see it listed as one you can add. Why?
 - A. The DC you are on is not the GC server, so the Schema Admin snap-in would not be available on that DC.
 - B. You are not a member of the Schema Admins group, so you cannot install the snap-in.
 - C. The DC you are logged on to doesn't serve the role of Schema Master, so the snap-in will not run.

- D. The schmmgmt.dll file has not been registered.
2. You just finished setting up a forest containing three DCs. Server DC1 is the forest root DC. Servers DC2 and DC3 will serve as DCs also. You want to assign the GC responsibility to DC2. How do you determine which DC is serving as the GC server now? (Choose all that apply.)
- A. You can look in the Properties of each *Server* object within the Active Directory Sites and Services administrative tool to determine if the server is the GC server.
 - B. You know that DC1 is the GC because the first DC set up in the forest automatically takes the role of GC.
 - C. You can look at the Properties of NTDS Settings under each *Server* object within Active Directory Sites and Services.
 - D. You know that DC3 is the GC server because the third DC takes role of GC away from the forest root server upon being added to the domain.
3. You have a new attribute that needs to be added to the GC. You have the Schema Admin snap-in open. How do you make sure an attribute is included in the GC?
- A. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Replicate this attribute to the GC** is selected.
 - B. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Allow this attribute to be shown in advanced view** is selected.
 - C. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Index this attribute for containerized searches in the Active Directory** is selected.
 - D. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Index this attribute in the Active Directory** is selected.
4. You recently made your new staff member a member of the Universal Group named Enterprise Admins. The new staff member is located at a branch office. When the user logs off and then back on, he notices that he cannot get to some of the Administrative tools. You recently added the user to the Universal Group and you have a 56K link between your branch office and your main office. Your GC server is at the main office. What could be the problem? (Choose all that apply.)

- A. You cannot add users to Universal Groups, only to Global and Domain Local groups.
 - B. You have Universal Group caching turned on and the cache information hasn't refreshed since this morning.
 - C. Transmission of GC data is failing across the WAN link.
 - D. GC replication doesn't support 56K links.
5. You have a network with a main office and a satellite office. The functional level of your network is Windows 2000 Native. The satellite office has a DC. The main office has a DC and a GC server. You encounter a problem with the link between the main office and the satellite office. You are concerned that users will not be able to log on at the satellite office because they cannot access the GC. To your surprise, they are still able to log on to the domain. How is this possible?
- A. The DC at the branch office could be set to cache Universal Group information, allowing clients to still log on.
 - B. The GC isn't required for logon, simply for searching the directory after you are logged on.
 - C. The DC at the satellite office is operating in the role of Schema Master and can authenticate without a GC server.
 - D. The users are logging on locally and not authenticating to the domain.
6. You have multiple locations that are part of the Default-First-Site-Name site. These locations are in Florida, Oregon, and Iowa. You have instituted GC servers at each location. While monitoring your network, you are noticing a lot of replication traffic between the locations. How can you remedy the amount of replication traffic and how that traffic is handled?
- A. Implement the use of *Subnet* objects
 - B. Implement the use of *Object* classes
 - C. Implement the use of sites
 - D. Implement the use of site connectors

Working with the Active Directory Schema

7. You are working with the Schema Admin snap-in and cannot make any changes. You created a network administrator equivalent account in the forest root domain but cannot modify the schema. Why?

- A. You must be a member of the Enterprise Admin group to modify the schema.
 - B. You must be a member of the Schema Admin group to modify the schema.
 - C. You must be a Domain Admins member in each domain in the forest to modify the schema.
 - D. Only the initial Administrator account during forest creation can modify the schema.
8. You are a network administrator and you want to modify an attribute that is associated with one of your user accounts. How do you do this?
- A. Open **Active Directory Users and Computers** and change to **advanced** view. This will allow you to modify the properties of the attributes in the user account for which you need to make the change.
 - B. Open **Active Directory Sites and Services**. Open the **Properties** for the site containing the attribute and make the modifications.
 - C. Open the **Schema** Snap-in, expand **Objects**, and select the *User* object to modify the associated attributes.
 - D. Open the **Schema** Snap-in, expand **Attributes**, and find the attribute you want to modify.
9. You are explaining the various attributes to a fellow network administrator. You are showing her the properties of a User account, and your new network administrator asks what the Other button means with regard to various attributes. What do you tell her?
- A. Those attributes are multivalued attributes.
 - B. Those attributes are single-value attributes.
 - C. Those attributes are actually *Object* classes.
 - D. Those attributes are Index attributes.
10. As a network administrator, you are responsible for making sure that various attributes are indexed for optimal performances for queries. What steps do you take to make an attribute indexed?
- A. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Index this attribute in the Active Directory**.
 - B. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Replicate this attribute to the GC**.
 - C. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Allow this attribute to be shown in advanced view**.
 - D. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Attribute is Active**.

11. You are working with *Schema* objects and you need one component that has to be supplied by a third-party. Which component is supplied by a third party so standards can be followed?
 - A. LDAP name
 - B. Common name
 - C. OID
 - D. Object GUID

12. You make a mistake while setting up new classes in your schema. You want to correct the mistake so you can have the appropriate name and configuration for the class. How do you do this?
 - A. You must deactivate the class that was added with the mistake and then rename it. You then can create a new class with the appropriate name and configuration.
 - B. You must delete the class that has the mistake and simply create the appropriate *Class* object.
 - C. You must wait 24 hours before you can delete any new classes in the schema. You can then delete the class and create the corrected *Class* object.
 - D. You can go in and fix the existing *Class* object without having to recreate the object.

13. You have an office with three locations separated by 56K WAN links. You are experiencing slow queries when looking for objects in the Active Directory. You have one GC server at your main office. What can you do to improve the query performance?
 - A. Add GC servers to your other two locations.
 - B. Add DCs that are not GC servers to your other two locations.
 - C. Add a DNS server for faster resolution at your other two locations.
 - D. Add another OU to the directory to separate the locations by OU.

14. You have been experiencing a large amount of processor utilization on your GC server. Your network consists of one location with 2500 users. You currently have three DCs for fault tolerance and load balancing. What can you do to help with your GC server processor utilization?
 - A. Add a fourth DC to the network.
 - B. Add another GC server to the network to offload some of the traffic.
 - C. Remove one DC from the network.
 - D. Split your network into three OUs with less than 1000 users each.

15. You are working on updating the schema and cannot associate an attribute with a class. What can you do to resolve this?
- A. Add yourself to the schema Admins group.
 - B. Makes sure the Schema Operations Master is online and reachable.
 - C. Reload the schema in the Schema admin tool.
 - D. Move the role of Schema Operations Master.

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|----------------|--------------|
| 1. D | 9. A |
| 2. B, C | 10. A |
| 3. A | 11. C |
| 4. B, C | 12. A |
| 5. A | 13. A |
| 6. C | 14. B |
| 7. B | 15. C |
| 8. D | |

MCSA/MCSE 70-294

Working with Group Policy in an Active Directory Environment

Exam Objectives in this Chapter:

- | | | | |
|-------|---|-------|--|
| 4 | Planning and Implementing Group Policy | 4.2 | Configure the user environment by using Group Policy. |
| 4.2.1 | Distribute software by using Group Policy. | 4.2.4 | Configure user security settings by using Group Policy. |
| 4.3.1 | Distribute software by using Group Policy. | 4.3 | Deploy a computer environment by using Group Policy. |
| 4.1 | Plan Group Policy strategy. | 4.3.3 | Configure computer security settings by using Group Policy. |
| 4.2.2 | Automatically enroll user certificates by using Group Policy. | 5 | Managing and Maintaining Group Policy |
| 4.2.3 | Redirect folders by using Group Policy. | 5.1 | Troubleshoot issues related to Group Policy application deployment. Tools might include RSoP and the gpresult command. |
| 4.3.2 | Automatically enroll computer certificates by using Group Policy. | 5.3 | Troubleshoot the application of Group Policy security settings. Tools might include RSoP and the gpresult command. |
| 4.1.2 | Plan a strategy for configuring the user environment by using Group Policy. | 4.1.1 | Plan a Group Policy strategy by using Resultant Set of Policy (RSoP) Planning mode. |
| 4.1.3 | Plan a strategy for configuring the computer environment by using Group Policy. | | |

Introduction

Group Policy is used to manage and control various features and components of the Windows Server 2003 network. Group Policy settings can be used to define users' desktop environments, to specify security settings, and to configure and control application behavior. Group Policy can be used to automatically deploy software to users and computers. You can also use group policies to assign scripts and redirect folders. Policies can be applied to a site, a domain, an organizational unit (OU) or a local computer.

Because Group Policy is used for so many important management functions, it is important for network administrators to be intimately familiar with how Group Policy works, and how they can use it for more flexibility and control of network components.

This chapter starts with the basics of Group Policy terminology and concepts, introducing you to user and computer policies and Group Policy Objects (GPOs). We discuss the scope and application order of policies, and you'll learn about Group Policy integration in Active Directory. We show you how to plan a Group Policy strategy, and then walk you through the steps of implementing Group Policy. We show you how to perform common Group Policy tasks, and discuss Group Policy propagation and replication. You'll also learn best practices for working with Group Policy, and we'll show you how to troubleshoot problems with Group Policy.

EXAM 70-294
OBJECTIVE
4
4.2.1
4.3.1

Understanding Group Policy

Group Policy is derived from the System Policies of the Windows NT days, and has been significantly enhanced, first in Windows 2000 and now again in Windows Server 2003. Implementing Group Policy in the Active Directory allows system administrators to control aspects of the user or service environment within the network from a global perspective.

You can use Group Policy to accomplish the following tasks, among others:

- **Assign scripts** You can specify scripts that will run at login, logoff, startup, shutdown, and other times.
- **Manage applications** You can designate applications that will be installed on, updated on, or removed from computers.
- **Redirect folders** You can specify alternate locations for system folders, such as My Documents, My Pictures, and others.
- **Change Registry settings** You can designate a set of Registry settings that will be applied to the local computer when a user logs on.

Gaining a full understanding of how Group Policy can impact the network requires a full understanding of the terminology and concepts.

Terminology and Concepts

You will encounter a number of terms, acronyms, and jargon when designing and implementing a group policy in your organization. Although some of the terms can be confusing at first, after you've had a chance to really work with policies, you will be able to navigate through even the most complex policy implementations.

Of course, when we refer to Group Policy, we are actually talking about the superset of all the individual components that make up the larger whole. You will find policy elements that affect only users or computers, policies that are set at the workstation level or applied to an OU in Active Directory, and ways to apply basic security to policies. Let's start with the basic terms used as the foundation of building Group Policy.

Local and Non-Local Policies

Group Policy allows you to set policies that will impact resources connecting to a specific computer or interacting with the entire directory. The terms *local policy* and *non-local policy* identify where the group policy settings originate. A local policy is stored on a specific computer (a workstation or a member server) and applies only to activities on that computer. For example, a local policy only affects a user object when the user logs on interactively on the server, either at the console or via terminal services. Local policies can also affect the way a user object accesses data from the specific server across the network. Generally, local policies should only be used on workstations; however, there are a few situations where local policies on a server would make sense.

Non-local policies are applied to group objects, primarily. These policies affect objects in the directory and are enacted when the object is active in the network. If a non-local policy affects a user object, its effect is applied every time that user object logs on, no matter what PC is used as the logon console. Group policies can apply to any of the following:

- A local computer
- An entire site
- A domain
- A specific OU

Group policies can be filtered through security settings, much like NTFS file and folder permissions control access to data on a server volume. As you will see shortly, there is a specific order in which policies are applied if local and group policies differ in a specific area, but the best practice for policies in general is to apply the policies at the group level, not at the local level.

User and Computer Policies

As you might have guessed, some policies apply to user accounts, and other policies apply to computer accounts. You can only apply policies to user and computer objects, not security

groups or other objects (however, policies can be filtered by security groups by setting the security group Access Control Entry on the GPO). These two types of policy application work as follows:

- User policies affect how user accounts interact with the network and are applied when a user logs on to the network.
- Computer policies affect how computer objects interact with the network and only apply to those computers that participate in the Active Directory.

You configure each of these types of policies in separate areas in the GPO Editor.

User and computer policies are divided into three groups: Software Settings, Windows Settings, and Administrative Templates.



NOTE

See the section titled *Implementing Group Policy* later in this chapter for instructions on opening and using the GPO Editor.

Software Settings

The primary use of this setting is to install, update, or remove software on computers on the network. The Software Installation node is located in this group, and other policy groups can be added in this area by other applications.



NOTE

The Software Installation node does not appear in local GPOs, as automated deployment of software through group policy can only be applied at the site, domain, or OU level, not at the local level.

Software policies set in this area under Computer Configuration apply to all users who log on to the computer where the policy applies. This policy setting could be used to designate a specific computer on the network where a particular application should be installed, no matter who logs on to the computer. Software policies set in this area under User Configuration apply to all computers that a particular user logs on to. This setting is useful if a particular user has a specific application that he or she needs to use, no matter where that user uses a computer in the organization. The policies can be set so that if an application is installed on a computer this way, only the user to whom the policy is applied is able to see or run the application.

Windows Settings

Policies applying to scripts, security, folder redirection, and Remote Installation Services, among others, are located in this area. There are significant differences between these policy settings depending on whether they are applied in the Computer Configuration or User Configuration node. Table 9.1 details some of the policy groups and whether they are applied to user or computer settings.

Table 9.1 Group Policies for Windows Settings

Policy	Location	Description
Scripts	Computer Configuration	Specifies startup and shutdown scripts to be run on the computer.
Scripts	User Configuration	Specifies logon and logoff scripts to be run by users.
Account policies	Computer Configuration\ Security Settings	Contains policies related to password and account lockout settings.
Folder redirection	User Configuration\ Security Settings	Contains policies to redirect certain user folders, such as Application Data, My Documents, and Start Menu, to alternate locations.
Internet Explorer maintenance	User Configuration\ Security Settings	Contains settings to modify defaults for Internet Explorer, such as user interface settings, favorites, connection settings, and security zone settings.
Public Key policies	Computer Configuration\ Security Settings	Contains policies related to system-level public key activities, such as Encrypted File System, Enterprise Trust, Autoenrollment settings, and Automatic Certificate Request settings.
Public Key policies	User Configuration\ Security Settings	Contains policies related to user-level public key activities, such as Enterprise Trust and Autoenrollment settings.

Administrative Templates

Policy settings that appear in the Administrative Templates node of the GPO Editor contain Registry settings to achieve each of the settings contained in the hierarchy. Policies for user configuration are placed in the HKEY_CURRENT_USER (HKCU) area of the Registry,

while those for computer configurations are placed in the HKEY_LOCAL_MACHINE (HKLM) area.

Administrative templates contain settings for Windows components such as NetMeeting, Internet Explorer, Terminal Services, Windows Media Player, and Windows update, to name a few. Other components common to both user and computer configurations include settings for user profiles, script execution, and group policy.

While the different policy settings between user and computer configurations are too numerous to list here, there are some key components available for the user configuration. These include the Start Menu, Taskbar, Desktop, Control Panel, and Shared folder settings.

Group Policy Objects

All group policy information is stored in Active Directory in GPOs. You can apply these objects at the site, domain, or OU level within the directory. Since the GPO is an object in the directory, you can set security permissions on the objects to determine who will access the policy settings stored in the GPO.



NOTE

It is important to differentiate between local GPOs, which are stored on the local computer, and non-local GPOs, which are stored on the DC(s) in Active Directory.

Because GPOs can impact a large portion of the directory, you should update GPOs infrequently. Each GPO update must propagate across the entire directory to take effect, and this could be a time-consuming process if the directory structure is very large. You should also restrict the number of individuals who make changes to GPOs that can impact the entire organization. Otherwise, you can run into the situation where two administrators make contradictory changes to a GPO in different locations of the tree, and the changes propagate differently around the tree, potentially causing problems until the directory has completely updated the GPO changes.

Scope and Application Order of Policies

A single object in the network can be subject to multiple policy settings, depending on how Group Policy is configured on the local machine and in the directory. Active Directory processes policy settings in a specific manner when an object connects to the network. Knowing this process will help you troubleshoot problems with policy settings as they arise.

Local, Site, Domain, OU

Group Policy settings are applied in the following order:

1. **Local settings** Each computer has its own local GPO, and these settings are applied before any others. There is only one local GPO per computer.

2. **Site settings** Group policies associated with the site in Active Directory are processed next. The system administrator can set a specific order in which the site policies are to be applied, if more than one policy is defined.
3. **Domain settings** Group policies associated with a domain object follow the completion of the site settings. If multiple domains are involved, the administrator can set the order of preference in which those settings will be applied.
4. **OU settings** Group policies associated with an OU are applied last in the processing order, but the processing starts with the OU highest in the directory structure. The remaining OU GPOs will be processed in descending order until the OU that contains the directory object is reached. If multiple policy settings are applied for a particular OU, the administrator can set the order in which the settings are applied.


Figure 9.1 details the order in which multiple policies are applied when a user object logs on to the domain. In the diagram, the user object exists in the OU 4 OU, which is in the OU 3 OU of Domain 1 of Site. When the user logs on, the local policy of the computer is applied, followed by any GPOs attached to Site, then Domain 1, then OU 3, and finally OU 4.



NOTE

User policies are applied at logon; computer policies are applied at bootup.

Figure 9.1 Processing Policy Settings at User Logon



Understanding Policy Inheritance

We saw in Figure 9.1 that when the user logged on, policies from the Site, Domain, and OUs were applied to the user object. The example indicated that any policies associated with OU 3 would be applied before the policies in OU 4. Through policy inheritance, the policies in OU 3 will apply to all objects in OU 3, OU 4, OU 5, and OU 6, even if no specific policies are assigned to OU4, OU5, or OU6.

Objects in child containers generally inherit policies from the parent containers within a domain. If a policy setting is enabled in OU 3 and that same policy setting is not configured in OU 4, then objects in OU 4 inherit the policy setting from OU 3. If a policy setting is disabled in OU 3 but that same policy setting is enabled in OU 4, then the policy setting is enabled in OU 4, as the GPO for OU 4 overrides policy settings from OU 3. This is the way it works by default.

However, administrators can block inheritance on group policy settings at the OU level. If you want to start with a clean slate at a particular OU, you can use the Block Policy Inheritance setting at that OU, and only the settings in the GPO for that OU will apply to objects in the OU. Blocking policy inheritance does not impact local computer policy settings, only Active Directory group policy settings.

In addition, policies set at a higher container can be marked as No Override, which prevents any lower container settings from changing the policy settings of the higher container. Going back to Figure 9.1, if the GPO for OU 3 is marked for No Override, and a policy setting in the GPO for OU 4 conflicts with a setting from OU 3, the setting in OU 4 will not take effect. You cannot block a policy that is set to No Override.

You should use great care in using the Block Policy Inheritance and No Override settings when configuring Group Policy. Changing the default way in which policy is applied can complicate troubleshooting of policy settings if problems are encountered.



EXAM WARNING

Be sure you have a complete understanding of how Group Policy is applied before taking the exam. You will need to be able to determine how and when policies are applied based on policy scope, order of processing, security settings, and implications of the No Override and Block Policy Inheritance settings. If you can develop a policy map like that shown in Figure 9.1, you should be able to correctly answer any questions about policy settings based on these factors.

Filtering Scope by Security Group Membership

As mentioned, you can further control which policies are applied to which objects by filtering policy application by security group membership. Similar to setting permissions on files and folders with NTFS security settings, you can set security on a GPO so that only certain groups can see the GPO, which means that only those groups will have the policies applied.

Looking back at Figure 9.1, the diagram assumes that there is no security filter on the GPOs at any level. Now let's suppose that the user object is a member of the Accounting group, and that the GPO in OU 4 has security permissions set. If the security permissions on the GPO in OU 4 do not give members of the Accounting group access to read the GPO, then the user will not have the GPO settings for OU 4 applied when he or she logs on.

If you find yourself needing to filter GPO settings based on group membership, you might need to set multiple GPOs on a container and adjust the security settings accordingly. Again, adding a number of GPOs to a container increases the complexity of the policy setting process, which can cause complications for troubleshooting.

How Much Is Too Much?

A word of caution about group policy: too much of a good thing can be a bad thing. Yes, you can use group policy to significantly detail the operations of your network environment, but it will come at a cost. Each additional GPO that applies to a user at logon increases the time needed to authenticate to the directory. If there are site or domain GPOs across slow network links, logon time will increase even more.

A key factor in minimizing the amount of time needed to process GPO settings at logon is minimizing the number of policies that are configured. In other words, avoid setting a policy at one level in the hierarchy and retracting that setting in a lower level. If not every object needs a policy set, only set the policy for a specific group or OU.

Keeping the number of GPOs to a minimum will also aid in troubleshooting policy problems. The greater the number of GPOs applied, the greater the chance of a misconfiguration, and the more places you will have to investigate to find the source of the conflict.

Group Policy Integration in Active Directory

As mentioned earlier, non-local group policy settings are stored in objects in the Active Directory. These objects are linked to specific containers: sites, domains, and OUs. Since GPOs are objects in the directory, they are subject to all the settings and rules of other objects.

Group Policy Propagation and Replication

Active Directory replication has an impact on group policy application in a large directory structure. Because GPOs are objects in the directory, they must be replicated to all copies of the directory partition on all domain controllers (DCs) before the settings can take effect in all circumstances. Replication is a concern for GPOs linked to a site or domain with multiple controllers.

When group policy is set for a domain, by default the actual object is tied to the server that has the primary domain controller (PDC) Emulator operations master token. The other DCs will receive the updated policy information as the token is passed around through replication. Users who authenticate to DCs other than the PDC might not receive the updated policies upon logon if the directory has not had ample time to replicate the settings.

You can specify a particular DC to be used for editing group policy by using the **DC Options** command in the **View** menu of the GPO Editor. As mentioned, the default is the DC with the PDC Emulator operations master token, but you can change this setting.

Sites that have multiple servers connected over slow WAN links have several issues related to policy propagation and replication. Obviously, a DC with an updated group policy is impacted by a slow WAN link when attempting to replicate the data across the link. Depending on how the directory is configured, DCs across the slow link can be set up to replicate much less frequently than those on a faster link.

Also of concern are users who authenticate to a DC across a slow WAN link. While the normal authentication process might not be all that network-intensive, more GPOs that have to be processed by the user significantly increases the time needed for full authentication.



Planning a Group Policy Strategy

You must consider a number of factors when planning the group policy strategy for your organization. Some of these factors include size of the organization, geography of the organization, structure of the organization, and so on. More importantly, you must determine the effective policy settings you want to have for each object in the directory.

One way to test your policy plan is to create the policies and then log on with user accounts from different locations of the directory and see how the policies impact the user experience. This is time consuming, cumbersome, and has a definite impact on the production network. Fortunately, Microsoft provides a way for evaluating the proposed policy environment without impacting the production system.

Using RSoP Planning Mode

The Resultant Set of Policy (RSoP) tool, included with Windows Server 2003, has a special planning mode that system administrators can use to evaluate the design of the group policy within the directory. The planning mode of RSoP can simulate a number of situations where group policy settings can be affected by a number of factors, including slow network links.

Opening RSoP in Planning Mode

To use RSoP in planning mode, you will need to run the **Resultant Set of Policy Wizard** from inside the Microsoft Management Console (MMC). You can follow these steps to open RSoP in planning mode to collect information for an RSoP report.

1. Open **Microsoft Management Console (MMC)** and add the **RSoP snap-in**.

1. Select **File | Add/Remove Snap-in**.
 2. Click **Add**.
 3. Select **Resultant Set of Policy** from the list.
 4. Click **Add**, and then click **Close**.
 5. Click **OK**.
2. Right-click on **Resultant Set of Policy** and select **Generate RSoP Data**.
 3. Click **Next** in the Resultant Set of Policy Wizard window.
 4. Click the **Planning Mode** option button, and click **Next**.

The RSoP wizard will walk you through the steps of gathering the data that can be collected and included in the RSoP report. On each page, there is a **Skip to the final page of this wizard without collecting additional data** check box. If you select the check box, only the data specified up to that point in the wizard will be included in the RSoP query. All other settings will take their default values.

The first page of the wizard collects user and computer information on which the query will run. This is the only data that is required in the wizard, as all subsequent pages can be skipped by clicking the **Skip to the final page of this wizard without collecting additional data** check box. Figure 9.2 shows this wizard page. On this page, you must select a specific user or user container, and a specific computer or computer container. You can use the **Browse** buttons to search for a user or computer or the parent container, or you can enter the information directly into the fields. After the information for the user and computer selections is complete, the **Next** button will enable and you can move to the next page of the wizard.

Figure 9.2 Entering User and Computer Information in the Resultant Set of Policy Wizard

The next page of the wizard allows you to specify any advanced simulation options. On this page, you can specify the report to simulate a slow network connection and loopback processing options, if any. You can also specify which site's policies to test, if there are multiple sites available.

If you specified a specific user or computer in the initial page of the wizard, the next page of the wizard will allow you to specify an alternate location for the object or objects specified. Changing the location of the object will let you test what changes would occur if you moved the object to a different location in the directory. If you only select containers in the initial page, this page will not display.

The next page of the wizard identifies the security groups for the user object selected. If a specific user is selected in the first page of the wizard, the security groups for that user are displayed. If a user container is specified, the Authenticated Users and Everyone groups are listed as defaults. You can add user groups to the list or remove groups from the list to see what changes will occur as a result.

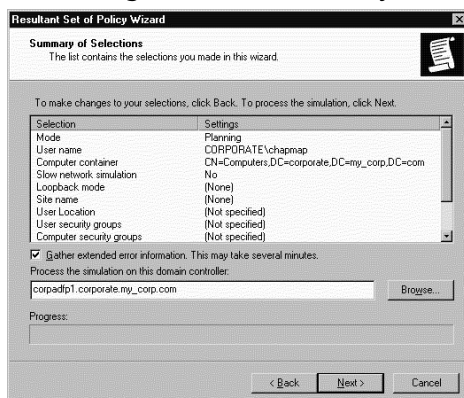
The next page of the wizard identifies the security groups for the computer object selected. As with the user selection in the previous page, you can specify which security groups to use when running the query.

The next options page of the wizard allows you to select the Windows Management Instrumentation (WMI) filters to use on the user object or container in the query. The default selection is for all linked WMI filters, or you can select only specific WMI filters.

The last options page of the wizard selects the WMI filters for the computer object or container. As on the previous page, you can accept the default selection of all WMI filters, or you can specify which filters to use.

After you complete all the pages of the wizard, or if you select the option to skip the remaining information pages, a summary page will display the options that will be used when running the query. Figure 9.3 shows the summary page and the information specified for a sample query. In this window, you can choose to gather extended error information or select a different DC to process the simulation. Clicking **Next** will start the query based on the information listed in this page.

Figure 9.3 Reviewing the Settings of the RSOP Query Prior to Execution

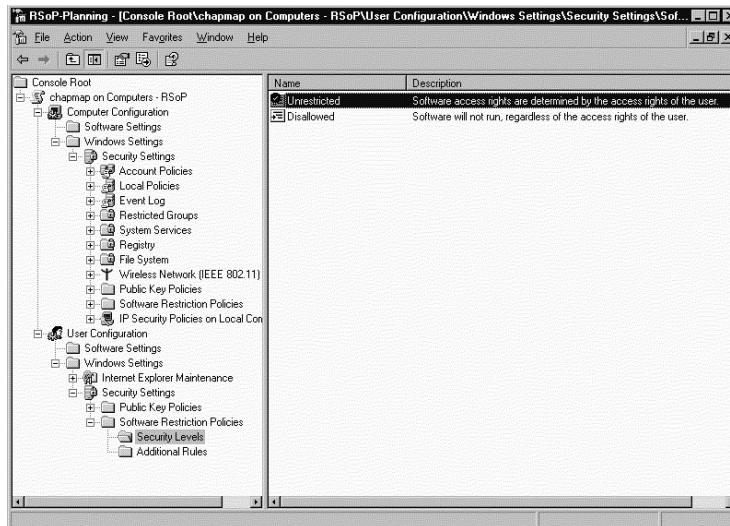


When the query has completed (which might take several minutes depending on the size and configuration of your environment), the wizard's finish page will display. Clicking **Finish** will close the wizard and return you to the MMC to review the RSoP report.

Reviewing RSoP Results

The results of the RSoP query displayed in the MMC will look similar to the Group Policy Object Editor window, with a few important differences. Figure 9.4 shows the RSoP results window in the MMC. This particular query was run on the user *chapmap* and the *Computers* container. When looking at the policy settings in the window, you only see the policies that will be in effect for the user when logged on to a computer in that particular container. You will also only be able to view the policy settings in this interface. You will not be able to change any policy.

Figure 9.4 Reviewing the RSoP Planning Results

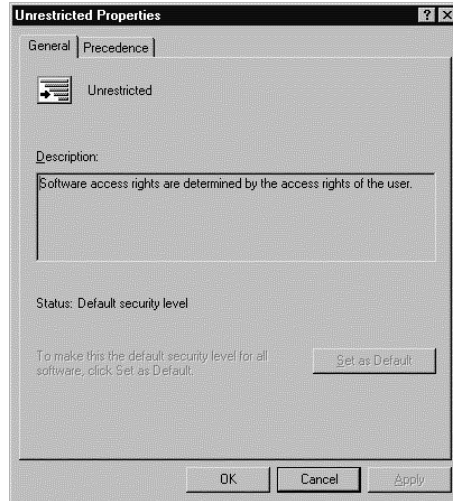


When you right-click either the **Computer Configuration** or **User Configuration** node in the tree and select **Properties**, you will find information about the policies that were processed to generate the results found in the report. You can select to view all GPOs and their filtering status to see which GPOs were processed and which were not, and why not if they were not. You can display revision information to see how many times a particular GPO has been modified, and you can display scope information that tells where the GPO resides. If you click the **Security** button, you can see the security permissions set for the GPO.

If you open a policy setting, you can view the properties of that policy setting. Figure 9.5 shows the properties of the setting selected in Figure 9.4. As shown in the figure, the option to set this particular setting as default is grayed out, because no changes can be made

in this interface. If you click on the **Precedence** tab, you will see a list of GPOs where this particular policy is set, including the order in which this policy was processed.

Figure 9.5 Viewing the Properties of a Policy Setting in the RSoP Report



You can run an additional query on a different set of user and computer objects from this interface by right-clicking on the RSoP result object in the left pane, **chapmap on Computers – RSoP** in this instance, and selecting **Change Query**. If you go in and make group policy changes that would impact the results of the query and want to see how those changes actually affect the system, you can right-click the RSoP result object and select **Refresh Query**. This second option will re-run the query with the same options

Strategy for Configuring the User Environment

When setting group policy at the user level, you are creating an environment that will follow the user around the network. No matter what computer the user logs on to, the group policy settings inherited by that user will apply. This section covers some of the “shoulds” and “should nots” related to the user environment.

One policy setting that will follow the user around no matter where he or she logs on is *roaming profiles*. Enabling roaming profiles for a user community will store all the user settings on a server rather than on the local computers. When a user logs on, all of his or her profile settings (Desktop items, My Documents, Registry settings, etc.) will be pulled off the server, ensuring that the user has the same environment on each computer he or she uses. This approach has many advantages, but it has disadvantages as well. Some profile settings are hardware-dependant, and if the computers used by the user do not have the same hardware, the user could encounter difficulty upon logon (video cards can be especially problematic in this regard).

Software Installation is another policy that can be of great benefit to the organization. If a certain group of users has a particular application that is critical to performing their work tasks, you can set up software installation policies that will download and install the application on any computer the user uses throughout the company. This policy also keeps unauthorized users from being able to run the application even though it is installed on the computer they are using. The same caveat applies to software installation as to roaming profiles. Not all software packages are compatible with other programs that might be installed on a computer. Before implementing this type of policy in the organization, you'll need to make sure that the applications being installed will work well with other programs that already exist on the computer. The last thing you want to do is to break one program on a system by installing another.

The vast majority of other group policy settings that you can apply to users in the directory have little chance of causing conflict with other settings on the local computer. Logon and logoff scripts, application settings, folder redirection, and environment configurations can help to standardize the user's computing experience across multiple machines, which can, in turn, ease the support burden on your IT staff.

Speed Is Key

One of the most significant factors to take into consideration when configuring group policy for the user environment is communication speed. In a small office where all systems are connected to the network at a fairly high speed, you might not encounter many problems with the group policy settings. If your organization is large or geographically diverse, there are additional items you should consider.

For example, consider a company that spreads across the entire North American continent with some additional sites in Europe. The majority of its facilities are manufacturing sites with a small number of computers and relatively slow WAN links connecting them. For the type of computing activities that take place at these facilities, the communication speeds are sufficient.

If this organization were to implement roaming profiles or software distribution through group policy, its IT personnel would encounter major difficulties. Just the number of policies applied at logon could significantly affect the logon time for the users. The company regularly has contacts from its corporate office, where roaming profiles and software distribution would be beneficial, visit these sites and access computing resources while there. As soon as one of these site visitors tried to log on and download his or her roaming profile or even a single software application, the lines at the Help Desk would start ringing.

Fortunately, there are ways to limit certain policies from taking effect across slow network links, but this must be taken into consideration early in the design process. It's generally not safe to assume that just because a certain user group has never performed a particular activity, such as traveling to a remote site and attempting to work across a slow WAN link, that the members of that group never will.

Strategy for Configuring the Computer Environment

When setting policy for the computer environment, the settings applied will impact every user who logs on to the computer. Unlike user settings, there are two places where computer policy is applied. The first is the local policy set at the computer each time it boots. These settings are applied first, and any subsequent policy that conflicts with the local settings will override the local settings. However, computer policy can also be set in Active Directory. These settings follow the same rules as user settings in terms of priority order. Any computer policies set at the site level will be overwritten by additional policy settings at the domain or OU level when the settings conflict.

One case where computer policy overrides user policy is when a GPO containing computer settings is configured to operate in *loopback mode*. Loopback mode is a special setting that is only used in cases where a very specific set of policies needs to be applied in a controlled environment. Loopback mode allows administrators to apply group policy based on the computer at which the user is logging on. In other words, this setting is used if a particular user should have different policies applied, depending on where he or she logs on. When loopback processing is enabled, the computer policies set for the system override any user policy settings applied during logon.

Loopback operates in two modes—*replace* and *merge*. When loopback is enabled, one of these two modes must be selected. Replace mode will eradicate any user policy settings applied at logon and only retain the computer policy settings. Merge mode will allow user settings that do not conflict with computer settings to be applied. If there is a conflict between the two, the computer settings override the user settings.



EXAM WARNING

You might encounter exam questions about loopback mode. Make sure you understand why you would want to use loopback mode, as well as the difference between the replace and merge settings.

The philosophy of “less is more” applies directly to the approach for setting computer policy in the domain. In general, you should try to have only one set of policies apply to computers. If you do have cases where you need different policy settings to apply to different sets of computers within the organization, set up the separate policy objects, but restrict access to those objects so that only the systems that need to be affected by the object will process the settings.

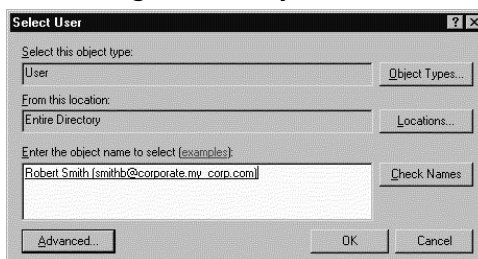
EXERCISE 9.01

RUNNING AN RSoP PLANNING QUERY

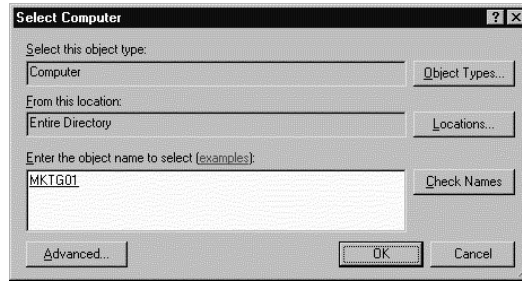
This exercise walks you through the process of generating an RSoP planning report based on changing a user object from one OU to another. For this exercise, we will use the user object for Robert Smith, which exists in the Marketing OU, and build an RSoP report showing the policy settings that would apply to the object if it were moved into the Accounting OU. As long as you have appropriate permissions to run an RSoP query on a system, you should be able to emulate the steps in this exercise on a system to which you have access, as you will not be changing any settings on the system in the process.

1. Open an MMC window and load the RSoP snap-in (see the steps outlined earlier in this section if needed).
2. Right-click the **Resultant Set of Policy** object in the console tree, and select **Generate RSoP Data**.
3. In the Resultant Set of Policy Wizard, click **Next**.
4. Select the **Planning Mode** option button, and click **Next**.
5. In the **User and Computer Selection** window, select the **User** option button, and click the **Browse** button in the **User information** frame.
6. In the **Select User** dialog box, choose a user object, as shown in Figure 9.6, and click **OK**.

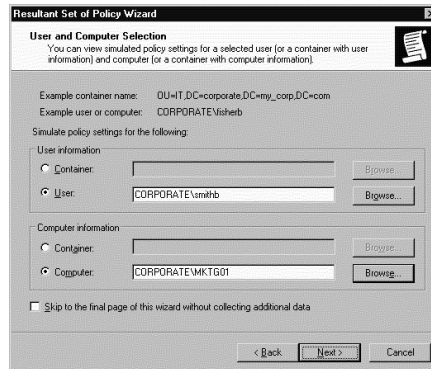
Figure 9.6 Selecting a User Object



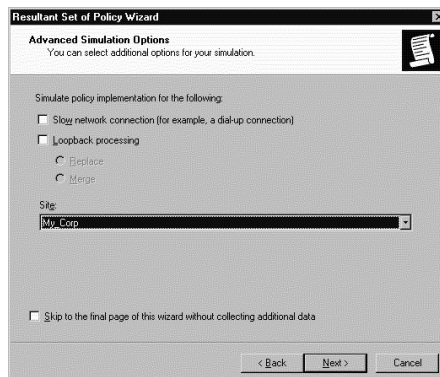
7. In the **User and Computer Selection** window, click the **Computer** option button, and click the **Browse** button in the **Computer information** frame.
8. In the **Select Computer** dialog box, choose a computer object, as shown in Figure 9.7, and click **OK**.

Figure 9.7 Selecting a Computer Object

9. The **User and Computer Selection** window should now appear as in Figure 9.8. Click **Next**.

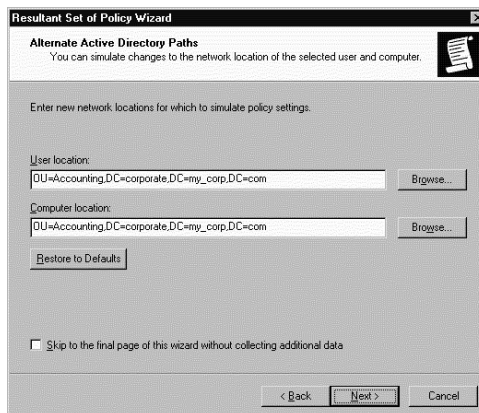
Figure 9.8 Specifying the User and Computer Objects in the RSOP Wizard

10. In the **Advanced Simulation Options** page, select the appropriate site from the **Site** drop-down list, shown in Figure 9.9, and click **Next**.

Figure 9.9 Specifying a Site in the Simulation Options Page

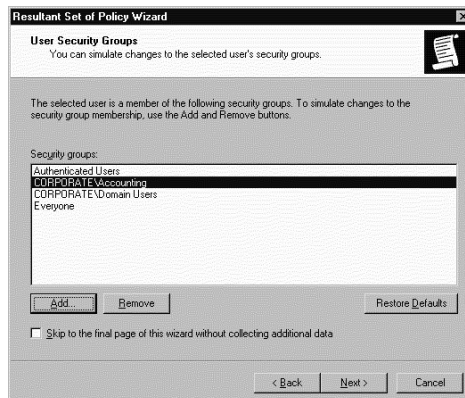
11. In the **Alternate Active Directory Paths** page, change the location for the user and computer objects. In Figure 9.10, the Accounting OU is selected for both the user and the computer, as the user account had been in the Marketing OU and the computer had been in the Computers container. When the new locations have been selected, click **Next**.

Figure 9.10 Setting the New User and Computer Locations



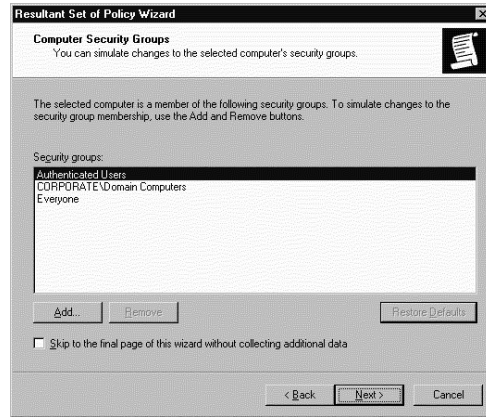
12. In the **User Security Groups** page, change the security groups to match those of the new location. Select the groups that the user would no longer belong to, and click **Remove**.
13. To add new security groups to the query, click the **Add** button and select the appropriate groups. When finished, the page should be similar to Figure 9.11. In this example, Robert had belonged to the Marketing security group, which was removed. He has been added to the Accounting security group, as shown in the figure. Click **Next**.

Figure 9.11 Selecting User Security Groups



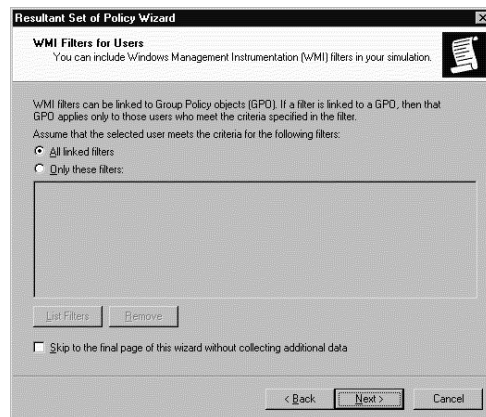
14. In the **Computer Security Groups** page, you can leave the security group setting as it is, or you can change group assignments by using the **Add** and **Remove** buttons. Figure 9.12 shows the settings used for this query. When complete, click **Next**.

Figure 9.12 Selecting Computer Security Group Settings



15. In the **WMI Filters for Users** page, select the **All linked filters** option button to include all WMI filters in the query, as shown in Figure 9.13, or select the **Only these filters** option button to specify which filters to use. When finished, click **Next**.

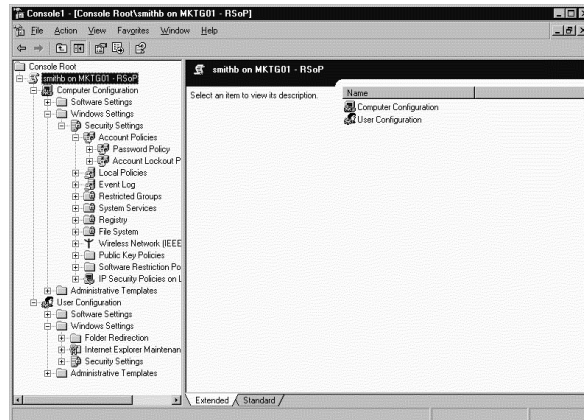
Figure 9.13 Selecting WMI Filters for Users



16. In the **WMI Filters for Computers** page, select the **All linked filters** option button to include all WMI filters in the query, or select the **Only these filters** option button to specify which filters to use. When finished, click **Next**.

17. Review the selections made in the Summary of Selections page, and click **Next** to start the query.
18. When the query has completed, click the **Finish** button to close the wizard and view the RSoP report, shown in Figure 9.14.

Figure 9.14 Viewing the RSoP Report



19. Browse through the report looking at the policies that would be enabled for user smithb on computer MKTG01. Close the MMC when done.

EXAM
70-294
OBJECTIVE
4
4.2.1
4.3.1

Implementing Group Policy

Now that you know how to evaluate the effects of group policy on the directory, it is time to start creating policy objects and applying policy to the environment. In this section, you will learn about the different places where you can create GPOs, and the tools to modify and manage them.

The Group Policy Object Editor MMC

The **Group Policy Object Editor** is a snap-in for the MMC. Because group policy can be applied at several locations, opening the Group Policy Object Editor can differ depending on where you want to apply group policy.

From within an MMC, you can select the **Group Policy Object Editor** snap-in from the **Add/Remove Snap-in** window. When selecting the Group Policy Object Editor from the list of stand-alone snap-ins, the Group Policy Wizard will open, allowing you to select the scope of the group policy to work with. Clicking the **Browse** button in this wizard will open the **Browse for a Group Policy Object** window, shown in Figure 9.15. The

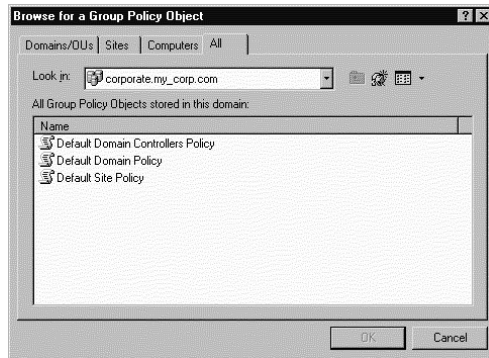
first three tabs in the window allow you to search for GPOs of a specific type: Domain/OU, Site, and Computer. The fourth tab, selected in Figure 9.15, displays a list of all policy objects in the domain, regardless of location. Local computer policy objects will not show in this listing, because they are stored on the computer, not in the domain.



NOTE

The local computer policy is the default in the Group Policy Wizard. To access it, simply click **Finish** on the first page of the wizard, instead of clicking the **Browse** button.

Figure 9.15 Viewing all Group Policy Objects in the Domain



To edit one of the existing GPOs stored in Active Directory, select the GPO from one of the tabs and click **OK**. This will take you back to the Group Policy Wizard. When you click **Finish** in the wizard, the Group Policy Object Editor will open in the MMC, and you can begin editing the GPO.

Creating, Configuring, and Managing GPOs

Loading the Group Policy Object Editor snap-in in an MMC will allow you to edit existing policies in the network. When the domain is first created, there are three default policies created:

- Default Site Policy
- Default Domain Policy
- Default Domain Controllers Policy

You will probably want to create new policies and associate them with specific areas of the directory.


Creating and Configuring GPOs

There are two ways to create new GPOs in the directory. You already know how to load the Group Policy Object Editor snap-in into the MMC, so let's look at how to create a new GPO from the Group Policy Wizard.

In Figure 9.15, you saw the **Browse for a Group Policy Object** window that opens when you click the **Browse** button in the Group Policy Wizard. Next to the **Look in** drop-down menu, you will find the **Create New Group Policy Object** button. When you click this button, a new GPO will be created in the scope you have selected in the **Look in** menu. Creating the GPO in this scope will automatically link the object to the container that was selected in the scope.

Another way to open the GPO Editor and create a new GPO is from within the **Active Directory Sites and Services** or **Active Directory Users and Groups** tools. Right-click the object in the container list where you want the GPO to be created, and select **Properties**. Then, select the **Group Policy** tab in the **Properties** window to see what policies are already linked to the container or to create a new object for the container. Figure 9.16 shows the **Group Policy** tab for the IT Management container. In this example, there is only one object tied to this container. To create and edit a new GPO, click the **New** button, give the policy a name, and then click **Edit** to open the Group Policy Object Editor for the new GPO.

Figure 9.16 Viewing the Group Policy Objects for a Container



Naming GPOs

All GPOs created in the directory should have unique names. Even though each GPO is associated with a specific container and could have the same name as another object in the tree, there will be much less confusion when troubleshooting if each GPO name is unique. GPO names can contain letters, numbers, and special characters, but the name cannot be

longer than 255 characters. Any GPO name longer than 255 characters will be automatically truncated to the 255-character maximum.

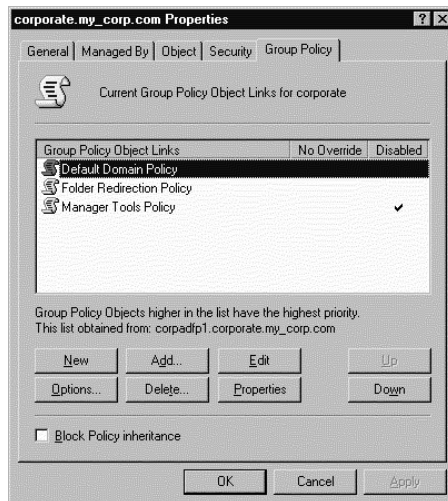
There are no other specific rules as to how to name each GPO. In the same way that you should name each object in the directory to match its function or purpose, you can consider the same approach when naming GPOs. If you have a set of policies that will impact a single container in the directory, such as an OU, you could include the name of the OU in the name of the GPO. If the policies contained in a GPO are going to be linked to a number of containers in the directory, you could name the GPO after the function its policies are designed to perform.

Managing GPOs

From the Group Policy tab of the container Properties window, you can perform a number of functions on the GPOs associated with the container. We have already covered creating and editing a new GPO from the interface. Now let's take a look at some of the other ways you can manage the GPOs from this interface.


Figure 9.17 shows the Group Policy tab of the root of the domain of My Corp. There are three GPOs stored within this container in Active Directory: Default Domain Policy, Folder Redirection Policy, and Manager Tools Policy. Based on the information displayed in the figure, the Default Domain Policy and the Folder Redirection Policy objects will be processed by objects logging on within this domain. The Manager Tools Policy will not be processed with the other two GPOs at this level because it has been disabled at this level, as indicated by the check mark under the Disabled column next to the policy object. We can also see that none of the GPOs have been marked as No Override.

Figure 9.17 Managing the Group Policies for the Root Domain



The No Override and Disabled settings can be set in two places. Clicking the **Options** button will open an Options window for the selected GPO. The Options window, shown in Figure 9.18, allows you to set the No Override and Disabled settings for the object. In addition, you can set both options by right-clicking the GPO in the list and selecting either the No Override or Disabled entries in the pop-up menu.

Figure 9.18 Reviewing the Options for a GPO



Clicking the **Properties** button in this interface will bring up the Properties window for the object. Within the GPO properties, you can modify a number of settings that will control who accesses the policy and how it is applied. These properties are covered in detail in the next section.

You can also click the **Delete** button to remove a policy from this container. When clicking Delete, you will be asked if you want to remove the link to the GPO from the container or if you want to permanently remove the GPO from the directory altogether. If the policy is linked to multiple containers and you only want to remove the link from the current container, select the **Remove the link from the list** option button, as shown in Figure 9.19. Otherwise, click the **Remove the link and delete the Group Policy Object permanently** button to completely eliminate the GPO from the directory.

Figure 9.19 Removing a GPO Link or Object



If there are multiple GPOs linked to a container, as there are in Figure 9.17, you can specify the order in which the GPOs are processed within the container. When multiple GPOs are present in the list, use the **Up** and **Down** buttons to arrange the order of the GPOs in the list.

Finally, you can block policy inheritance for the container by enabling the **Block Policy inheritance** check box. If the container is a child object in the directory, turning

on this option will prevent the container from inheriting any policy settings from parent containers. The only time that the Block Policy inheritance setting can't prevent settings from inheriting is if a parent container has a policy with the No Override option set. It should also be noted that the Block Policy inheritance setting applies only to the container and not the specific GPOs associated with the container.

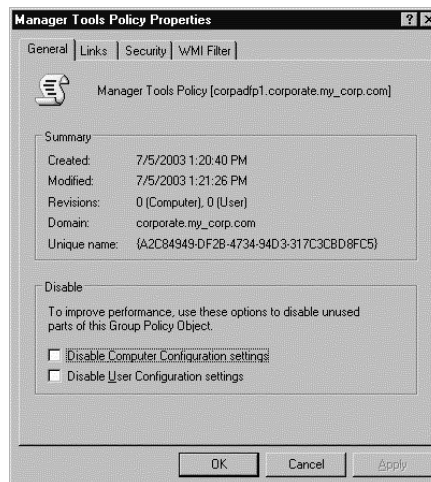
Configuring Application of Group Policy

Placing a GPO in a container enables the policy settings within the object on all objects that log on as part of that particular container. There are times when you will not want all objects associated with a container to have the policy settings applied, either for security or performance reasons. This section deals with ways of governing access to the settings within a GPO from within the GPO properties interface.

General

Figure 9.20 shows the General tab of the GPO Properties window. This is the view that is opened by default when the Properties window is opened. This view provides system information about the GPO and allows you to exclude certain portions of the policy from application. If the policy object only contains user configuration policies, you can check the **Disable Computer Configuration settings** check box, and the Computer Configuration settings will be ignored when the object is processed. This will help to cut down on processing time at bootup if there are no policies specifically set in the object for computer settings. Alternately, you can check the **Disable User Configuration settings** check box to prevent the user configuration settings in the GPO from processing at logon.

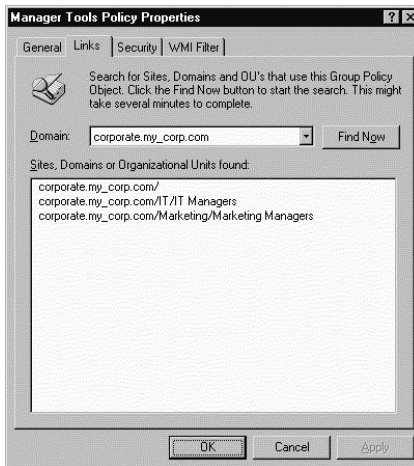
Figure 9.20 Viewing the General Properties of a GPO



Links

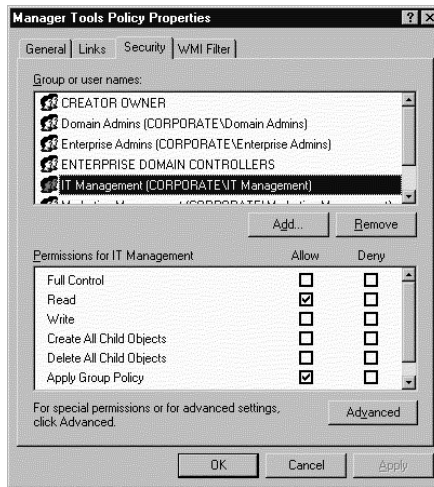
Figure 9.21 shows the Links tab of the GPO Properties window. In this view, you can search for all the places where the specific GPO has been linked. By default, this window is empty when first opened. Select the appropriate domain from the drop-down list and click **Find Now** to search for all the containers where the GPO is linked. In Figure 9.21, this step has already been done. This particular GPO is linked to three containers. While you cannot change any of the settings for the GPO in this view, you can find all the places where the GPO is enabled when troubleshooting policy problems.

Figure 9.21 Viewing the Links of a GPO



Security

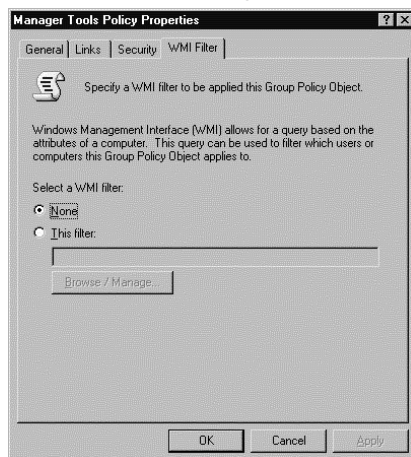
Figure 9.22 shows the Security tab of the GPO Properties window. In this view, you can set all the security permissions necessary to govern how the policy will be applied and managed. In this example, the Authenticated Users group is not listed. When a new GPO is created, the Authenticated Users group is given Read and Apply Group Policy permissions on the GPO. Those two permissions are the minimum needed to be able to have policy settings applied to a group.

Figure 9.22 Viewing the Security Settings on a GPO

In this case, the IT Management security group has been given the Read and Apply Group Policy permissions to the object, so members of the IT Management security group can see and process the policy settings contained within this object. However, these permissions alone do not allow members of this security group to process the policy object. Members of this group have to log on within the context of a container that was linked to this GPO. In Figure 9.21, we saw that the object was linked to the root domain container, the IT Managers container, and the Marketing Managers container. That means everyone in the IT Managers security group will process the GPO, right? Wrong. Only members of the IT Managers security group who are located in the IT Managers container process the GPO, because in Figure 9.17, we saw that the Manager Tools Policy object was disabled in the root container of the domain.

WMI Filter

Figure 9.23 shows the WMI Filters tab of the GPO Properties window. In this view, you can set WMI filters to further restrict who does and does not have access to the GPO for processing. You can use WMI queries to further filter application of a GPO beyond what you can achieve with security settings. WMI filters are written in the WMI Query language (WQL) and are generally used for exception processing.

Figure 9.23 Viewing the WMI Filter Settings on a GPO

Placement of Group Policy Objects

There are a number of different ways to choose how and where to locate GPOs in Active Directory, but one method we have found to be very useful is to place the GPO in the same location as the container to which it will be linked. If we have a particular OU for which we need to create policy settings, we create the GPO at that container level in the directory and link the GPO to the container. We also frequently include at least a portion of the container name in the GPO name so we can track down the location of the GPO quickly.

When working with a GPO that will be linked to multiple containers, we generally store that GPO at the root of the domain where the policies will be applied. We make sure that the GPO is not linked to the domain directly, but only to the containers that need the settings. In this case, we generally name the container based on the specific tasks or functions that the GPO settings affect.

The philosophy behind this practice is ease of locating GPOs for editing. By keeping the GPOs at the same level as the linked container or at the root of the domain, we do not spend time trying to hunt down the particular GPO when changes or a review of settings are needed.

Whatever scheme you choose for the placement of your GPOs in the directory, you will help yourself and any other support personnel by being consistent. Should you choose to store all your GPOs in a single location in the directory for one domain, do the same for all domains. In addition, document where and how you created the GPO structure. You might thank yourself later for doing so, and there's less of a chance that people coming into your role after you will have problems locating and troubleshooting policies.

Delegating Administrative Control

You might be in an Active Directory environment where one group or organization only controls a small portion of the directory. Because Active Directory allows you to delegate control of parts of the directory tree, you might find yourself needing to delegate control over Group Policy as well. By default, only the following are allowed to create and manage GPOs in the directory:

- Domain administrators
- Enterprise administrators
- Members of the Group Policy Creator Owners group.

Granting the ability to create and manage GPOs to a non-administrator user takes two steps:

1. The user must become a member of the Group Policy Creator Owners group. Membership in this group will allow users to create GPOs in the area of the directory where they have access. When a member of this group creates a new policy object, he or she will become the owner of the object and will have full control over the object through the Group Policy Creator Owners permissions.
2. The user must be given permissions to a container in the directory where he or she will be managing group policy. This is done through delegation of control. When in Active Directory Users and Computers, right-click the designated container from the console list and select the **Delegate Control** item. Work your way through the Delegate Control Wizard to select the users who should be given control in the container. Add the Manage Group Policy Links item from the Permissions list, and then finish the wizard.

After these two steps have been performed, the user will be able to create new GPOs in the container where control was given. If you want the user to be able to edit the policies in other objects, you can give the user explicit permissions on the GPO in the directory. The user will only be able to create GPO links in containers where he or she has been granted that permission.

Verifying Group Policy

After you have created and linked GPOs in the directory, you should verify the correct operation of the policy settings before allowing the policy to be processed by users. To do this, you can use the Resultant Set of Policy tool in *logging mode* instead of *planning mode*. Access the RSoP tool just as you did for *planning mode*, but in the first page of the wizard, select the **Logging mode** option button instead of the **Planning mode** button. The settings for generating an RSoP report in *logging mode* are different from those in *planning mode*. The next few paragraphs detail the wizard pages and the settings needed to generate the report.

The first data page of the wizard allows you to select which computer to generate the report for. Your options are the current computer or another computer on the network. You can also select not to include computer configuration settings in the report.



NOTE

Note that only computers with operating systems newer than Windows 2000 will be able to generate RSoP logging data.

The next page allows you to select the user for which the report will be run. You can select the current user or identify a different user from the directory. If you do not want to include user configuration data in the report, you can select the option to only include computer configuration information.

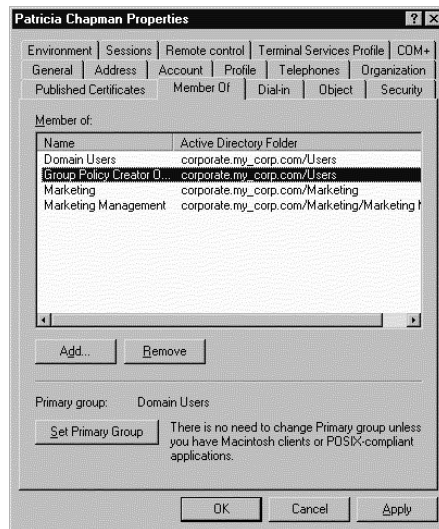
After completing the wizard, you can browse through the policy settings that will be in effect for the user once he or she logs on. There are fewer options needed for the logging mode of RSoP because the tool is not generating any “what if” information in the report. Instead, this report looks at the existing user and designated computer and reviews the policy settings that will be in effect for the user when he or she logs on.

EXERCISE 9.02

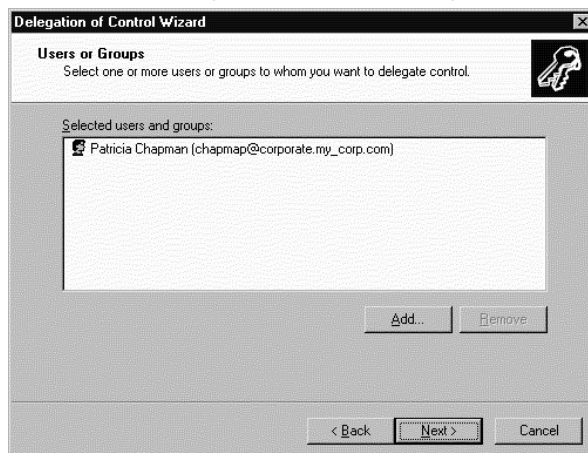
DELEGATING CONTROL FOR GROUP POLICY TO A NON-ADMINISTRATOR

In this exercise, we walk through the process of setting up a non-administrator user to create and manage group policy in a specific container. For this example, one of the managers in the Marketing department will be given permissions to create GPOs in the Marketing container of the directory. After we have her permissions configured, we will log on as the user and create a simple GPO for the container.

1. Open **Active Directory Users and Computers**.
2. Find the user object in the tree, and open the **Properties** for the user.
3. Click the **Member of** tab, and click **Add**.
4. Enter **group policy** in the object name field, and click **Check Names**.
5. The Group Policy Creator Owner group will be recognized. Click **OK**.
6. The group should now be listed in the groups list as shown in Figure 9.24.

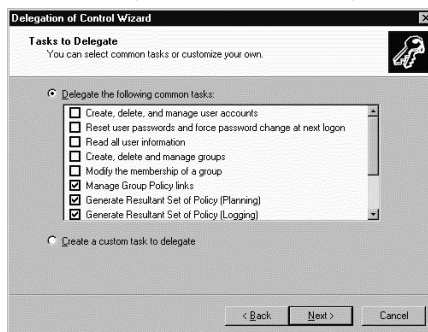
Figure 9.24 Adding the Group Policy Creator Owner Group

7. Click **Apply**, and then click **OK** to close the user Properties window.
8. Right-click the appropriate container in the console tree (in this case, the Marketing container) and select **Delegate Control** from the menu.
9. In the Delegation of Control Wizard, click **Next**.
10. In the Users or Groups window, click **Add**.
11. Select the username from the directory and click **OK**. Repeat the process to add more users if necessary.
12. When the user list is complete, as shown in Figure 9.25, click **Next**.

Figure 9.25 Entering Users in the Delegation of Control Wizard

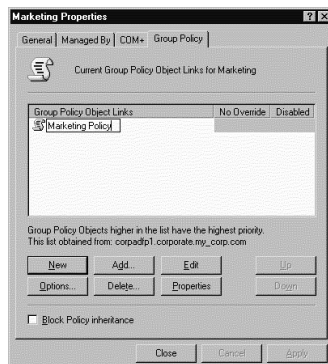
13. In the **Tasks to Delegate** page, click on the **Manage Group Policy links** check box as shown in Figure 9.26. If you would like the user to be able to work with RSOP, enable the **Generate Resultant Set of Policy (Planning)** and **Generate Resultant Set of Policy (Logging)** items as well. Click **Next**.

Figure 9.26 Enabling the Group Policy Settings



14. Click **Finish** to close the wizard.
15. Log on as the user and run Active Directory Users and Computers, or use **Run As** to run the tool as the user.
16. Active Directory Users and Computers should open to the container to which the user was just added.
17. Right-click the container and select **Properties**.
18. Click the **Group Policy** tab.
19. Click **New**.
20. Type in the name of the policy object. In Figure 9.27, the object has been named Marketing Policy.

Figure 9.27 Naming the New Policy Object



21. Click **Edit**.

Now you can go through and set the appropriate policy items as needed.

EXAM
70-294OBJECTIVE
4
4.2.1
4.2.2
4.2.3
4.3.1
4.3.2

Performing Group Policy Administrative Tasks

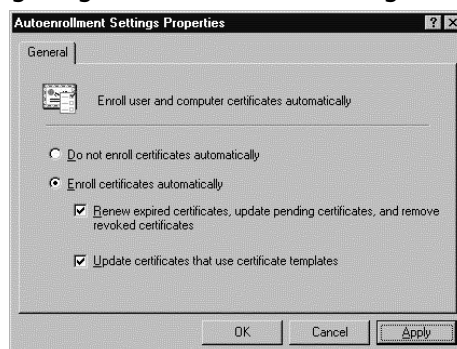
A number of tasks can be performed with group policy settings. This section of the chapter covers some of the more typical administrative tasks that you might perform in setting up group policy for your organization.

Automatically Enrolling User and Computer Certificates

If your organization is using Certificate Services to manage user and computer certificates, you might want to enable autoenrollment of the certificates. Your certification authorities (CAs) need to be configured to support autoenrollment, but without enabling this setting in policy, users have to go through a manual process to enroll.

You will set the autoenrollment policy in both the user configuration and the computer configuration of the GPO. Since you will probably want the settings to apply to all systems in the organization, enable the settings in the Default Domain Policy object at the root of each domain in the organization. Follow these steps to enable this security setting:

1. Open **Active Directory Users and Computers**.
2. Right-click the domain container in the console tree and select **Properties**.
3. Click the **Group Policy** tab and select the **Default Domain Policy**.
4. Click **Edit** to open the Group Policy Object Editor.
5. Expand the **Computer Configuration** object, and then the **Windows Settings** object.
6. Expand the **Security Settings** object, and then select the **Public Key Policies** object.
7. Double-click the **Autoenrollment Settings** object in the right-hand pane.
8. Click the **Enroll certificates automatically** option button.
9. Enable the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box.
10. Enable the **Update certificates that use certificate templates** check box. Your settings should now appear as shown in Figure 9.28.

Figure 9.28 Configuring Autoenrollment Settings

11. Click **Apply**, and then click **OK**.
12. Expand the **User Configuration** object in the console tree, and then the **Windows Settings** object.
13. Expand the **Security Settings** object, and then select the **Public Key Policies** object.
14. Double-click the **Autoenrollment Settings** object in the right-hand pane.
15. Click the **Enroll certificates automatically** option button.
16. Enable the **Renew expired certificates, update pending certificates, and remove revoked certificates** check box.
17. Enable the **Update certificates that use certificate templates** check box.
18. Click **Apply**, and then click **OK**.

If your organization has multiple domains, repeat this process for each domain in the environment. Remember that only systems running Windows 2000 or later will be able to participate in autoenrollment of certificates.

Redirecting Folders

Another feature that is becoming increasingly popular is folder redirection, especially since group policy makes this an easy task to perform. Through group policy, you can specify folder redirection for the following four system folders on the user system:

- Application Data
- Desktop
- My Documents
- Start Menu

Folder redirection can be seen as a subset of roaming profiles. By specifying an alternate location for these folders on a network share, the user has access to these folders no matter which computer he or she uses to log on.

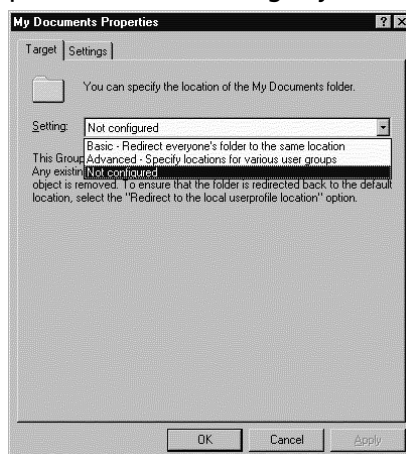
Of the four folders that can be redirected, setting the My Documents folder for redirection is probably the most advantageous. Not only will the user have his or her data available at any computer, but storing this data on the server allows the data to be easily backed up to tape or other offline storage media. As an administrator, you can also set quotas on server storage, helping to keep the size of the My Documents folder in check. You can also take advantage of the offline folders feature of Windows 2000 and Windows XP to keep the data available to users when they are not on the network.

When setting up folder redirection, you should allow the system to create the folders in the location where the data will be directed. A number of permissions must be set correctly to maintain security on the redirected folders. Your best bet is to let the system handle this part of the process.

Folder redirection settings are located in the User Configuration area of the GPO under Windows Settings. To enable redirection of one of the four folders, follow these steps:

1. Right-click the folder name and select **Properties**.
2. In the **Target** tab of the window, you can select the setting to use for redirection, as shown in Figure 9.29. You can select between two options for the location of the redirected folder. The basic option redirects the folder to the same folder path for all users. For the Application Data and Desktop folders, there are three options for the folder location:
 - Creating a directory for each user in the path specified
 - Redirecting all users to the same location
 - Pointing the folder to the local user profile location

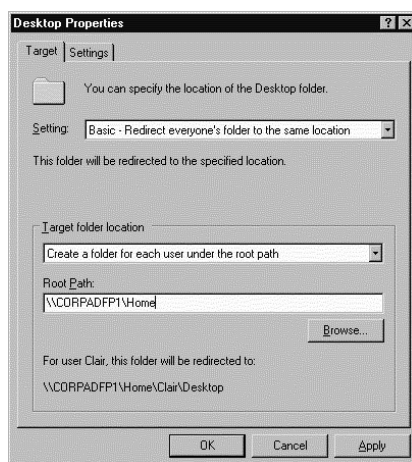
Figure 9.29 Selecting Options for Redirecting My Documents



If you choose to point the folder to the path in the user profile, the folder will point to the default location as if no redirection had been applied. Redirecting the folder to a specific location will create that location either on the network or on a local path, and all users who have this policy applied will point to the same folder. For the Start Menu and Desktop folders, this might be a beneficial setting, as you can centrally control the appearance and contents of those folders in one location, but you need to be aware of the security settings on the folder.

The primary choice for this setting will probably be to create a folder for each user in a location specified, as shown in Figure 9.30. As you can see, when the root path is specified, the dialog box gives you an example of what the folder path will be.

Figure 9.30 Setting the Folder Location for Redirection



The Start Menu and My Documents folders have slightly different options for redirection. When redirecting the Start Menu, you do not have the option of specifying a unique path for each individual user. Whether setting up basic or advanced redirection of the Start Menu, you can only specify one common location for all users or redirect the folder back to the local user profile.

The Start Menu options are simpler than the Application Data and Desktop folder settings, but the My Documents options are more complex. When redirecting the My Documents folder, there are four location options for storing the folder. As with the Application Data and Desktop folders, you can store the My Documents folder in the local user profile, a common directory for all users, or have the system create a folder for each user in a common location. There is a fourth option, however, for My Documents. That option allows you to redirect the My Documents folder to the user's home folder on the network. This option will not create a My Documents folder *in* the user's home folder. It will simply point the My Documents folder to the user's home directory on the network.

There are a few items you should pay attention to if you consider implementing this option. First, you must have implemented the home folder settings for all users, and you must have created those folders prior to implementing this option. Second, the security settings on the home folder are not changed by the folder redirection policy, so you need to be aware of the settings applied to the user home folder on the network. Finally, you have the choice of including the My Pictures folder with the redirected My Documents folder, or having the My Pictures folder stored in a different location. This might be advisable if server disk storage is a concern. If you choose this option, the My Pictures item in the My Documents folder will be a shortcut pointing to the correct location for the actual folder.



EXAM WARNING

It is likely that there will be exam questions related to folder redirection. Be sure you are able to spell out the differences in folder locations for the My Documents and Start Menu folders. You should also be able to explain how folder permissions are set in each of the different redirection settings.

The advanced option allows you to select the folder location based on security group. This is one way to specify a different target location for the folder for different groups of users. You can set multiple security groups to have different target locations within a single GPO in the domain. Another way to accomplish this, especially if you only have a small set of users whose folders should be redirected, is to set folder redirection GPOs at other locations within the directory and filter access to those GPOs based on security.

When selecting the advanced redirection option, you can add the individual security groups for redirection, and have the same choices for folder location as with the basic option. Setting advanced folder redirection is functionally equivalent to setting up multiple GPOs with basic redirection settings and security filtering. The difference is that there is only one GPO to manage instead of several.

Configuring User and Computer Security Settings

EXAM
70-294

OBJECTIVE
4.1.2
4.1.3
4.2
4.2.4
4.3
4.3.3

When browsing through the Group Policy Object Editor, you might have noticed that there are security settings for both the user configuration and computer configuration. Some of these settings are the same for both configurations, such as the Autoenrollment Settings for certificates discussed earlier. There are many differences between the two options, however, and we cover some of those differences in this section.

Computer Configuration

With these security settings, you can provide additional control and management over objects in the directory. The settings contained in this area can govern how users authenticate to computers and other resources on the network, can provide additional permissions or restrictions for resources in the directory, can control audit settings, and can alter group

membership. The settings in this area of group policy are primarily used to specify alternate settings for specific computers on the network.

Table 9.2 lists the main option groups under Security Settings in the Computer Configuration in the Group Policy Object Editor, along with a description of the security setting.

Table 9.2 Security Settings for Computer Configuration

Security Setting Collection	Description
Account Policies	Contains setting groups for password policy settings, account lockout settings, and Kerberos policy settings.
Local Policies	Contains setting groups for auditing policy settings, user rights assignment settings, and security options settings.
Event Log	Contains settings for application, system, and security event logs.
Restricted Groups	Contains groups for specific security restrictions.
System Services	Contains settings for controlling startup and permissions for system services.
Registry Keys	Contains Registry keys and permissions to add.
File System	Contains files or folders and permissions to add.
Wireless Network Policies	Contains policies governing specific wireless network connections.
Public Key Policies	Contains setting groups for Encrypted File System policy settings, Automatic Certificate Request settings, Trusted Root Certification Authorities settings, and Enterprise Trust settings.
Software Restriction Policies	Contains settings, when enabled, for restricting access to certain software, such as 16-bit applications.

User Configuration

There are fewer options for configuring security settings in the User Configuration area of group policy. The two groups of policies in this area are listed in Table 9.3.

Table 9.3 Security Settings for User Configuration

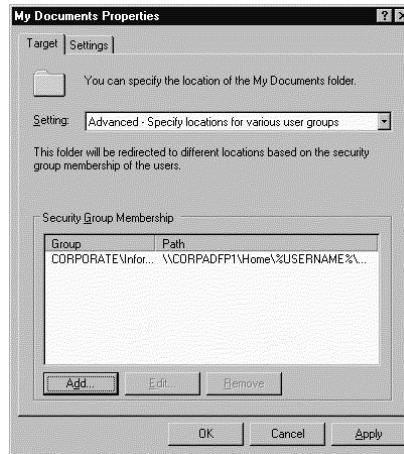
Security Setting Collection	Description
Public Key Policies	Contains settings for certificate autoenrollment and Enterprise Trust.
Software Restriction Policies	Contains settings that identify, through various means, applications that are authorized to run on a system.

EXERCISE 9.03

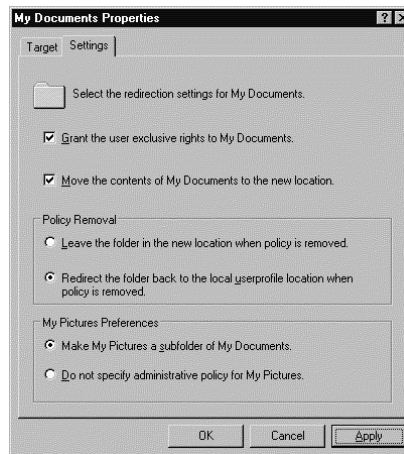
REDIRECTING THE MY DOCUMENTS FOLDER

In this exercise, we walk through the process of redirecting the My Documents folder for a specific group of users in the directory. We will take the Information Technology group and redirect their folders to a shared location on the network, and use advanced redirection to limit folder redirection only to members of that group. We will point the My Documents directory to a common location and use the network's home directory path as the root folder for the redirected folder.

1. Open **Active Directory Users and Computers**.
2. Right-click the domain container and select **Properties**.
3. Click the **Group Policy** tab and click **New**.
4. Name the policy **Folder Redirection Policy** and click **Edit**.
5. Under **User Configuration**, expand **Windows Settings**.
6. Expand **Folder Redirection**.
7. Right-click **My Documents** and select **Properties**.
8. In the **Setting** drop-down menu, select **Advanced – Specify locations for various user groups**.
9. In the **Security Group Membership** pane, click **Add**.
10. In the **Security Group Membership** pane, enter the name of the security group, or click **Browse** and find the group in the directory. This example uses the Information Technology group.
11. In the **Target Folder Location** pane, select **Create a folder for each user under the root path** from the drop-down menu.
12. Enter the UNC path to the desired folder in the **Root Path** field, or click **Browse** to find the desired path. This example uses the path `\\COR-PADFP1\Home` for the root path.
13. Click **OK**.
14. The **My Documents Properties** window should now appear as shown in Figure 9.31. Click the **Settings** tab.

Figure 9.31 Viewing the Redirection Settings for My Documents

15. Make sure the check boxes for **Grant the user exclusive rights to My Documents** and **Move the contents of My Documents to the new location** are enabled.
16. Click the **Redirect the folder back to the local user profile location when policy is removed** option button.
17. The Settings tab should appear as shown in Figure 9.32. Click **Apply**, and then click **OK**.

Figure 9.32 Configuring the Settings for My Documents Folder Redirection

Now when members of the Information Technology group log on, their My Documents folders will be created in the network share, and the data from their existing folders will be moved into the new folders.

Using Software Restriction Policies

One of the relatively new challenges facing system administrators today is the significant increase of malicious code. Not only are more and more individuals writing malicious code, such as viruses, but with the ever-increasing use of e-mail and the Internet, these programs are being spread faster and faster. Some organizations are struggling with the proliferation of other programs, not specifically malicious in nature, but productivity killers nonetheless. Or users might download and install programs that cause conflicts with existing programs, generating additional support calls to your help desk.

Making use of software restriction policies will allow you to place controls on “untrusted” code within your organization. Through a combination of rules, you can identify specific applications or types of applications that are either allowed to run or prevented from running. These rules are powerful and complex, but by themselves cannot provide full protection against malicious code. Use of software restriction policies will augment the protections you might already have in place, but you should not plan to rely solely on these policies to completely protect your environment.

Setting Up Software Restriction Policies

The settings for Software Restriction are located in the Security Settings area of Group Policy. You might have to enable software restriction policies before you can make changes, as most systems do not have these policies enabled by default. When software restriction policies are enabled, you will find a number of settings in the area. The *Enforcement* policy determines if the software policies will apply to all files or exclude library (.dll) files. This policy also identifies whether the policies apply to all users of the system or just to non-administrators. If you want to exclude administrators from software restriction policies, this is where you would set that option.

Another policy you will find here is the *Designated file types* policy. You can specify which file types, based on file extension, to which software restriction policies should apply. You can add additional file types by adding the file extension to the list in this policy window.

The third policy you will find here is the *Trusted Publishers* policy. In this policy setting, you can specify what user level is allowed to enable trust for software publishers, and how to check for expired certificates for those publishers.

The other area located here is the *Additional Rules* folder, where the specific rules you will create for your system will be located.

Software Policy Rules

There are four types of rules that you can use to identify applications to which policy should apply: the hash rule, certificate rule, path rule, and Internet zone rule. Each rule identifies a different way to identify files that should have rules applied. Within the rule, you can set the security setting for the resulting file or files to either Disallowed or Unrestricted. A Disallowed setting in a rule will prevent the user from accessing the file or files. An Unrestricted setting will allow the user to access the file or files.

Hash Rule

When you create a hash rule, you identify a specific file to which you want the rule to apply, and the system generates a hash on the file, including attributes such as date and time of creation and file size. After the policy is in place, the system performs a hash on each file accessed, and if the hash matches the hash in the rule, the rule is applied.

Certificate Rule

When you create a certificate rule, you identify a set of files that are signed by a specific certificate. In creating the rule, you select the specific certificate for the rule. When the system processes a file request, it will check the certificate settings on the file to check for a match against the certificate in the rule, and will process the rule if there is a match. Certificate rules do not apply to .exe and .dll files, but will apply to all other file types listed in the Designated File Types policy.

Path Rule

When you create a path rule, you identify a file or set of files based on their location on disk. The path can identify the path to a folder, a specific file, or a set of files based on a wildcard. When the system processes a file request when path rules are in place, it will compare the file requested to the path rules, and process the rule if there is a match.

Internet Zone Rule

When you create an Internet zone rule, you specify settings based on the Internet zones identified in Internet Explorer: Internet, Local Computer, Local Intranet, Restricted Sites, and Trusted Sites. Internet zone rules only apply to Windows installer packages. If a user downloads an installer package from a site in one of the zones, the zone settings will determine if the user will be able to execute the installer.

Precedence of Policies

Since several rules can be applied to the same program, there is an established order of precedence that is applied. A rule based on a higher precedence will override a conflicting rule applied with a lower precedence.

1. Hash rule
2. Certificate rule
3. Path rule
4. Internet zone rule

Based on this order, if a program is unrestricted based on a hash rule but disallowed based on a path rule, the program will run, as the hash rule has precedence over the path rule. For path rules, there is an additional order of precedence based on the path specified. If there are conflicting path rules, the more restrictive path rule will apply. The following list identifies the precedence of paths from most restrictive to least restrictive.

1. Drive:\Folder1\Folder2\filename.extension
2. Drive:\Folder1\Folder2*.extension
3. *.extension
4. Drive:\Folder1\Folder2\
5. Drive:\Folder1\

When similar rules are applied, such as multiple path rules, the more restrictive rule applies. For example, if a program is set to Disallow in one path rule and set to Unrestricted in another, access to the program will be denied, as Disallow is the more restrictive setting.

Best Practices

The following items include some of the recommendations for implementing software restriction policies.

- **Test, test, test** Never implement software restrictions without testing, especially when applying a Disallow setting. Placing restrictions on certain types of files can negatively impact the operation of your computer and/or the network environment.
- **Couple software restriction policies with access control restrictions** Using access control in conjunction with software restriction makes a more complete restriction solution.
- **Use anti-virus software** Software restriction policies are not a sufficient substitute for a solid anti-virus package. The tools used in conjunction can increase the security of the system, but do not plan on using software restriction in place of anti-virus tools.
- **Use Disallow as default with great caution** If you take the approach of using Disallow as the default and identifying specific applications to allow, be sure you test the system thoroughly. Some applications can launch other applications in normal course of operation. As stated in the first item, test your implementation thoroughly before unleashing it on an unsuspecting audience.

EXAM 70-294
OBJECTIVE
4
4.2.1
4.3.1
5

Applying Group Policy Best Practices

If you have been reading straight through this chapter, you've seen that there are a vast number of ways that group policy can be implemented in an Active Directory environment. How should you approach a group policy implementation in your environment? This section covers some of the best practices related to implementing group policy.

- **The fewer, the better** Keep the number of policies defined as small as possible. Since each user policy the user encounters must be processed at logon, you can keep user logon delays to a minimum by reducing the number of user policies. In addition, a smaller number of policy objects means fewer places for you to look for problems or conflicts when troubleshooting group policy issues. Computer policies are processed at boot time, so reducing the number of these will speed the boot process of the computer.
- **Avoid conflicting policies whenever possible** Although you can set up a lower-level policy to override a higher-level one, you should avoid doing this unless necessary. Again, simplicity should be the rule.
- **Filter out unnecessary settings** If you set up a policy object that only contains user policy settings, set the properties on the object so that only the user configuration portion is processed. This will help cut down on unnecessary processing time.
- **Avoid nonstandard group policy processing whenever possible** Even though you can use Block Policy Inheritance, No Override, and loopback processing options, you should only do so for special cases. Because each of these options alters the standard way in which policy is applied, they can cause confusion when attempting to troubleshoot policy problems.
- **Keep policy objects contained within the domain** It is possible to link a container to a GPO that resides in another domain, but it is unwise to do so. Pulling a GPO from a different domain slows the processing of policy settings at logon time.
- **Use WMI filters sparingly** This suggestion relates to processing time. The more WMI filters there are to process, the longer it takes to apply policy at logon.
- **Keep policy object names unique** If you name each policy object to describe its function, this should not be a difficult practice to adopt. Even though the directory can support multiple GPOs with the same name, it could get very, very confusing for you when trying to troubleshoot a policy problem.
- **Link policies to a container only once** You can link the same GPO to a container more than once, but you shouldn't. The system will attempt to process each policy linked to a container, and even if there are different options set on each instance of the policy link, it can still yield unexpected results.

EXAM
70-294
OBJECTIVE
4
4.2.1
4.3.1
5.1
5.3


Troubleshooting Group Policy

Even the most experienced system administrator is going to encounter times when he or she has misapplied policy or inadvertently created a policy conflict where it was not expected. Fear not, however, because our good friend Resultant Set of Policy and its sidekick, gpresult.exe, can help us out of these jams. Along with a few guidelines, these tools can help you resolve even the stickiest policy problems.

The first step in troubleshooting policy problems is *mapping*. Ideally, when you first start developing a plan for group policy, you will map out policy settings as they apply to your Active Directory environment. Not only will a policy map help you to understand how policy settings will impact the network during planning, but an up-to-date map can help you know where to go looking for problems when they occur. If you do not have a policy map, you should draw one up before you get too far into your troubleshooting process. It might take some extra time up front, but it can save you time and headaches later.

Figure 9.33 shows a sample policy map drawn up based on information used in the examples in this chapter. The diagram was created in Visio, but you can use any diagramming tool (including a pencil and paper) that will help you understand the layout of your policy settings. In the diagram, solid lines indicate a logical connection of Active Directory containers, specifically a domain and its associated OUs. The dashed lines indicate links between containers and GPOs. The policy object is located on the level where it is defined. In Figure 9.33, the Manager Tools policy was created at the domain level, but because the policy was disabled at that level, it is not linked in the diagram.

Figure 9.33 Viewing a Sample Policy Map



Let's walk through a couple of quick scenarios. A user whose user object is in the Marketing container will have group policy applied in the following order upon logon: Local Computer policy, Default Domain policy, Folder Redirection policy, and Marketing policy.

A user whose user object is in the IT Managers container will have group policy applied in this order: Local Computer policy, Default Domain policy, Folder Redirection policy, IT Manager policy, and Manager Tools policy.

A user whose user object is in the Accounting container will have group policy applied in this order: Local Computer policy, Default Domain policy, and Folder Redirection policy. Therefore, if the user in the Accounting container is supposed to have folders redirected by the Folder Redirection policy, but the folders are not being redirected, You should look at the Folder Redirection Policy object and see what options or permissions are on the object that would prevent the user from having the policy applied, and so on.

EXAM
70-294
OBJECTIVE
4.1.1

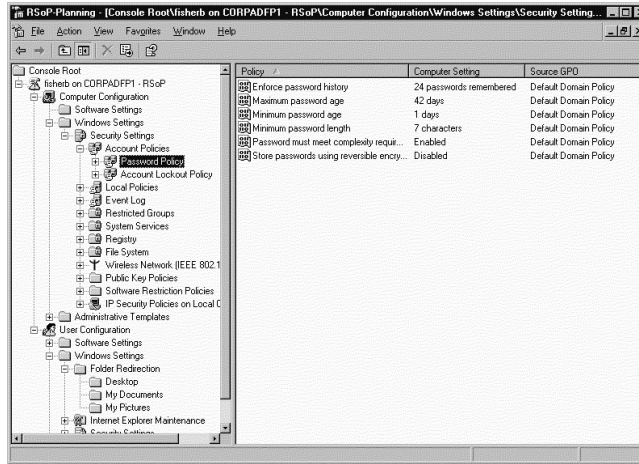
Using RSoP

Just having a policy map will not help you identify the location of policy conflicts in all cases. That's where RSoP comes in. Previously, we've used RSoP to plan our policy environment and test the environment prior to implementation. You can also use RSoP to discover what policy is applied to a user object and where the policy setting came from. To do this, add the RSoP snap-in into the MMC and generate a report based on the user and computer in question as described earlier in the chapter.

Let's say that a user is attempting to change his password, but he continually gets an error saying that his password is too short. You seem to recall that you had set a policy that allowed six character passwords as a minimum, but the user continually gets a message that his password must be at least seven characters. You run an RSoP report on the user and get the result shown in Figure 9.34. Remembering that password settings are a part of computer configuration, you open to that portion of the report and find the minimum password length policy. Sure enough, it's set to seven characters. However, in the window you also see that the minimum password length setting came from the Default Domain Policy object. Therefore, either your recollection of setting a minimum password length of six characters was faulty, or you set that policy in a GPO that was not processed by this user, and now you can find out why.

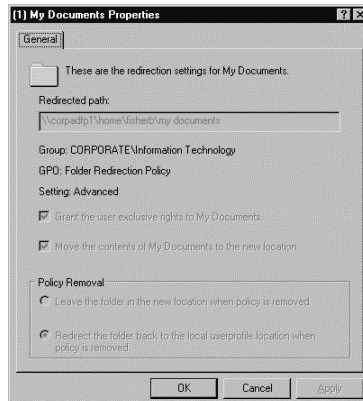
In fact, because the password length setting will apply to all users of the domain, this is an expected result. There is only one set of account policies that is defined for the entire domain, and domain controllers always get the policy from the Default Domain GPO, even when account policies are defined on GPOs linked to OUs. Therefore, account policy settings, such as password length, account lockout, and so on, must be specified in a Default Domain GPO.

Figure 9.34 Viewing RSoP Results



Other settings do not display as clearly in the MMC window. Let's take a quick look at a folder redirection setting that was applied to this same user. When you click the **My Documents** folder under **User Configuration | Windows Settings | Folder Redirection**, the MMC will display information about the Redirection policy. However, the default format of the display does not show all the information at one time, so you can double-click the entry to bring up the **Properties** window shown in Figure 9.35. You can see in the Properties window the location where the My Documents folder has been directed and the settings of the policy that caused this user's folder to be redirected. In this case, the GPO that triggered the redirection is the Folder Redirection Policy object, and it was created in advanced mode, with the user matching the Information Technology group membership. In addition, you can see the settings enabled for this particular redirection policy in the grayed-out check boxes and option buttons.

Figure 9.35 Viewing the Folder Redirection Policy Properties



Using gpresult.exe

Sorting through the information provided by RSoP can be a little daunting, especially if there is a lot of customization occurring through group policy settings. Some types of information can be easier to track down using a different tool—gpresult.exe. gpresult is a command-line tool that produces a text report of the resultant set of policy. Table 9.4 lists some of the command-line parameters that can be used with gpresult.exe. By default, running gpresult.exe with no parameters will generate an RSoP report for the current user on the local computer. Different options can be used to specify alternate users and different computers, as well as limiting the scope of the query.

Table 9.4 Command-Line Parameters for gpresult.exe

Parameter	Description
<i>/s Computer</i>	Identifies the location of a remote computer for the query. <i>Computer</i> can be the name or IP address of the remote system. Do not use backslashes in the computer name.
<i>/u Domain\User</i>	Identifies the user to run the program as, in case the current user does not have permission to generate RSoP data.
<i>/p Password</i>	Identifies the password to use for the user object identified with the <i>/u</i> parameter.
<i>/user TargetUserName</i>	Identifies the user for which RSoP data is to be generated.
<i>/scope {user computer}</i>	Identifies the specific scope, user or computer, for which the RSoP report should be run.
<i>/v</i>	Generates verbose policy information.
<i>/z</i>	Displays all available information about the policy settings. This setting generates much more output than the <i>/v</i> parameter.

The output of gpresult.exe is grouped into several different sections. The first section of the output gives basic information about the user and computer analyzed in the query. This output is shown in Figure 9.36. One of the items of interest in this section is the indication of a slow link connection, listed in the last line of the figure.

Figure 9.36 Viewing the Results of gpresult.exe

RSOP data for CORPORATE\fisherb on CORPADFP1 : Logging Mode

```
-----

OS Type:                               Microsoft(R) Windows(R) Server 2003,
                                         Standard Edition
OS Configuration:                       Primary Domain Controller
```

Continued

Figure 9.36 Viewing the Results of gpreresult.exe

```

OS Version:                5.2.3790
Site Name:                  My_Corp
Roaming Profile:
Local Profile:              C:\Documents and Settings\fisherb
Connected over a slow link?: No

```

The next section of output contains information about the computer settings of the resultant policy. This output is listed in Figure 9.37. The output lists the directory path to the computer objects, the last time policy was applied to the computer, and the object from which the policy was applied. The output also lists the specific GPOs that were applied to generate the resultant policy. Following that list is a section containing security group information for the computer.

Figure 9.37 Viewing the Computer Configuration Output from gpreresult.exe

COMPUTER SETTINGS

```

-----
CN=CORPADFP1,OU=Domain Controllers,DC=corporate,DC=my_corp,DC=com
Last time Group Policy was applied: 7/6/2003 at 3:03:06 PM
Group Policy was applied from:      corpadfp1.corporate.my_corp.com
Group Policy slow link threshold:   500 kbps
Domain Name:                        CORPORATE
Domain Type:                         Windows 2000

```

Applied Group Policy Objects

```

-----
Default Domain Controllers Policy
Default Domain Policy
Folder Redirection Policy
Default Site Policy
Local Group Policy

```

The computer is a part of the following security groups

```

-----
BUILTIN\Administrators
Everyone
Cert Publishers

```

Continued

Figure 9.37 Viewing the Computer Configuration Output from gpresult.exe

```

BUILTIN\Pre-Windows 2000 Compatible Access
BUILTIN\Users
Windows Authorization Access Group
NT AUTHORITY\NETWORK
NT AUTHORITY\Authenticated Users
This Organization
CORPADFP1$
    Domain Controllers
NT AUTHORITY\ENTERPRISE DOMAIN CONTROLLERS
Cert Publishers

```

The next information contained in the output is a breakdown of the user configuration settings. Listed in Figure 9.38, you will find the directory container for the user object, the GPOs that were applied, and a listing of the security groups to which the user belongs.

Figure 9.38 Viewing the User Configuration Information from gpresult.exe

```

USER SETTINGS

-----

CN=Byron Fisher,OU=IT Managers,OU=IT,DC=corporate,DC=my_corp,DC=com
Last time Group Policy was applied: 7/10/2003 at 6:51:05 PM
Group Policy was applied from:      corpadfp1.corporate.my_corp.com
Group Policy slow link threshold:   500 kbps
Domain Name:                        CORPORATE
Domain Type:                        Windows 2000

Applied Group Policy Objects
-----

IT Managers Policy
Default Domain Policy
Folder Redirection Policy
Default Site Policy

The user is a part of the following security groups
-----

Domain Users
Everyone

```

Continued

Figure 9.38 Viewing the User Configuration Information from gpresult.exe

```

BUILTIN\Administrators
BUILTIN\Users
BUILTIN\Pre-Windows 2000 Compatible Access
NT AUTHORITY\INTERACTIVE
NT AUTHORITY\Authenticated Users

    This Organization
LOCAL
Information Technology
IT Management
Domain Admins

```

Running gpresult.exe in verbose mode (*/v*) or *really* verbose mode (*/z*) will give you additional information about the specific policy settings that apply to the user/computer combination. One such entry is listed in Figure 9.39.

**NOTE**

Microsoft recommends that, when you use the */z* switch with the *gpresult* command, you should direct the output to a text file because of the amount of information that is returned. To do so, type: **gpresult /z >filename.txt**.

Figure 9.39 Viewing a Sample Policy Listing from a gpresult.exe Verbose Output

Folder Redirection

```

-----
GPO: Folder Redirection Policy
  KeyName:          InstallationType:  basic
  Grant Type:       Exclusive Rights
  Move Type:        Contents of Local Directory moved
  Policy Removal:   Redirect the folder back to user profile
                   location
  Redirecting Group: Everyone
  Redirected Path:  \\corpadfpl\home\fisherb\desktop

```

EXERCISE 9.04

RUNNING AN RSoP QUERY IN LOGGING MODE

In this exercise, we walk through the steps required to generate an RSoP query in logging mode to produce a report on actual policy settings for a user in the directory. The steps in the exercise will use the sample user and computer information, but you can run this report for any user and computer in your environment, provided you have access to the tools. Running this query will not impact a production system.

1. Open the **Microsoft Management Console**.
2. Select **File | Add/Remove Snap-in**.
3. In the **Standalone** tab, click **Add**.
4. Scroll through the list until you find the **Resultant Set of Policy** item, and then click **Add** and **Close**.
5. Click **OK** to return to the MMC window.
6. Right-click the **Resultant Set of Policy** object in the tree and select **Generate RSoP Data**.
7. In the RSoP Wizard, click **Next**.
8. Make sure the **Logging mode** option button is selected, and then click **Next**.
9. Click the **Another Computer** option button, and then click **Browse**.
10. Find a computer in the directory and select it.
11. In the **Computer Selection** window, after selecting the computer, click **Next**.
12. In the **User Selection** window, click the **Select a specific user** option button.
13. Select one of the users listed, and then click **Next**.
14. In the **Summary of Selections** window, click **Next**.
15. Click **Finish** to close the wizard.
16. Browse through the policy settings in the MMC window.

Summary of Exam Objectives

Group Policy is easier to work with in Windows Server 2003 than in previous implementations. A solid understanding of how Group Policy works is still the best preparation for managing group policy in an Active Directory environment. Group Policy that exists on the local computer is processed before any other policy. As additional policies are processed, policy settings in the most recently processed Group Policy Object (GPO) will override conflicting settings already applied. Group Policy is applied in the following order: local policy (on Windows 2000 and Windows XP computers), site policy, domain policy, then OU policies are applied starting at the highest OU level descending to the lowest. Policies set in a parent OU are inherited by all lower OUs, unless a lower OU has the Block Policy Inheritance setting applied. This setting can be bypassed by applying the No Override setting in a parent OU. GPOs will only be processed by objects that have security rights to view and apply the settings. Unneeded GPOs can be filtered out through these security settings.

Before you implement any group policy settings, you should plan your group policy strategy. You should determine the policy settings you want to enforce and where they should be enabled. You can use the Resultant Set of Policy (RSoP) MMC snap-in to see how policy settings will affect users in the directory. Running the RSoP tool in *planning mode* simulates policy settings for users in the directory without actually enabling those settings. The results of the query show the settings that would be applied to a specific configuration under specific circumstances. When planning policy settings for the user configuration, you apply settings that will affect users no matter where they log on to the network. If there are a large number of user policy objects that have to be processed, user logon times can increase to the point of trouble. You should strive to keep the number of processed policy objects as low as possible, especially over a slow network link. When planning policy settings for the computer configuration, remember that the settings will apply to every user who logs on to the computer receiving the policy settings. When there are special-case systems that need specific settings applied (for example, when a user should have different settings depending on which computer he or she logs on to), you can use computer policy in loopback mode, which will force the computer configuration policies to override user policy settings in all cases.

When you are ready to implement a group policy plan, use the Group Policy Object Editor. The Object Editor can be launched directly from the MMC, or you can create and edit policy objects within Active Directory Sites and Services or Active Directory Users and Computers. Creating a new GPO inside a directory container automatically links that GPO to the container. A GPO can be linked to more than one container, but each GPO should be linked to a specific container only once. GPOs can have names up to 255 characters, which can include spaces, numbers, and special characters. GPO names should reflect their location in the tree or special functions. GPOs can be filtered through security groups, or they can be set to only allow the user or computer configuration part to process. While administrators are the only directory users who have permissions to create, edit, and manage GPOs by default, these tasks can be delegated to other users in the directory.

Some of the common group policy tasks you can enable include setting autoenrollment of security certificates, redirecting folders, and configuring security settings. When redirecting folders, you can set the Application Data, Desktop, My Documents, and Start Menu folders to point to a network location instead of the local computer. These folders can be redirected to a unique folder for each user, or you can set a network share that all users will use, such as for the Start Menu or Desktop folders. Folder redirection settings can be basic, in which case all users who process the policy object will have folder redirection enabled, or advanced, in which case users will have folders redirected based on security group settings.

When troubleshooting problems with group policy, you should have a firm understanding of how your policy environment is supposed to look. Trying to troubleshoot policy problems without this understanding can be like trying to find a small town in Texas without a road map. You can also use RSoP in logging mode to discover which policy settings are actually in effect for a particular user or computer, and find the specific GPOs that applied those settings. You can also use the `gpresult.exe` command-line tool to generate RSoP logging data in a text format.

Exam Objectives Fast Track

Understanding Group Policy

- ☑ Group Policy replaces and improves upon System Policy from Windows NT 4.
- ☑ Group Policy can be used to assign scripts, manage applications, redirect folders, and change Registry settings.
- ☑ Policy settings can be applied to user configuration or computer configuration.
- ☑ Policy is applied in the following order: local, site, domain, OU.
- ☑ Group Policy takes time to apply at logon or startup. The fewer policy settings that must be applied, the shorter the time needed to complete the logon or boot process.

Planning a Group Policy Strategy

- ☑ Use the Resultant Set of Policy (RSoP) tool in planning mode to see the effects of group policy settings prior to implementation.
- ☑ Take slow network links into consideration when developing the group policy strategy.
- ☑ Changes to user configuration settings will apply to users no matter which computer they use to log on.
- ☑ Changes to computer configuration settings will apply to every user who logs on to the computer.

Implementing Group Policy

- ☑ The Group Policy Object Editor is used to make changes to group policy objects (GPOs).
- ☑ New GPOs can be created within the Active Directory Sites and Services tool or the Active Directory Users and Computers tool.
- ☑ GPOs are linked to a container to which they should be applied. A single GPO can be linked to more than one container.
- ☑ Non-administrative users can create and manage group policy if you delegate control to the users.

Performing Group Policy Administrative Tasks

- ☑ Group Policy can be used to set autoenrollment of certificates for users and computers.
- ☑ The My Documents, Desktop, Application Data, and Start Menu folders can be redirected to alternate locations through Group Policy.
- ☑ The My Documents folder can be redirected to a user's home folder on the network, if home folders have been set up in the network.
- ☑ Most security settings are found in the computer configuration settings rather than the user configuration settings.

Applying Group Policy Best Practices

- ☑ Keep the number of GPOs to a minimum.
- ☑ Avoid setting up policies with conflicting settings.
- ☑ Filter out unnecessary GPO processing with security filters or by turning off the user configuration or computer configuration processing on a GPO.
- ☑ Avoid overriding the standard processing order of policies whenever possible.
- ☑ Keep WMI filtering to a minimum.
- ☑ Keep policy object names unique.
- ☑ Link a GPO to a container only one time.

Troubleshooting Group Policy

- ☑ Develop a policy map along with an Active Directory map.

- ☑ Use RSoP in logging mode to view the actual policy settings for a user on a computer.
- ☑ Look for the policy object where the specific policy setting was applied in the RSoP report to help find the place where policy settings might not be correct.
- ☑ Use gresult.exe to produce a text report of RSoP data.

Exam Objectives

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

Q: How many GPOs should I create?

A: That will depend on your environment. There will always be at least one policy that applies to every user, the local computer policy. There are default policy objects for sites and domains, but those objects can be filtered through security groups. Just remember that the more policy objects your users have to process, the longer it will take them to log on to the network each time.

Q: Why should I use group policy? Can't I just do what I want in logon scripts?

A: Logon scripts are good for certain network authentication tasks, but group policy settings can impact a greater scope of network settings than can be configured through a script.

Q: Can I run the RSoP tool on all the computers in my organization?

A: No. The RSoP tool will only collect data for computers with operating systems newer than Windows 2000.

Q: How can I control access to group policy settings?

A: Access to group policy settings can be filtered through security group permissions. By default, the Authenticate Users group has Apply Group Policy permissions enabled for any default or newly created GPOs. Changing security group permissions for the GPO will allow you to control who can process the settings in the GPO.

Q: If I have group policy settings I want to apply to more than one OU, do I have to create a GPO for each OU with the same settings?

- A:** No. You can create a single GPO and link it to multiple OUs to have the same set of policies applied to each OU. This approach to applying policy settings to multiple groups will help ease administration overhead if you need to change the policy settings for any reason. This way, you will change only one object for the settings to take effect. Otherwise, you would have to locate and edit each GPO when making changes.
- Q:** What are the advantages and disadvantages of delegating administrative control for group policy?
- A:** If you have a very large Active Directory implementation with several domains and OUs with different needs, delegating group policy administration to others can ease your system administration headache. It is very difficult for a single person or small organization to keep up with the various needs of a very large organization. Delegating control will allow others more closely in touch with the various subordinate organizations to manage the policy needs more efficiently. The major drawback here is one of organization. Once you hand over the reins to another person or group, you can lose touch with the changes they are making. If you are ultimately responsible for troubleshooting problems throughout the entire organization, being unaware of changes made in GPOs can impede your troubleshooting speed. Delegating control can save you administration time up front, but it could also increase your troubleshooting time on the back end.
- Q:** Is redirecting folders always a good idea?
- A:** Some people think that just because you can do something means that you should. In setting up a user environment within an organization, this is not always the case. In some organizations, redirecting folders will have significant benefits: increased data security by redirecting folders to network shares that are backed up regularly, enforcing standards by pulling Desktop and Start Menu folders from a common, write-protected location, and so on. However, there are just as many reasons not to redirect folders. Users who travel from site to site in a geographically large organization could find themselves waiting very long times to access data when logging on from a location that is across a very slow WAN link to where their redirected folders are located. Plus, how would users respond if the server on which their redirected folders are stored has a catastrophic hardware failure and they are not able to access data for several hours or days while the problem is repaired? As with any networking technology, research and planning should be done to assess the impact of redirecting folders in your specific environment before implementation.

Q: Can I use software restriction policies to prevent the spread of computer viruses in my organization?

A: Yes, you can, but not as a complete solution. With software restriction policies, you can limit the execution of certain types of software and you can look for specific program codes to block. As with any good virus protection software, these “rules” must be updated regularly to keep up with the changes in the types of software that you want to prevent from running. The worst mistake a system administrator can make is to create a set of rules for software restriction and then never look at them again. Part of the task of setting up software restriction policies is to maintain them regularly.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Understanding Group Policy

1. You have just set up a Windows Server 2003 Active Directory network, and you want to use group policies to control user configuration. You have configured local policies on some of the machines in your domain, and you also want to configure some site and OU policies for more granular control, but you are concerned about policies at different levels “canceling each other out.” Which of the following types of GPOs will override settings applied at the domain level? (Choose all that apply.)
 - A. Local
 - B. OU
 - C. Site
 - D. Domain
2. You have been asked to set up a group policy environment in a new Windows Server 2003 Active Directory network. Your supervisor has asked if local computer settings will override settings applied in a domain GPO. You explain to him that policies applied later in the processing order generally take precedence over policies set earlier. In what order are group policies applied?
 - A. OU policies, domain policies, site policies, local policies
 - B. Site policies, domain policies, OU policies, local policies
 - C. Local policies, site policies, domain policies, OU policies
 - D. Local policies, OU policies, domain policies, site policies

3. Your department has just hired a new junior system administrator and has asked you to train him. The trainee has worked some with Active Directory, but has never used Group Policy before. He has been running RSoP in planning mode to get an understanding of where different group policy settings are stored, but he keeps getting confused because he is not seeing the same groupings between the computer settings and user settings in the report. What are the main types of policies for user and computer configurations he should see in the report, as represented by nodes in the console tree?
 - A. Assign scripts, Manage applications, Redirect folders, and Change Registry settings
 - B. Software settings, Windows settings, and Administrative templates
 - C. Security settings, Account settings, and Software settings
 - D. Local settings, Site settings, Domain settings, and OU settings

4. You work for a large company that has just acquired another company in a merger. The acquired company has merged its Active Directory structure into your forest. The new group wants to maintain control over their portion of the directory, but you want to make sure certain that domain policy settings are not changed by GPOs applied at the OU level. How will you achieve this?
 - A. Set the No Override option on the domain GPO.
 - B. Set the Block Policy Inheritance option on the domain GPO.
 - C. Set the Disable Domain Inheritance option on the domain GPO.
 - D. Unlink the domain GPO from the domain container.

Planning a Group Policy Strategy

5. You have been asked to implement group policy for a large, geographically diverse company. The users in the company are used to being able to log on very quickly, and you do not want to slow the logon process significantly when adding group policy settings. Which of the following are ways to reduce the processing time for group policy when a user logs on? (Choose all that apply.)
 - A. Apply the Block Policy Inheritance setting on the OU closest to the logon object to keep all other policies from processing.
 - B. Set the Disable Computer Configuration Settings or Disable User Configuration Settings options in the GPO options.
 - C. Filter access to the GPO with WMI settings.
 - D. Filter access to the GPO with security group permissions.

6. You have been asked by your supervisor to duplicate the group policy settings of the Sales department for the Marketing department. A coworker suggests that instead of creating a new GPO for the Marketing OU, you can just link the existing Sales GPO to the Marketing OU. What are the guidelines for linking GPOs to a container?
 - A. Each GPO can be linked to only one container.
 - B. Each GPO must be linked to a container within the same domain.
 - C. Only one GPO can be linked to the root domain container.
 - D. Each GPO should be linked to a single container only one time.
7. You are the administrator for the corporate Active Directory network. There are four business units that are separated into individual domains that are rather large. How should you approach managing group policy for the corporation?
 - A. Limit each business unit to one Default Domain Policy object in the root of each domain, and apply all policy settings for the domain in that object.
 - B. Identify one or more users in each domain and delegate control to them to create and manage group policy for the domain while retaining the ability to manage policy for each domain.
 - C. Give all users rights to manage group policy for themselves.
 - D. Only allow the administrator to manage group policy for the company.

Implementing Group Policy

8. You just took over as network administrator for a company. Your network consists of a single domain. The previous administrator had set up a group policy for the domain that allowed six unsuccessful logon attempts before an account would be locked out. A series of new computers has been purchased and deployed in the environment, and the local policy on these systems is set to allow three unsuccessful logon attempts before locking an account. You decide that you want to enforce account lockout to occur after three unsuccessful logon attempts across the company. How would you achieve this?
 - A. Set the local policy on each PC to lock out accounts after three attempts, and set No Override on the local policy.
 - B. Set group policy in a domain GPO to lock out accounts after three unsuccessful logon attempts.
 - C. Set the Block Policy Inheritance on the group policy.
 - D. Remove the local policies from each PC.

9. You need to create a new GPO to enable settings for a particular OU. You open Active Directory Users and Computers and select the OU in the tree. What is the next step in the process of creating a GPO for this OU?
- A. From the **Actions** menu, select **Create New GPO**.
 - B. Right-click on the OU and select **Create New GPO**.
 - C. Right-click on the OU and select **Properties**.
 - D. From the **Actions** menu, select **Group Policy Object Editor**.

Performing Group Policy Administrative Tasks

10. You want to enforce minimum password lengths for all users in a particular domain. What is the best approach to doing this?
- A. Set the minimum password length policy in Computer Configuration | Windows Settings | Security Settings | Account Policies in the Default Domain Policies GPO.
 - B. Set the minimum password length policy in User Configuration | Windows Settings | Security Settings | Account Policies in the Default Domain Policies GPO.
 - C. Set the minimum password length policy in User Configuration | Windows Settings | Security Settings | Account Policies in the local policy for each computer on the network.
 - D. Set the minimum password length policy in User Configuration | Windows Settings | Security Settings | Account Policies for each OU in the network.
11. You have been asked to set up folder redirection for a particular set of users. Upper management wants these particular users to have a consistent interface on their computers, specifically the appearance of the Desktop and Start menu. These users will not be contained in a separate OU, and management does not want a separate policy created for this function. How will you accomplish this task?
- A. Set up Basic folder redirection settings in an existing GPO for the Desktop and Start Menu folders, and filter access to the redirection settings based on security group.
 - B. Set up Basic folder redirection settings for the Start Menu, and Advanced folder redirection settings for the Desktop folder.
 - C. Set up Advanced folder redirection settings for the Start Menu, and Basic folder redirection settings for the Desktop folder.
 - D. Set up Advanced folder redirection settings for both the Desktop and Start Menu folders, specifying the specific security groups that should have the folder redirections.

Applying Group Policy Best Practices

12. You have been asked by your project team to draft a policy document for managing group policy within your Active Directory environment. This policy document needs to include a summary of the best practices for implementing group policy. Which of the following statements would you include in your policy document? (Choose all that apply.)
- A. Keep the number of GPOs being processed to a minimum.
 - B. Change Registry settings through Group Policy wherever possible.
 - C. Assign security permissions on GPOs to individual users.
 - D. Maintain standard processing order whenever possible.
13. One of the best practices for redirecting the My Documents folder is to let group policy create a folder for each user in a common path. Why should you avoid redirecting the My Documents folder to the user's home folder on the network? (Choose all that apply.)
- A. You cannot set exclusive rights on the user's home folder through group policy.
 - B. After you redirect the My Documents folder to the user's home folder, you will not be able to change the folder redirection settings.
 - C. You cannot redirect the user's My Pictures folder to the home folder.
 - D. Users must belong to the Redirected Folder Users security group, a setting that is often overlooked by system administrators.

Troubleshooting Group Policy

14. You have been asked to create a special policy environment for testing. You have been given the following requirements: Create a GPO called Test Settings in the root domain container. The settings of the Test Settings GPO should not apply to any users in Active Directory. You should be able to apply and remove the settings to/from an OU with minimal effort. Which of the following options meets these requirements? (Choose all that apply.)
- A. Set No Override at the domain level.
 - B. Rename the Test Settings GPO to break the link to other containers.
 - C. Set Block Policy Inheritance at the domain level.
 - D. Remove the link to the Test Settings GPO from the domain container.

15. A user complains that when he tries to save files to his My Documents folder, he keeps getting an error that he does not have permissions to write to the folder. He also tells you that when he looks at the files in his My Documents folder, he doesn't see any files that he recognizes. The domain policy you created redirects the My Documents folder to a secured share on the network. You suspect that someone has made a change to group policy elsewhere in the domain. How can you find the policy that is impacting folder redirection? (Choose all that apply.)
- A. Run an RSoP logging query for the user with his computer and look in the results for the policy objects applied to the computer.
 - B. Run an RSoP logging query for the user's OU and look in the results for the policy objects applied to the user.
 - C. Run an RSoP logging query for the user and his computer and look in the results for the policies applied to the user.
 - D. Run an RSoP planning query for the computer, ignoring the user settings, and look in the results for the policy objects applied.

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|----------------|-----------------|
| 1. B | 9. C |
| 2. C | 10. A |
| 3. B | 11. D |
| 4. A | 12. A, D |
| 5. B, D | 13. A |
| 6. D | 14. D |
| 7. B | 15. C |
| 8. B | |

MCSA/MCSE 70-294

EXAM
70-294
OBJECTIVE
4.2.1
4.3.1

Deploying Software via Group Policy

Exam Objectives in This Chapter:

- 4.2.1 Distribute software by using Group Policy.
- 4.3.1 Distribute software by using Group Policy.
- 5.2 Maintain installed software by using Group Policy.
 - 5.2.1 Distribute updates to software distributed by Group Policy.
 - 5.2.2 Configure automatic updates for network clients by using Group Policy.

- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

In the preceding chapter, you learned what Group Policy is and how to work with Group Policy Objects (GPOs). One of the most important functions of Group Policy in an enterprise-level network is the ability to automate software deployment throughout the organization, saving network administrators and users a great deal of time and trouble.

In this chapter, you will learn about Group Policy's software installation feature. We'll provide an understanding of the terminology and concepts behind software installation, and we'll show you how to use the components of software installation: Windows installer packages, transforms, patches, and application assignment scripts. You'll find out how to deploy software to users and to computers by assigning or publishing applications.

After covering the concepts, we walk you through the steps of preparing for Group Policy software installation, working with the GPO Editor and setting installation options. You'll find out how to upgrade applications, configure automatic updates, and remove managed applications. We'll also cover how to troubleshoot problems that can occur with Group Policy software deployment.

Understanding Group Policy Software Installation Terminology and Concepts

When Active Directory was first introduced in Windows 2000, one of its heralded features was the ability to distribute software via Group Policy. Although this was a welcome feature, there were many skeptics. However, experience has shown that IntelliMirror technology (of which Group Policy software installation is a part) makes an administrator's job much easier when it comes to managing a large pool of users and workstations.

Maintaining the correct applications, service packs, and so forth on users' workstations can be a daunting task, but with Group Policy, software can be distributed, configured, and maintained in a centralized fashion. From the applications users need to complete their work, to patches and updates that fix bugs or enhance security, software deployment is a very powerful feature.

To take full advantage of the software deployment component of Group Policy, you need an understanding of how it works "under the hood." The first step in understanding is to review some of the basic terminology.

Some of the terms associated with Group Policy software deployment may be unfamiliar if you haven't used this feature before. For example, we'll be talking about two types of deployed applications: *published* and *assigned*. A published application is made available to users through the Add/Remove Programs applet in Control Panel. Each user has the option to install the application, or not, when it is published. An assigned application is "pulled" down to the user's computer or the computer itself. During startup or logon, Group Policy assignments are checked. If software is part of a group policy linked to the organizational unit (OU), domain, or site, then the software is "*advertised*" to the user or to the computer. *Advertising* refers to making the application ready for installation when a

triggering action occurs (the user clicks the application shortcut, the user attempts to open a document associated with the application, or the computer starts up).

Another term with which you'll need to be familiar is *software package* or *Windows Installer package*. A package is a file with the .msi extension that contains a database with all the instructions and information necessary to install the application. We'll talk about *transforms*, which are files with the .mst extension that make modifications to the database contained in the .msi file.

If you don't know the basic concepts, you can easily misconfigure software installation policies, and that can create problems on your network. Before implementing a new feature such as software installation, you should first ensure that you understand both the concepts and the procedures involved. Then, you can start to develop a software deployment plan. When you have a viable plan in place, you can begin to put the software installation feature to work for you on your network. In the next section, we will provide more detailed information about Group Policy software installation concepts.

Planning for Software Deployment

You should plan your software deployment strategy carefully before configuring software installation in Group Policy. This will save time and allow you to target the specific users and computers that need the software you are deploying. Best practices include the following:

- You can deploy software at the site, domain, or OU level. Microsoft recommends that you deploy the software as high in the Active Directory hierarchy as possible, because this will prevent you from having to create numerous GPOs deploying the same software, for individual domains or OUs.
- Rather than use separate GPOs to deploy multiple applications, it is easier to administer multiple applications from the same GPO. This also speeds up logon, since fewer GPOs have to be processed.
- If your organizational needs dictate that there are a number of different groups of users or computers that need different software deployed, you can create OUs for software management and place the appropriate users or computers in them, and then apply a different GPO to each OU.

If you have several GPOs that apply to the same user or computer, remember that Group Policy is applied in the following sequence: at the site level, then at the domain level, and then at the OU level.

Group Policy Software Installation Concepts

You can use Group Policy to deploy software within a domain environment by editing an existing GPO or creating a new one. The GPO must be applied to a domain, OU, or site in Active Directory. When you open a GPO that is applied to one of these units, you'll see two nodes labeled **Software Installation** in the left pane of the Group Policy Editor console: one that is under the **Computer Configuration** node and one that is under the **User Configuration** node.



NOTE

If you open the **Local Group Policy** object on a Windows XP or Windows Server 2003 computer that is a stand-alone computer or member of a workgroup, you will see that there are no **Software Installation** nodes under the **Software Settings** folder in either **Computer Configuration** or **User Configuration**. That's because Group Policy software installation is supported only in a Windows 2000 or Server 2003 domain environment. You can use Group Policy to deploy software to computers running the following operating systems only: Windows 2000 Professional or Server, Windows XP Professional, and Windows 2003 Server. The computers must be members of an Active Directory domain.

As mentioned earlier, Group Policy software installation deals with two basic types of software deployment: assigning and publishing. Which of these you choose determines when the software will actually be installed on the user's workstation.

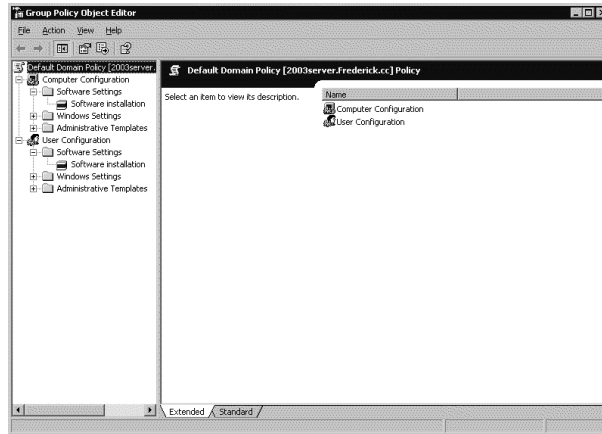
In the following sections, we will look at exactly how each of these options works, and help you determine which is most appropriate for a given situation.

Assigning Applications

The first option is to assign an application. You should assign applications if you want selected users to have the applications available regardless of which computer they are logged on to. An assigned application will “follow” the user from computer to computer within the domain environment.

Applications can be assigned to a user or to a computer by using the appropriate Software Installation node in Group Policy, as shown in Figure 10.1. Using the **Software Installation node** under **Computer Configuration** | **Software Settings** in the left pane of the Group Policy Editor console will allow you to assign the application to a computer. Using the **Software Installation node** under **User Configuration** | **Software Settings** in the same console tree will allow you to assign the application to a user.

Figure 10.1 Group Policy Software Installation



After determining that you want to assign applications (rather than publish them), next you must decide whether to assign applications to users or to computers. Assigned applications are configured based on use. If a particular user will require a word processing or spreadsheet application, you can assign the application to that user. If you will be installing a particular application on every computer in the organization, or to specific computers (for example, all the computers in the Financial department), you can assign the application to the computer objects in Active Directory.



EXAM WARNING

You are likely to see questions on the exam that test your ability to work with the GPO Editor interface, so get as much hands-on experience as possible to ensure that you can answer these point-and-click questions.

When an application is assigned to a user, the application will show up as a shortcut, on which the user can click. This shortcut does not mean that the application is installed, however. The shortcut can be configured to show up in the Start menu or on the desktop. There are also file association changes made to the workstation. This shortcut will “follow” the user, so that it appears on whichever computer the user uses to log on to the network. When the user clicks the shortcut, the application is then deployed to the workstation where the user is logged on. This ensures that users will have the appropriate software, regardless of which workstation they are logged on to.

When an application is assigned to a computer, the software is deployed when it is safe to do so (that is, when the operating system files are closed). This generally means that the software will be installed when the computer starts up, which ensures that the applications are deployed prior to any user logging on. Large application deployments can be done this

way so users won't have to click and wait. Applications that are assigned to computers are available to any user who logs on to that computer. Often, administrators will do large deployments to computers during off hours so when users arrive the next day, they have the updated and installed software ready for use.

Publishing Applications

When an application is published, it is advertised to users through the **Add/Remove Programs** applet in **Control Panel**. This allows users to control when (and whether) the applications will be deployed. Applications that are not required, but which you want to make available as an option for users, are generally deployed this way. If an application isn't used by everybody but might be useful for some to complete a project or task, it can be published for the users to install when and if they need it.

Publishing an application also allows users to uninstall the application from their workstations. This gives users more control over their workstations, whereas assigned applications maintain themselves as installed applications even if the user manually deletes the files.

Figure 10.2 shows the matrix between assigning and publishing software to users and computers.




EXAM WARNING

For the exam, it is important to remember that applications can be assigned to either users or computers, but can be published only to users. If you publish the application, the advertisement attributes are stored in the Active Directory. No changes are made to the Registry until the application is actually installed. When an application is either assigned or published, an application assignment script (with the file extension .aas) is created to hold the advertisement information and the configuration information for the application. This .aas file is stored in the GPO.

Document Invocation

Whether you assign or publish an application, file association changes can be made in the Registry on the workstation where the new application is installed. *Document invocation* refers to the ability of the system to install an application in response to the user's attempt to open a document that is associated with that application. This is also referred to as *file extension activation*. You can control whether applications will be automatically installed by file extension activation. This selection is made by checking a check box on the **Deployment** tab of the **Properties** sheet of the application. You will learn more about editing the Properties options later in the chapter.

For example, if Microsoft Word has been assigned to a computer or user but has not yet been installed, and a user receives a Word document and attempts to open it by double-clicking it, the Installer will immediately install the application and then open the

Figure 10.2 Assigning and Publishing Software Matrix

document with it. It is not necessary for the user to install it via the desktop or Start menu icon, or (in the case of an application assigned to the computer) reboot the computer. The same thing happens if the application has been published, but the user has not chosen to install it via **Add/Remove Programs**. When the user attempts to open the documents, it will be installed automatically. This is also called *on-demand installation*.

What happens if more than one application is associated with the same file extension? Normally, the associated application that was most recently installed on the computer is the one that is used to open the file. You can configure the GPO to set priorities on file extensions, so that you can ensure that the published application that installs when users try to open a file with a specific extension is the right one. This is done by editing the Software Installation Properties of the **User Configuration** or **Computer Configuration** node in the GPO Editor. You will learn more about editing these options later in the chapter.

Application Categories

To make it easier for users to find applications, you can put software into categories. With a large number of applications, users must scroll through the entire list of programs in

Add/Remove Programs to find the applications they want. To simplify the process, you can categorize the applications you assign or publish.

Categories are not predefined and thus need to be set up by the administrator. Grouping common applications together will assist your users in finding the software they need. You can group applications by department, by job function, or in other ways that are logical and meet the needs of your organization's structure. For example, all members of a particular department might need to use the same application, or all secretaries—regardless of department—might need a particular software application. It is not necessary to define categories for each individual GPO; instead, you create categories that will apply to the entire domain.

Group Policy Software

Deployment vs. SMS Software Deployment

Software deployment via Group Policy differs from software deployment via Systems Management Server (SMS). The one simple difference is that SMS is a more controlled software distribution environment. With Group Policy, you set up the deployment as either assigned or published and that is it. With SMS, you can control configuration of items such as bandwidth usage, load balancing, scheduling, and so forth. To accomplish load balancing with Group Policy, you would have to introduce a Distributed File System configuration. Scheduling and bandwidth throttling are available through SMS only, not through Group Policy.

Another key difference between using SMS and using Group Policy is that one is a *pull* model and the other is a *push* model. Software deployment through Group Policy is a *pull* configuration, meaning that the client pulls the software down to a workstation. SMS uses a *push* model where the SMS servers take the responsibility along with the agents to determine what software is needed and the best time to copy the package.



NOTE

SMS installer packages can be converted to Windows Installer packages, using the Installer Step-up tool (ISU) available from Microsoft at www.microsoft.com/smsserver/downloads/20/tools/installer.asp. This utility allows you to create .msi files from SMS setup packages.

Group Policy Software Installation Components

Now that we have discussed the concepts of when and how software should be deployed, let's look at the components involved in using Group Policy to deploy software. In Windows 2003 as in Windows 2000, the Windows Installer technology is the driving force behind this feature.

You will become familiar with four file types as you work with software installation:

- The *application package* is the first and basic file type you will encounter.
- The *transform* gives you the ability to make changes to a package, or transform the package.
- *Patches* are available for many software programs, and you can deploy these with Group Policy.
- The *application assignment script* stores the information regarding assignment or publishing of the application.

In the following sections, we will discuss each of these in more detail.

Windows Installer Packages (.msi)

In the early days of Windows computing, you could use a third-party installation and packaging tool to simplify software deployment (including Microsoft's SMS). Beginning with Windows 2000, the new Windows Installer technology became available, this provides a native packaging and distribution tool for Windows operating systems, and Group Policy provides a way to distribute software without buying a distribution product.

The Installer technology is made up of the following components:

- The Installer service, which is an operating system service that uses Windows Installer packages to perform software installation, modification, and uninstallation.
- The .msi file, which is a group of files compressed together along with the appropriate scripting to install and configure the software. It is essentially a relational database containing a number of tables that holds information about the application. The package can be configured to handle upgrades as well as new installations.
- The application programming interface (API) by which applications interface with the Installer service.



NOTE

The Installer service works with Windows 9x, NT 4.0, 2000, and XP/2003. However, software deployment via Group Policy is only available with Windows 2000 and later operating systems. To use Windows Installer with Windows 9x and NT 4.0, you need to download the *instmsi.exe* file from Microsoft's Web site.

A big advantage of Windows Installer is its ability to “roll back” to the former state if problems occur during an installation. The Installer service can also monitor the state of installed “self-repairing” applications, and detect missing or corrupt program files. The service can then automatically restore the damaged or missing components so that the application will work properly again.

Availability of Installer Packages

Windows Installer packages can be created using packaging tools, but many vendors have their own packages available for download. As with anything downloaded off the Internet, testing should be done prior to full deployment.

Modern Microsoft software comes with Installer packages on the installation CD-ROM. Office 2000 was the first Microsoft application that came with .msi files for software deployment and maintenance. Many software vendors and developers have followed Microsoft's lead and include .msi files with their applications.

Companies can create Installer packages for their proprietary (inhouse) software as long as they have the source code, executables, DLLs, and knowledge of the Registry entries and shortcuts used by the program. Veritas WinINSTALL LE, InstallShield, and other repackaging tools are available from Microsoft and third parties to help you create Installer packages and repackage existing packages.

The database design of the Installer package makes it fast to query and provides for smaller file sizes. The information in the tables includes data that will allow for different installation scenarios, so that there is a set of information about how to install the application clean for the first time, how to install it over a previous version, and so forth. Because the Installer service tracks the installation of the application's features and components, it makes it easier to remove the application completely, without leaving remnants that can cause problems later.

Transforms (.mst)

Packages provided by vendors with their applications are configured and ready to install. This means there isn't much room for flexibility if, for example, you want different installation options for computers in different departments. This is where *transforms* come into play. Transforms are also called *modifications*. That's because a transform is a record of changes to the original package file that allows you to customize the installation by including or excluding particular features. A transform is applicable to a specific Windows Installer package.

Transforms are especially important when you are doing silent or unattended installations. The ability to add or remove certain features or make Registry changes in applying your package makes configuration easier for the administrator. Installing and configuring applications via Group Policy cuts down on the time spent by the administrator in setting up, configuring, and troubleshooting applications for users.

Transforms customize the installation features at the time you assign or publish the application. You can create transforms using the authoring and repackaging utilities we discussed earlier, or the utilities included with applications themselves. Office 2000 included a Custom Installation Wizard to create transforms for making modifications to the application's package when deploying it in your organization. It is often easier to apply a transform rather than repackage an application to make changes.

You associate your transforms with the application when you configure software installation for the application. In the new package that you add via the GPO Editor, you need to select **Advanced published or assigned** in the **Deploy software** dialog box that begins the software deployment process.

Patches and Updates (.msp)

There are times when an application has to be updated because of fixes or new features that are available through a service pack, patch, or other update software. An .msp file is a special type of modification that is used to update an existing Windows Installer package with new information. This allows for easy updates of users' workstations and application of important security patches and other fixes.

With an .msp file, only the updated information needs to be distributed to users. This cuts back on the time and effort required to deploy updates and patches, and cuts down on the amount of network traffic generated by application updates.



NOTE

You should always test patches and upgrades before deploying them throughout the organization in a production environment.

To deploy an .msp update, you generally should advertise the package again to those to whom the original application was assigned or published. Sometimes, the software vendor will provide an entire new .msi package if changes are extensive. In this case, you should just replace the old .msi file with the new one. Otherwise, use the .msp file to make the changes to the original .msi file, and then redeploy the application. This will automatically install the new version for those users and computers where the original was installed, and make it available instead of the old version for those users to whom the original was published, who haven't yet installed the application.

Note that .msp files are not able to make certain changes. For example, they cannot be used to remove Registry keys, or remove or change the names of shortcuts and files. They cannot be used to change product codes, and you can't use them to remove features. These tasks require the use of an .mst transform or a new .msi package.



EXAM WARNING

Remember that .mst (transforms) and .msp (updates) files cannot be deployed by themselves. They must be associated with an existing .msi (application package) file.

Application Assignment Scripts (.aas)

When you set up Group Policy Software Installation and publish or assign applications, an Application Assignment Script (with the file extension .aas) is generated automatically. The Application Assignment Script is stored in the GPO in Active Directory. The script contains information regarding the configuration of the Software Installation. Advertisement information is also stored within the assignment script.

Configuring Deployment: Users or Computers?

How do you decide whether software should be deployed to users or to computers? In many environments, deploying software to users makes the most sense. This is especially true if you want the software to be advertised to particular users, regardless of what computer they are logged on to. If you have employees who move from one workstation to another frequently, and you need to ensure that they always have the proper software available, you should deploy the software to users.

You should also deploy software to users if you want to make certain applications available for users to install optionally if they need it, but do not want it installed if it's not necessary. Because you can publish applications to users (but not to computers), it makes more sense to deploy to users in this situation. An application you assign to a computer will be installed the next time the computer is rebooted, whether any of the users working at that workstation need it or not.

In other situations, it makes more sense to deploy the software to computers. If you have a department where you want to ensure that certain applications are available at every computer, or you need to have an application installed on a specific computer regardless of who uses it, you should deploy the software to the computer(s). Other reasons for deploying to workstations rather than users could be based on keeping software up to date with patches. When software is assigned to a computer, installation does not require a user to be logged on and can happen during startup. This may make more sense for software deployment of patches or software updates.

Deploying Software to Users

GPOs can be linked to a site, domain, or OU (or to a local computer). With that in mind, we will now discuss deployment of software to user objects in Active Directory. Because software installation cannot be done through local group policies, we will be concerned with deploying software at the site, domain, or OU level. The easiest way to deploy software to a specific group of users is to use the OU that contains the user objects. A link can be made to an existing GPO, or you can create a new GPO for this purpose.

Remember that when you deploy software to users, it might be installed soon after they log on. This is determined by whether you assign the software or publish it. If the software is assigned, the software will be installed when the user attempts to run the application from the shortcut or clicks on an associated file. Large installations might make users think that the workstation is locked or froze up, so you have to be careful about whether you assign, publish, or deploy to the workstation instead.

If the application is published, the user can install the application from **Add/Remove Programs** in **Control Panel**. This makes it more likely that the user will know what's going on, since he or she will have chosen to install the application. However, the published application will be installed via document invocation if file associations were set up within the package, which can result in the same problem of a user not realizing an installation is taking place and thinking there is a problem with the computer.



EXAM WARNING

Be sure to have a good understanding of packages, transforms, patches, and application assignment scripts. These items are key pieces to a good software deployment plan. You will need to know how they fit into the big picture with software installation.

Deploying Software to Computers

Most of the same rules discussed in regard to deploying software to users also apply to deploying software to computer objects in Active Directory. However, you need to remember that you can only assign software to computers; there is no publishing to computer objects. Software installation policies can be applied like any group policy to sites, domains, or OUs. In Active Directory, by default each computer object is added to the **Computers** container in the root domain. You will most likely want to set up software deployment to computers by creating an OU, but this depends on your Active Directory design.

When software is deployed to computer objects, the installation generally takes place when the computer boots, prior to the appearance of the **Ctrl + Alt + Del** screen. This means the user cannot log on until all of the software has been installed. This must be considered prior to designing or assigning software installation packages. Assigning too many applications at the same time can cause the workstation to take a long time to start up.



TEST DAY TIP

Be sure you are comfortable with the differences and similarities between assigning versus publishing applications with the Software Installation component of Group Policy.

Using Group Policy Software Installation to Deploy Applications

Now that you know the basics of software installation, let's look at the details and step-by-step procedures involved in completing the process. We will look at the interface used to add software installation packages: the GPO Editor MMC snap-in.

In this section, we will review the Microsoft Windows Installer technology and packages, in the context of how they are used in the process of software deployment. We will also look at how to create your own Windows Installer packages using Veritas WinINSTALL LE. Because the configuration of legacy applications is often an issue in real-world deployment scenarios, we will show you how to deploy software when you don't have a Windows Installer package and do not want to create one. Finally, we will discuss how to set up distribution points.

Preparing for Group Policy Software Installation

Determining which applications you plan to distribute with Group Policy Software Installation is an important first step in the deployment process. Because the GPOs used to deploy software can be linked to a site, domain, or OU, some planning is required. You must take into consideration your Active Directory design and the application needs of your organization.

Some departments will require a particular application, whereas there is no need for that application in other departments. For example, the Financial department may need accounting software that is not used elsewhere. In other cases, an application is required for all those in a particular job function. For example, all project managers may need a particular project management application, regardless of department. There are also times when an application must be distributed throughout the entire enterprise. For example, the software that is used to open and read personnel policies or security policies that apply to all employees will be needed by everyone, regardless of department or job function. Your Active Directory design and organizational needs will ultimately determine your plans for where you will configure Software Installation within Group Policy.

Creating Windows Installer Packages

Although Microsoft provides Installer packages with most of their software programs, the situation is not quite as simple when you have third-party software to install. Then, you may not have the convenience of having a Windows Installer package available, but when this happens, you can use a utility to create an Installer package. One such tool that has been available since Windows 2000 is WinINSTALL. The original version of WinINSTALL LE (Limited Edition) was included on the Windows 2000 Server installation CD-ROM. The software is no longer included on the Windows Server 2003 CD-ROM, but a free MSI repackager, WinINSTALL LE 2003, can be downloaded at no cost at the OnDemand Software Web site at www.ondemandsoftware.com/FREELE2003/. Alternatively, you can

Determining Deployment Methodology

To truly understand how to properly organize your software installation plan, you must first look at your Active Directory structure. Remember that the GPOs used to deploy software are linked to a domain, site, or an OU. Different Active Directory layouts will determine different application deployment plans.

If your directory consists of one domain with OUs that divide your users and computers by location, domain-level group policies probably won't be appropriate for software installation. However, if you have multiple domains for separate geographic locations and your OUs are used for different departments, you have a quandary. Software installation might occur over a slow link if you are not careful, which could result in a great deal of network congestion. This means that your distribution points will need to be carefully planned to prevent this from happening.

download a trial version of the full WinINSTALL product. The full product can be used in environments where deployment needs are more complex, and provides features such as hardware and software inventory, conflict assessment, MSI validation, and multicast replication—many of the same features offered by Microsoft's SMS.



NOTE

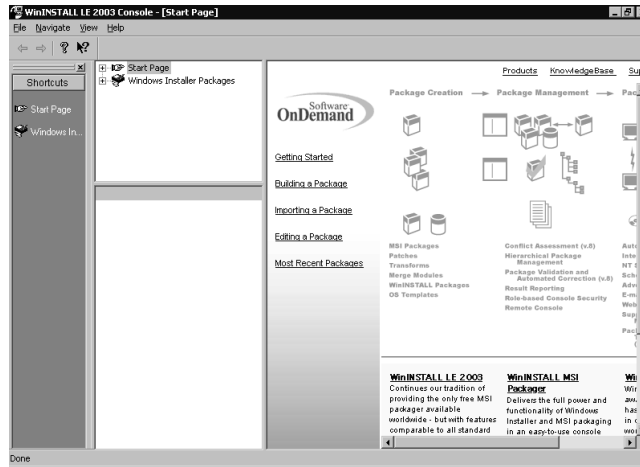
If you have a Windows 2000 Server installation CD-ROM, you can still use the version of WinINSTALL LE to create .msi packages. On the CD-ROM, navigate to the **VALUEADD** directory and you will find the software in **valueadd\3rdparty\mgmt\winstle**. However, WinINSTALL LE 2003 supports the most up-to-date MSI schema (version 2.0).

Figure 10.3 shows the WinINSTALL LE interface.

Before you begin to create your own packages, you should configure a freshly installed workstation to use for this purpose. This will ensure that you have a clean Registry and standard configuration of the operating system. Using a workstation that has had software installed and removed and other changes made to it can cause problems with package deployment.

If you cannot dedicate a workstation for creating .msi packages, you must use a computer that is as close as possible to the configuration of the workstations that will receive the package. The workstation on which you create the packages should be running the same version of the operating system as the computers on which the packages will be deployed. As simple as package creation seems, if configuration steps are not followed closely, you might spend more time troubleshooting problems than successfully deploying software.

Figure 10.3 WinINSTALL LE 2003 Console



EXAM WARNING

Although exam questions might not deal directly with the use of the WinINSTALL interface, successfully answering some questions might be dependent on your understanding of the concept of using a third-party tool to create the appropriate Windows Installer packages. You can use WinINSTALL LE and other packaging programs to both create new installer packages and view the properties of existing ones. You can make changes to the summary information table, although Microsoft recommends that you not change other components—such as required files, shortcuts, and Registry settings—unless you were the author of the original package. Such changes are better made via a transform.



NOTE

An important advantage of using .msi packages to install software is that Windows Installer uses elevated privileges. This means that a user can install an application that is published or assigned to him or her without having to have the user rights that are normally required to install applications.

Using .zap Setup Files

It is possible to publish applications that don't have .msi packages by using the application's Setup program. If you want to deploy software via Group Policy, do not have an .msi file, and do not want to create one, you can instead create a .zap file for the program. The key

point to remember in using .zap files is that they can only be published to users; you cannot assign software to users or to computers by this method. This type of software deployment also has some additional limitations when compared to Windows Installer packages, including the following:

- These installations cannot take advantage of elevated user privileges. This means that if the application requires an account with administrative privileges to be used to install it, users who don't have administrative privileges won't be able to install it even though it is published to them.
- The programs cannot be installed on first use by double-clicking a shortcut, as with Windows Installer packages.
- The system does not automatically repair or remove an application, and you cannot roll back a failed deployment.
- You cannot install features upon first use of the feature, as you can with .msi packages.

If these limitations don't present a problem with the application you want to deploy, the first step is to create a .zap file for the application being deployed. To create a .zap file, you must follow the format prescribed by Microsoft. The .zap file is a text file and can be created in any text editor (for example, Notepad). A sample is available to use as a guide. The following is an example from the Microsoft Knowledge Base article Q231747:

```
[Application]
; Only FriendlyName and SetupCommand are required,
; everything else is optional.

; FriendlyName is the name of the program that
; will appear in the software installation snap-in
; and the Add/Remove Programs tool.
; REQUIRED
FriendlyName = "Microsoft Excel 97"

; SetupCommand is the command line used to
; Run the program's Setup. If it is a relative
; path, it is assumed to be relative to the
; location of the .zap file.
; Long file name paths need to be quoted. For example:
; SetupCommand = "long folder\setup.exe" /unattend
; or
; SetupCommand = "\\server\share\long _
```

```

; folder\setup.exe" /unattend
; REQUIRED

SetupCommand = "setup.exe"

; Version of the program that will appear
; in the software installation snap-in and the
; Add/Remove Programs tool.
; OPTIONAL
DisplayVersion = 8.0

; Version of the program that will appear
; in the software installation snap-in and the
; Add/Remove Programs tool.
; OPTIONAL
Publisher = Microsoft

```

As you can see in the sample file, only two items are required to be completed for a working .zap file. As long as *FriendlyName* and *SetupCommand* are filled in with a Program Name and a string for executing the Setup program, the .zap file will work. The [Application] section is required, and you can also include an [Ext] section; the latter is the file extension section where the application is associated with a file extension in Active Directory. The [Ext] section is optional.

The .zap file is created in a text editor such as Notepad.



NOTE

Creating a .zap file requires less programming knowledge than repackaging an application as an .msi file, making this a popular choice for administrators without extensive programming experience.

After you create the .zap file, you have to add it to your Software Installation configuration within Group Policy. Exercise 10.01 walks you through the steps of publishing an application with a .zap file.

EXERCISE 10.01

PUBLISHING SOFTWARE USING A .ZAP FILE

When publishing software with a .zap file, you first need to determine which GPO you want to edit. After you determine whether to use a GPO that is

applicable to a site, a domain, or an OU, open the appropriate GPO (see the section titled *Working with the GPO Editor* later in this chapter) and make the appropriate addition by following these steps:

1. In the GPO Editor's left console pane, expand **User Configuration**, and then expand **Software Settings**.
 2. Right-click **Software Installation**, select **New**, and then select **Package**.
 3. Change the **Files of type** field to **ZAW Down-level applications package (*.zap)**.
 4. In the **Open** dialog box, navigate to the location of your .zap file or type the path in the **File Name** field.
 5. Click the .zap file you created and click the **Open** button.
 6. Click **Published** as the deployment method in the **Deploy Software** dialog box, and click **OK**.
-

Creating Distribution Points

To distribute software, you must ensure that the users are able to access the needed files from the network. As a network administrator, you must create shared folders on the network known as *distribution points*, to hold the necessary files for installing the deployed applications. A distribution point can be part of a Distributed File System (Dfs) hierarchy or any share point that is available to all users who will need to install the software.

Each share point needs to be configured with the appropriate NTFS permissions to allow access to those who will install the software. This will allow you to control the software that can be installed. If a user doesn't have permissions to access the folder where a package is stored, the software cannot be deployed to that user.

In most cases, it is preferable to control who is able to receive the software through their association and permissions to the GPO itself, but the NTFS permissions must be at least Read and Execute for the distribution point and its subfolders.

Working with the GPO Editor

For those who have worked with the Windows NT 4.0 System Policy Editor, learning to use the Active Directory GPO Editor should be relatively easy. However, deploying applications via Group Policy can be a bit complex. There are many different options to configure when you are setting up a package for deployment. You can deploy software for fresh installations, manage the upgrade of previously installed packages, and remove software from workstations by forcibly uninstalling the software. Every tool is available for managing software within your organization.

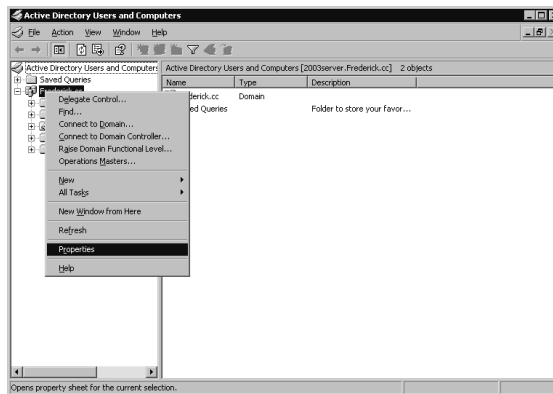
In the following sections, we will show you how to use the GPO Editor to set installation options, assign and publish applications, upgrade applications, and remove managed applications.

Opening or Creating a GPO for Software Deployment

The first step in deploying software via Group Policy is to create a new GPO or open an existing GPO that applies to the site, domain, or OU to which you want to deploy the software. You can open an existing domain policy by following these steps:

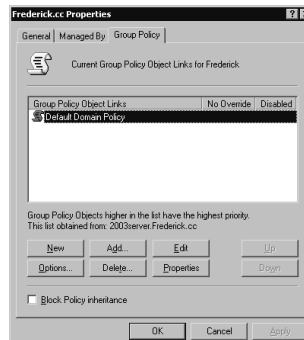
1. Click Start | All Programs | Administrative Tools | Active Directory Users and Computers.
2. In the left console pane of the ADUC tool, right-click the name of the domain and select **Properties** as shown in Figure 10.4.

Figure 10.4 Selecting the Domain Properties for Group Policy



3. Click the **Group Policy** tab as shown in Figure 10.5.

Figure 10.5 Configuring Group Policy



4. Select the policy you want to edit under **Group Policy Object Links**. Click the **Edit** button as shown in Figure 10.6. This will open the policy in the GPO Editor.

Figure 10.6 Editing the Policy



To deploy software at the OU level, follow the same steps except, in step 2, expand the node for the domain, right-click the name of the OU to which you want to deploy the software, and then click **Properties**.

If you want to deploy software at the site level, follow these steps:

1. Click **Start | All Programs | Administrative Tools | Active Directory Sites and Services**.
2. In the left console pane, expand the **Sites** node.
3. In the right details pane, right-click the site to which you want to deploy the software, and click **Properties**.
4. Click the **Group Policy** tab.
5. Select the policy you want to edit under **Group Policy Object Links**. Click the **Edit** button. This will open the policy in the GPO Editor.



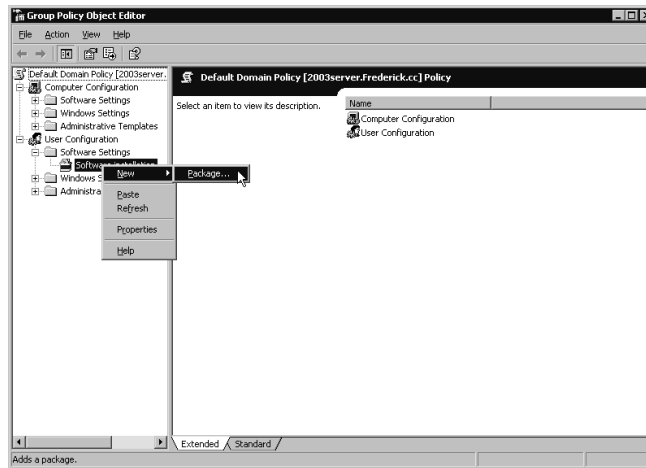
NOTE

To create a new group policy at any of the levels discussed, follow steps 1 through 3 in the instructions for editing an existing policy, and then click the **New** button to create a new GPO.

Assigning and Publishing Applications

Earlier we discussed the concepts of assigning versus publishing applications. Now we will look at the GPO Editor console's interface to become more familiar with the step-by-step process. After you open the GPO Editor, right-click on **Software Installation** under either **Computer Configuration** or **User Configuration** (depending on whether you want to assign the software to computers or assign or publish it to users) and choose **New Package** from the right context menu as shown in Figure 10.7.


Figure 10.7 Configuring a New Package



A dialog box will open asking you for the package you want to use. Navigate to a network location where the .msi file for the software you want to deploy is located. Package files should be stored in a central location. This central location is your distribution point for your software packages. Software packages can generally be downloaded from the manufacturer. Some organizations choose to create their own with other third-party software products.

When you choose a new package, it should be located on a network share. Otherwise, you will receive a message informing you that clients will not be able to install the package, as shown in Figure 10.8.

Figure 10.8 Error Message When Selecting Drive Letter




Next, a prompt will ask you if you want to assign or publish the application, or use the Advanced method, as shown in Figure 10.9.



NOTE

If you are deploying the software from the Computer Configuration node, the selection for Published will be grayed out because software can only be published to users.

Figure 10.9 Creating a New Package



If you choose **Advanced**, you will be presented with the Properties window for your new package. We discuss the options that you can configure in this Properties box in the section titled *Configuring Software Installation Properties* later in this chapter.

EXERCISE 10.02

ASSIGNING SOFTWARE TO A GROUP

This exercise will walk you through the steps of assigning software to an Active Directory group at the OU level. This gives more granularity to the configuration, and this exercise will give you some good hands-on practice in using the interface.

1. Ensure that you have a distribution point (a shared folder containing the .msi package) set up with the appropriate NTFS permissions assigned.
2. Log on as a Domain Administrator.
3. Open **Active Directory Users and Computers** from the **Administrative Tools** menu and right-click the OU to which you want to deploy the software. Select **Properties**.
4. Click the **Group Policy** tab and choose **New** to create a new GPO.

5. Select the new GPO in the list and type a distinguishing name for it.
 6. Click **Edit** to make changes to the GPO.
 7. In the GPO Editor, highlight **Software Settings** under **User Configuration**.
 8. Right-click, select **New**, and then select **Package**.
 9. Navigate to the location of your .msi package. This is the distribution point that you shared earlier. Enter the UNC path so the workstations can find the software.
 10. Next, you are prompted to select whether to publish or assign the application or choose the **Advanced** option. Select **Assigned**.
 11. Click **OK**. The software package name should show up in the right details pane of the GPO Editor.
 12. Close the GPO Editor window. In the OU's **Properties** dialog box, select the GPO under **Group Policy Object Links** and click the **Properties** button.
 13. In the GPO's **Properties** dialog box, click the **Security** tab.
 14. Remove **Authenticated Users** on the **Security** tab and add the appropriate group that contains the users to whom you want to assign this application.
 15. Click **OK** and the application should be ready for deployment.
-

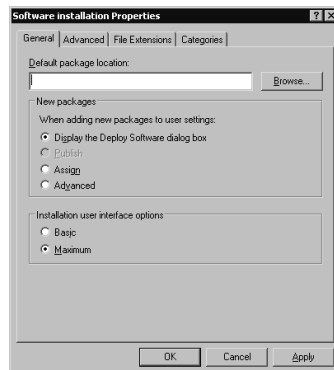


TEST DAY TIP

The more familiar you are with the interface, the better off you will be on the interactive questions you will run across on the exam. More and more exams are going to hands-on or lab type questions, so the more practice you get, the better off you will be. Get to know your interface for deploying software through Group Policy.

Configuring Software Installation Properties

When you first open the GPO Editor, expand **Computer Configuration** or **User Configuration** (depending on whether you want to deploy the software to computers or users), and then expand **Software Settings**. Under **Software Settings**, right-click **Software Installation** and choose **Properties**. You will see a window similar to Figure 10.10.

Figure 10.10 Software Installation Properties

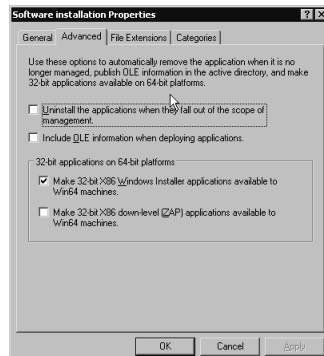
There are four tabs within the Properties of Software Installation. In the following sections, we will discuss the configuration options that can be made with each of these tabs.

The General Tab

On the General tab, you can specify the default location of all packages. Under the New Packages section on that same tab, you can specify the default value for publishing or assigning. The default is to prompt the user to decide at the time of object creation. The last item to be configured on this tab is the User Interface options. This setting determines how much of the installation the user sees. The *Basic* option only shows minimal screen display during software deployment. The *Maximum* option shows all the installation screens as the installation happens.

The Advanced Tab

The Advanced tab has options to be configured such as how to handle 64-bit machines as well as OLE information being published in Active Directory. Figure 10.11 shows the Advanced tab.

Figure 10.11 Advanced Tab of Software Installation

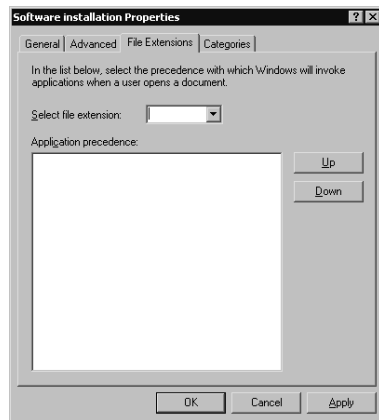
The first option in this window is **Uninstall applications when they fall out of the scope of management**. This means that if a software program was installed with Group Policy and later the account was moved to a different OU, the software could be uninstalled automatically.

You can also choose to have Object Linking and Embedding (OLE) information stored in Active Directory. OLE can be a key part of user interaction and collaboration.

The File Extensions Tab

The File Extensions tab is where you can associate documents and other file types to a specific application that is configured for deployment as shown in Figure 10.12.

Figure 10.12 File Extensions Tab



When you select an extension, you also have to consider some type of order since there are applications that have the same extension for the main file. The Up and Down buttons determine application preference.

The Categories Tab

The Categories tab has the option to create categories so that published applications will be easier to find in the **Add/Remove Programs** applet from **Control Panel**. Figure 10.13 shows the Categories tab.

The Add button allows you to specify new categories. Categories help in finding software installations for users. This is especially helpful when software is published so that users do not have to scroll through the entire list of available software.

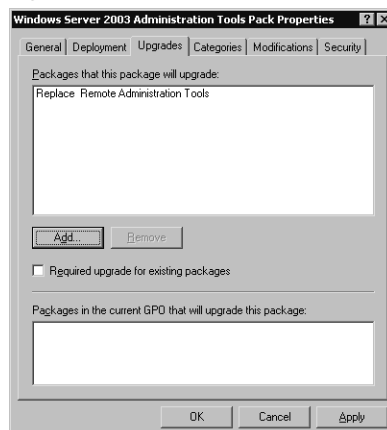
Figure 10.13 Categories Tab**EXAM WARNING**

Be sure you are comfortable with the different tabs and options available for the GPO Editor. These will come up on some exam questions that may require you to be able to point and click on the various options.

EXAM
70-294
OBJECTIVE
5.2.1

Upgrading Applications

For most applications, there will occasionally be upgrades released to address issues with the existing version. The software deployment tools available with Group Policy allow you to maintain control over upgrades by linking the upgrade package together with the original application package. Figure 10.14 shows the Upgrades tab in the Properties of an application.

Figure 10.14 Software Upgrades Tab

The Upgrades tab shows you packages that this package will upgrade, while the bottom pane shows other packages that will be affected by this package. Use the **Add** button to associate this package with the package it is replacing. A good rule of thumb is to use version numbers or exact names with application upgrades to keep things easy to administer. Generally, when software is deployed as an upgrade, the user is prompted to install the upgrade or the user can select to wait until later if he or she is busy and wants to delay the installation.

As we saw earlier, most software installation packages will come from the software manufacturer. These are known as natively authored packages. With natively authored packages, there can be a *declared upgrade relationship* between a package that is an upgrade and other packages. This is part of the database information that makes up a package. The package will know what previous versions it can upgrade and how to handle issues such as files that need to be deleted or kept.

The one catch is that a declared upgrade relationship only works with natively authored packages. With repackaged applications, you have to manually create the upgrade relationship using the Upgrades tab. This is done by clicking the **Add** button on the **Upgrades** tab and selecting the previous versions of those repackaged applications. Active Directory and Group Policy can use this information to upgrade the appropriate users or workstations.



NOTE

It is important to note that upgrading a repackaged application (as opposed to a natively authored application) usually results in removal of the existing application. When the new version of the application is installed, user preferences and other configured settings might be lost.

Automatically Configuring Required Updates

You can use the Upgrades tab to specify whether an upgrade is required or optional. If you want to force users to use the most recent version of an application, you can put a check in the **Required upgrade for existing packages** box. This will automatically upgrade the users' workstations the next time they run the application, or when the computer next reboots if the application is assigned to the computer. A required upgrade is performed whether or not the user wants to upgrade. This is good for applications such as service packs, virus updates, patches, and so forth, and is desirable for productivity applications such as Office if you want to ensure that all users have the same version to make it easier to support and troubleshoot the application.

EXAM
70-294

OBJECTIVE
5.2.2

Removing Managed Applications

In some situations, you may want to discontinue the use of a particular software application in your organization. This might occur because you want to replace the application with a comparable product from a different vendor, and do not want to have some users working with one vendor's product and some with the other's.

Group Policy Software Installation gives you the ability to easily remove software that was deployed with Group Policy. In the GPO Editor, locate the existing package in the right pane and select **Software Installation** in the left pane either under **Computer Configuration** or **User Configuration**. Right-click the application name and choose **All Tasks | Remove**. This will invoke the Remove Software dialog box, as shown in Figure 10.15.

Figure 10.15 Remove Software Dialog Box



There are two removal methods available:

- If you choose **Immediately uninstall the software from users and computers**, the software will be removed the next time the computer reboots (if the application is assigned to the computer) or the next time the user logs on (if the application is assigned to the user). This is called *forced removal*, and automatically removes the software regardless of users' wishes.
- If you want to leave the software on users' workstations but prevent new installations of it, select the **Allow users to continue to use the software, but prevent new installations** option. Users who have it installed will still be able to use it, but no one will be able to install it.

You can select to have the application automatically removed if the GPO no longer applies to a user. To do this, you need to edit the **Deployment** tab of the application's **Properties** dialog box. Check the check box labeled **Uninstall this application when it falls out of the scope of management**.

There is one other thing to remember about software removal. If you have a legacy application that requires the use of a .zap file, you will not be able to take advantage of the removal feature described previously. For the removal feature to work, you must use Windows Installer (.msi) packages to deploy the software.



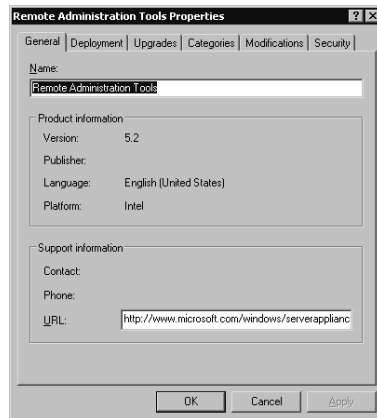
TEST DAY TIP

Make sure you understand how .zap files differ in terms of features and available options from Windows Installer packages, and know which options are available with .msi packages that are not available with .zap files.

Managing Application Properties

After packages are configured, you generally will not have to do much with them. However, there might be occasions when you need to edit an application's properties. To do this, double-click the package in the right details pane of the GPO Editor, with **Software Installation** selected in the left pane, and select **Properties**. Figure 10.16 shows the resulting dialog box.

Figure 10.16 Application Properties

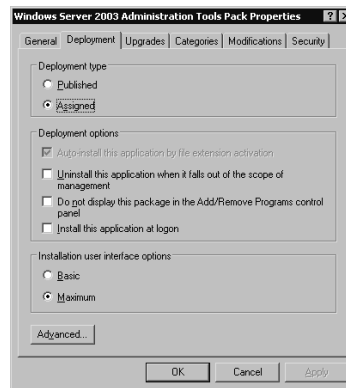


You are presented with six tabs that are used to configure various features, as follows:

- **General** Allows you to rename the package display name and add a URL for support information if desired. Programmers can put contact and telephone information into the package, which will be displayed in those fields. This tab also provides information about the software, including a version number, the publisher's name, language, and the platform on which the software is designed to run.
- **Deployment** As discussed earlier, this tab indicates whether the software is assigned or published as shown in Figure 10.17. This is also where you can select whether the application is to be installed by file extension activation (document invocation); this option is selected by default. Other deployment options include the ability to have the system automatically uninstall the application when it falls out of the scope of management, and the ability to prevent the package from being dis-

played in the Add/Remove Programs applet in Control Panel. You can also select to have the package installed at logon. This tab also allows you to choose the interface that the user will see during installation (basic or maximum). The Advanced button allows you to ignore language when deploying the package, and you can also select to make a 32-bit x86 application available to 64-bit Windows machines. Some advanced diagnostic information, including the product code, deployment count, and script name/path, are also provided in the Advanced Deployment Options dialog box.

Figure 10.17 Deployment Tab



- **Upgrades** As discussed previously, this tab contains upgrade information, including the name(s) of the package(s) that this package will upgrade, whether the package is to be a required upgrade that will be deployed regardless of the user's wishes, and packages in the GPO that will upgrade this package.
- **Categories** This tab allows you to associate the application with a category that is already configured as shown in Figure 10.18. This is especially useful when you publish applications, as they make it easier for users to find the applications within the list in the Add/Remove Programs applet. However, both published and assigned applications can be categorized.
- **Modifications** This tab is used to associate transforms with the package, and control the order in which the transforms are applied to the package, as described in the section titled *Adding and Removing Modifications for Application Packages* later in the chapter.
- **Security** This tab is used to control which users and groups are able to see and use the object in Active Directory, and define the level of access each has. Figure 10.19 shows the Security tab.

Figure 10.18 Categories Tab

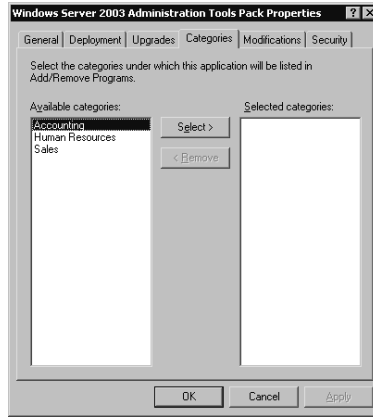
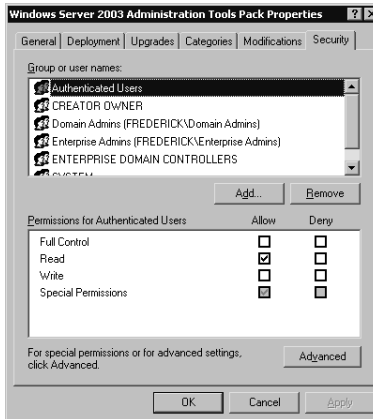


Figure 10.19 Security Tab



By default, the permissions shown in Table 10.1 will apply.

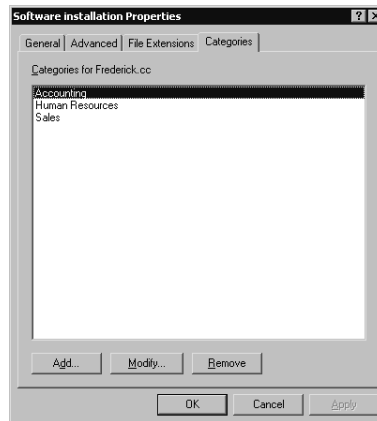
Table 10.1 Default Active Directory Permissions When Adding Packages

User or Group	Default Permissions
Authenticated Users	Read
Creator/Owner	Read/Write
Domain Admins	Full Control
Enterprise Admins	Read/Write
Enterprise Domain Controllers	Read
SYSTEM	Full Control

Categorizing Applications

We mentioned that you can set up categories for your applications to make it easier for users to find the software they need. Categories are set up first. This is done within the Properties of Software Installation. If you right-click on **Software Installation** and go to **Properties**, there is a **Categories** tab as shown in Figure 10.20.

Figure 10.20 Application Categories



Administering categories is simple. The Add button allows you to create new category. You can name it however you want. Many organizations use department names or division names as part of their naming plan.

The Modify button allows you to select an existing category and make modifications. The Remove button will remove a category.

Once the categories are created, the Properties of a package that is already set up will have a Categories tab also. This was shown in Figure 10.17 earlier. There is a list on the left of available categories, and the list on the right tells you what categories this application is setup for.

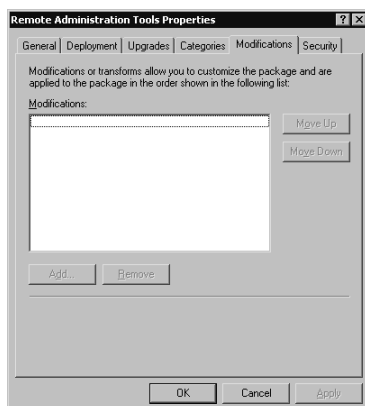
Adding and Removing Modifications for Application Packages

Often you will need more than one version of an application in use on the network, or even on a single machine. You may also need different features enabled for different users. Instead of creating a different package for each unique configuration of an application, you can use modifications, or transforms, to customize the package. To make a transform or modification, you must have the appropriate software. The packaging programs discussed earlier also can be used to create transforms based on a package.

To add and remove modifications, open the application's **Properties** dialog box and click the **Modifications** tab.

As you can see in Figure 10.21, you can assign multiple modifications to a package. Use the **Add** and **Remove** buttons to add the appropriate .mst file to the list or to remove it, and use the **Up** and **Down** buttons to organize the various transforms within the package and control the order in which they will be applied.

Figure 10.21 Software Modifications



EXERCISE 10.03

WORKING WITH SOFTWARE MODIFICATIONS

When working with packages, you might have to apply a transform or modification to the original installation in order to customize the package. This can be because of .ini file changes, Registry changes, or other customization required by your organization. To complete this exercise, you need an existing .msi file and an .mst file. In this exercise, we will apply a transform to a package that is deployed to users at the domain level.

1. Open **Active Directory Users and Computers** and right-click the domain name. Click **Properties**.
2. Select the **Group Policy** tab, select the **Default Domain Policy**, and click **Edit**.
3. In the **GPO Editor**, navigate to the **Software Installation** node under **User Configuration** in the left console pane.
4. Right-click **Software Installation**, select **New**, and then select **Package**.
5. In the **Open** dialog box, navigate to the package (.msi file) you chose for this lab and select it. Click the **Open** button.
6. Select **Advanced** when asked about published or assigned. Click **OK**.

7. Click the **Modifications** tab.
 8. On the **Modifications** tab, click the **Add** button and browse to the .mst file you chose earlier.
 9. Click **OK** to apply the transform to the package.
-

The tricky part about working with modifications is that you must use the Modifications tab when you are initially setting up the package within the group policy. When you select **Advanced** when setting up a package, you are presented with the Properties dialog box for your application. If you select the **Modifications** tab, you will have the opportunity to click the **Add** button.

If, however, you select **Assigned** or **Published** and are not immediately presented with a configuration dialog box, you will not be able to add modifications. The Apply button will be grayed out as we see in Figure 10.20.



TEST DAY TIP

Questions about modifications are likely to show up on your exam. Understanding the purpose and concept is important to answering some of the more difficult questions you will encounter.

Troubleshooting Software Deployment

An important part of any administrator's job is troubleshooting. With software deployment, as with any other aspect of networking, sometimes things go wrong, and when they do, you need to know how to track down the source of the problem and correct it. The Application log in Event Viewer can be a helpful first step in diagnosing some common problems.

Various types of Event Log Error messages might be observed here:

- If you see a series of *MsiInstaller* messages in Event Viewer, you are experiencing a problem with the Windows Installer service. These errors can range from a permissions issue on the distribution point to a problem with the version of Windows Installer you have running on the workstation.
- Watch for *Application Management* messages. These sometimes can indicate the reason why an application didn't deploy properly.
- *Userenv* is another source to look for that may give clues to why software installation failed.

Some common problems encountered with Group Policy software installation and possible methods of resolutions include:

- **Published application doesn't show up in Add/Remove Programs**
 - Check the Group Policy Object link. If there are filters configured or permissions have been changed, the policy may not be getting applied to the user resulting in the software not showing up in **Add/Remove Programs**.
 - Use tools such as `gpresult.exe` and `GPOTool.exe` to further troubleshoot the group policy settings and their application to the workstation or user.
 - Check to see what categories are displaying. All Categories will show all available software for installation.
 - There is an option to mark that does not allow the application to display in Add/Remove Programs control panel.
 - Directory replication not being synchronized can cause software not to show up until all domain controllers are up to date.
- **Software installation not completed when assigned**
 - Check the Group Policy Object link. Make sure there aren't any conflicting Group Policy settings. When planning software installation via Group Policy you need to be careful when you have other policies at higher levels like Domain or OU. This can be especially true if you have the same package configured within multiple GPOs.
 - Check permissions on the GPO. Users or Computers must have Read and Apply Group Policy permissions on the GPO. To check permissions, you must right-click on the site, domain, or OU, select **Properties**, go to the **Group Policy** tab, and click **Edit**. On the **Properties** window of the group policy, the **Security** tab has the permission entries.
 - Check permissions on the distribution point (the shared folder where the .msi package is located). Users need Read and Execute permissions to the distribution point hierarchy.
 - Make sure the software is on a Windows 2003 Share. If you are in a mixed environment, putting shares on Windows NT 4.0 servers is not supported.
 - Double check that if the group policy is set for Computers that it is associated with the appropriate OU. The same thing applies to the group policy set for Users; if software installation is assigned to a user and linked to an OU with only Computer objects, then of course the result would not be successful.



TEST DAY TIP

Troubleshooting is always a big part of each Microsoft exam. Be sure to review all exam objectives and think through the possible troubleshooting scenarios presented here.

■ Name resolution problems

- Name resolution is necessary for users to access the shares where packages are located, whether they are stored on a regular share or within a Dfs hierarchy. Ensure that your DNS servers are running properly. One possibility you should always check when it comes to Active Directory is whether the DNS server is overloaded. Ensure that it is responding to client name resolution requests properly. If name resolution ceases to function, many components of Active Directory will not work properly.

Verbose Logging

When you are experiencing serious application deployment problems, you can turn on verbose logging. This will create a special log file that records information about software installation and group policy application. To turn these features on, you must make a Registry change. For software installation, make sure the following entry exists in the Registry:

```
HKLM\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics
DWORD Appmgmtdebuglevel = 0000009b
```

Turning this feature on will create an `appmgmt.log` located in the `%windir%\debug\usermode\` folder. Only turn this feature on as needed, as it can create a large amount of overhead on the network.

You can also turn on logging for Windows Installer services. You also make a change in the Registry to turn this feature on:

```
HKLM\Software\Policies\Microsoft\Windows\Installer
DWORD Debug = 00000003
```

Setting this value will cause logging to happen for Windows Installer actions. The log created depends on the action itself. Two types of log files can be created: deployment-related or user-related.

With a deployment-related action you will see a log created in `%windir%\temp%\msi*.log`. This is because deployment-related actions run in the system context so the system temp folder is used.

User action would include using add/remove programs to install software. These log files end up in the user's temp folder. The path would be `%temp%\msi*.log`.

Software Installation Diagnostics Tool

Another tool that you can use comes from the Resource Kit. This tool is Software Installation Diagnostics and can be used to gain additional insight into problems you may be experiencing. This tool is also a command-line tool; you have to open a command prompt to run it.

The file is called `addiag.exe`. You can type `addiag.exe /?` and receive a list of commands to become familiar with the tool. You can use this tool to print out information possibly related to problem deployments. It will also generate Event Log entries related to software installation.

Summary of Exam Objectives

Using Group Policy Software Installation, applications can be assigned to users or computers, or published to users. When software is assigned to a user, it will be deployed the first time a user tries to launch that application by clicking the desktop or Start menu shortcut that is installed when you assign the application. Software can also be installed by clicking on a document that has a file extension recognized by Active Directory as being part of an application. This is known as *document invocation*. When software is assigned to computers, it is installed the next time the computer is rebooted, prior to logon. Large software deployments can cause a user to have to wait to log on.

When you publish software, the system advertises the software to be installed but does not install it automatically. The advertisement shows up in the Add/Remove Programs applet in Control Panel. This gives the user the ability to install the appropriate software if he or she chooses to do so. You can set up *categories* to organize the applications so they will be easier for users to find and install.

Software Installation uses the Windows Installer technology. This technology uses Windows Installer Packages or .msi files, which are databases that contain all the information needed to install the software in different scenarios. Software can be customized with configuration changes using an .mst file, also called a *transform* or *modification*. Software Installation uses .msp files to apply patches, updates, and service packs. Assignment scripts are files with an .aas extension that contain information about how the software is advertised on the network and other optional information such as file extension associations.

Deployment, maintenance, and updating of the organization's software infrastructure is an important task for network administrators, and Group Policy Software Installation makes it easier by automating much of the process. A mastery of the concepts and procedures involved in managing applications via Group Policy will serve you well both on the exam and on the job.

Exam Objectives Fast Track

Understanding Group Policy Software Installation Terminology and Concepts

- ☑ Software Installation uses Windows Installer Packages (.msi files).
- ☑ Publishing an application allows the user to install the application from Add/Remove Programs if desired.
- ☑ Assigning an application causes that application to be installed on first use, or when the computer reboots.

- ☑ Transforms are used to make modifications to an existing package. Modifications can only be added when initially setting up the package.
- ☑ Document invocation is the process of installing the software based on clicking on a file with an associated extension to the package.
- ☑ Application categories can be set up to better organize your applications if you have a large number of them. This is especially helpful with published applications; they can be viewed by category, which makes them easier to locate.

Using Group Policy Software Installation to Deploy Applications

- ☑ You can deploy software to either users or to computers.
- ☑ Applications can be assigned to either users or computers, but can be published only to users.
- ☑ A distribution point should be set up to hold the various application packages. This is a shared folder on the network, and users must have the appropriate permissions of Read and Execute to the folder hierarchy.
- ☑ Use the GPO Editor to maintain what packages are related to what group policy.
- ☑ Users must have Read and Apply Group Policy permissions to use the GPO to take advantage of software installation.
- ☑ You can set up categories by right-clicking on the **Software Installation** node and selecting **Properties**. The Categories tab is where you can add custom categories.
- ☑ Upgrades can be set up for existing packages and can be forced to install on users' workstations.
- ☑ Software Installation can be used to install *and* uninstall software .

Troubleshooting Software Deployment

- ☑ All the users associated with a group policy must have permission to the GPO for the settings to take effect.
- ☑ The distribution point can be a problem if NTFS permissions are not set properly.
- ☑ Ensure that name resolution is working so all necessary files can be located during the deployment process.

- ☑ Use the Application log in the Event Viewer to look for MsiInstaller, UserEnv, or Application Management messages and software deployment related messages.
- ☑ If you are in a mixed environment, make sure that your distribution point is not on a Windows NT4.0. Windows 2000 and Server 2003 have this fix already.
- ☑ You can turn on verbose logging for either Software Installation or the Windows Installer services by making a Registry change.
- ☑ The Resource Kit contains a diagnostics tool called Software Installation Diagnostics. It is a command-line tool.

Exam Objectives

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

- Q:** Do I have to use a Windows Installer package to deploy my software via Group Policy?
- A:** No, you can take advantage of a .zap installation, which allows you to use the original setup or installation files that came with the software. Just remember to create the .zap file prior to setting up the software in Group Policy. In addition, be aware of the limitations of .zap files discussed in the text.
- Q:** When does it make sense to assign rather than publish software?
- A:** The determining factor is how you want the software distributed. If you want automatic installation upon first use, then assign the application to users. If you want the user to have the choice to install the software when he or she is ready, publish the application; then users can use **Add/Remove Programs** in **Control Panel** to install the software. If you want to deploy the software to specific computers instead of specific users, you must assign the software because you can't publish to computers.
- Q:** There are multiple shared network locations within our organization. Where should I set up my distribution points?
- A:** That depends on the locations and the links between them. Ideally, you do not want software installation to occur over WAN links because these links tend to be much slower than LAN links. It is best to set up distribution points that will be local to the users/computers to which the software is being deployed.

- Q:** I want to publish a large number of applications to users. How can I organize these multiple applications so users can easily find them in Add/Remove Programs?
- A:** Set up categories to better organize applications, especially when you have many different applications to deploy along with constant updates. Both assigned and published applications can be categorized, but this feature is especially useful when publishing software because it makes it easier for users to locate and install the applications.
- Q:** An application I set up isn't distributing properly. What do I do?
- A:** First check to determine which GPO should be distributing the software. Check for such things as permission problems, assignment problems, or conflicting settings in other group policies.
- Q:** The Published option is grayed out when I try to set up a package for a computer. Why is this?
- A:** You can only assign, not publish, applications to computers.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Understanding Group Policy Software Installation Terminology and Concepts

1. As a network administrator, you need to deploy software to your user's workstations. What must you create on the network so the workstations will be able to locate, read, and execute the .msi file associated with an application?
 - A. A distribution point
 - B. An Admin C\$ Share
 - C. A share named "SFDeployment"
 - D. Nothing has to be created; everything necessary for software deployment is created by default when you select to deploy a package.
2. When you assign an application to a user, which of the following actions on the part of the user will cause the software to be installed? (Choose all that apply.)

- A. Clicking on the shortcut that represents the application in the Start menu or Desktop.
 - B. Opening **Add/Remove Programs** in Control Panel and selecting the application. from the list.
 - C. Double-clicking a file that has an extension that is associated with the application.
 - D. Contacting an administrator to request pre-staging of the user's workstation.
3. What term describes what happens when a user double-clicks on a file with an associated extension that launches the installation of a package configured in Group Policy?
- A. Folder redirection
 - B. Document invocation
 - C. Blocking inheritance
 - D. No override
4. You have configured Group Policy Software Installation to deploy several assigned and published applications. Which of the following is created automatically for each deployed application and stored in the GPO in Active Directory to contain advertisement information about the application configuration?
- A. Microsoft Installer Package
 - B. Logon Script
 - C. Application Assignment Script
 - D. Microsoft Software Transform

Using Group Policy Software Installation to Deploy Applications

5. You want to add a new package to deploy to users in the Marketing OU. What are the steps required to complete this task?
- A. Open **Active Directory Users and Computers**. Navigate to the Marketing OU, right-click, and select **Properties**. Click the **Group Policy** tab and select **Edit**. In the GPO Editor window, expand the **User Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.
 - B. Open **Active Directory Users and Computers**. Navigate to the Marketing OU, right-click, and select **Properties**. Select **Edit**. In the GPO Editor window, expand the **User Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.

- C. Open **Active Directory Users and Computers**. Navigate to the domain node, right-click, and select **Properties**. Click the **Group Policy** tab and select **Edit**. In the GPO Editor window, expand the **User Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.
 - D. Open **Active Directory Users and Computers**. Navigate to the Marketing OU, right-click, and select **Properties**. Click the **Group Policy** tab and select **Edit**. In the GPO Editor window, expand the **Computer Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.
6. What steps should you take to set up a new category for software distribution?
- A. Right-click **Software Installation** in the GPO Editor and select **Properties**. Go to the **Categories** tab. Click **Add**.
 - B. Right-click **Software Installation** in the GPO Editor and select **Properties**. Go to the **General** tab. Click **Add**.
 - C. Right-click the package in the right pane of the GPO Editor, select **Properties**, and then click the **Categories** tab. Click **New**.
 - D. Right-click the package in the right pane of the GPO Editor, select **Properties**, and then click the **General** tab. Click **New**.
7. Which of the following tabs and options are used to force applications to update when deploying upgrades, regardless of whether the user wants to upgrade the application?
- A. “Required upgrade for existing package” on the Upgrades tab
 - B. “Mandatory upgrade” on the Upgrades tab
 - C. “Required upgrade” on the Modifications tab
 - D. “Mandatory upgrade” on the Modifications tab
8. You are a network administrator and you have a number of legacy applications that need to be repackaged. You will be using WinINSTALL LE 2003 to create .msi packages for these applications. You have decided to set up a workstation that will be dedicated to creating these application packages. Which of the following is the best type of machine to use for this purpose?
- A. An existing computer configured with an optimized operating system that has been in service for at least two years.
 - B. A newly installed computer with a clean Registry and default operating system configuration.
 - C. An existing computer that has all of your organization’s applications installed.
 - D. A new computer running only the critical applications but the original Registry settings.

9. You need to deploy an older application to users on your network, and you want to use Group Policy Software Installation to make the deployment easier. However, you are undecided as to whether you should use a program such as WinINSTALL to repackage the application or use a .zap file to install it. In making this decision, you consider the limitations of using .zap files and whether these limitations will affect your deployment. Which of the following are limitations of using .zap files that should be factored into your decision? (Choose all that apply.)
- A. Applications deployed with .zap files cannot be installed by double-clicking a shortcut.
 - B. Automatic repair and removal doesn't work for applications deployed with .zap files.
 - C. Creating .zap files requires more programming expertise than repackaging applications does.
 - D. .Zap files install applications with elevated user privileges.
10. You are the network administrator for a medium-sized financial services company. The users in the Accounts Receivable department all use a popular spreadsheet application that was deployed via Group Policy Software Installation. A new version of the spreadsheet program has been released and includes features that will be useful for some of your users; however, these new features are not required by all users. You want users who are comfortable with the old version and don't need the new features to be able to continue using their current version, but there are several new employees coming on board and you want them to start out with the latest version of the application. How can you disable new installations of an application but not remove the old application from users' workstations?
- A. Right-click **Software Installation** in the left pane of the GPO Editor and click **Removal**.
 - B. Right-click the application in the right pane of the GPO Editor and select **Remove**.
 - C. Right-click the application in the right pane of the GPO Editor and select **Uninstall**.
 - D. Right-click **Software Installation** in the left pane of the GPO Editor and select **Delete**.
11. You are the network administrator for a company that manufactures housewares. You need to deploy a particular software application to all members of the Sales department. All members of this department are already members of the Sales group in Active Directory. Now that you are setting up distribution of software to this same group of people, you would like to use their membership in this security group to define to whom the software will be deployed, if possible. How can you ensure that only the members of the Sales group will receive the software?

- A. You can move all their accounts to a newly created OU for deployment purposes.
 - B. You can use the Security tab on the GPO to configure the appropriate permissions for the Sales group only.
 - C. You can associate the GPO used to deploy the software with the domain so that new users will also receive the application.
 - D. You can remove the existing policy and create a new one that is applied to the group in question.
12. You are the network administrator for a medical billing company. You want to deploy a new billing program to all the users in your organization. It is mandatory that every user have this new software installed. You decided to deploy the application via Group Policy Software Installation, but when you initially added the new package, you mistakenly configured the application to be published. You realize this will not accomplish your purpose because users can choose not to install the application. You want to change the application to assigned status. What is the best way to accomplish this?
- A. Right-click the application in the right pane of the GPO Editor, select **Properties**, and edit the **General** tab.
 - B. Right-click the application in the right pane of the GPO Editor, select **Properties**, and edit the **Deployment** tab.
 - C. Right-click **Software Installation** in the left pane of the GPO Editor, and click **Assigned**.
 - D. You will have to remove the package and add it again to change the status from published to assigned.

Troubleshooting Software Deployment

13. You are a network administrator who handles software installation for your organization. You are having problems with software installation for all of your applications. How would you turn on verbose logging for Software Installation when troubleshooting problems?
- A. Make a change to the GPO by selecting **Verbose Logging** on the **Security** tab.
 - B. A Registry change is required to HKLM\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics.
 - C. Make a change to the package by selecting **Verbose Logging** on the **General** tab of the package.
 - D. A Registry change is required to HKLM\Software\Policies\Windows\Installer.

14. You need to deploy an application to all the users in an OU via Group Policy Software Installation. As part of the deployment process, you have set up a distribution point in which to place the .msi package that will be used to install the software. Users are having trouble installing the program, however, and you suspect that it is because of a permissions issue on the share. You need to ensure that you have set the appropriate permissions on the share point for the users. What are the minimal permissions that should be configured for users to be able to connect to the share and install the software?
- A. Read
 - B. Read and Execute
 - C. Modify
 - D. Full Control
15. You are one of several network administrators who have deployed various software packages and updates within your domain. You have a user who is not receiving software installation as other users are in his department. He can log on and perform normal network functions. You are looking through your configuration to determine the possible cause. Which of the following are potential causes for this user not receiving the software that is associated with the group policy linked to his OU? (Choose all that apply.)
- A. Look for group policy conflicts that would cause this user to not receive the software.
 - B. Check permissions on the GPO to ensure that the user has Read and Apply Group Policy permissions.
 - C. Check that the user account isn't disabled in Active Directory.
 - D. Check permissions on the location of the software itself. The user needs at least Read permission to the distribution point.

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|----------------|--------------------|
| 1. A | 9. A, B |
| 2. A, C | 10. B |
| 3. B | 11. B |
| 4. C | 12. B |
| 5. A | 13. B |
| 6. A | 14. B |
| 7. A | 15. A, B, D |
| 8. B | |

MCSA/MCSE 70-294

Ensuring Active Directory Availability

Exam Objectives in This Chapter:

- 2.5.3 Diagnose and resolve issues related to the Active Directory database.
- 2.4 Restore Active Directory directory services.
 - 2.4.1 Perform an authoritative restore operation.
 - 2.4.2 Perform a nonauthoritative restore operation.
- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

“High availability” is a buzzword in today’s networking world, and that’s because having network services available when users need them is so important to the functioning of the organization. An important part of the administrator’s job is ensuring that the Windows Server 2003 network’s directory services are highly available, and preventing performance problems and downtime issues that can interrupt users’ ability to perform their work on the network.

This final chapter deals with how to maintain high availability of your Active Directory services. You’ll learn about the Active Directory database, and the importance of system state data to Active Directory availability. We’ll discuss fault tolerance plans as well as Active Directory performance issues. You’ll find out how to perform necessary maintenance tasks, such as defragging the database, and you’ll learn how to monitor or move the database.

Next, we’ll address backup and restoration of the Active Directory, and show you the different restoration methods that you can use and when each is appropriate. We’ll walk you through the steps of performing both an authoritative and a normal restore. Finally, you’ll learn how to troubleshoot Active Directory availability problems.

Understanding Active Directory Availability Issues

In this section, we look at the core components of the Windows Server 2003 Active Directory service. First, we discuss the structure and type of database, including its files and related components. You’ll also learn how updates are written to the Active Directory database, and how it recovers in the event of a failure during the update process. We will review the key system configuration components that comprise the system state data, and look at fault tolerance and performance issues involving Active Directory.

The Active Directory Database

The Active Directory service is based on a transactional database system. The word *transactional* refers to the transaction logs that enable the system to have robust recovery and data tracking in the event of unscheduled hardware outages, data corruption, and other problems that can arise in a complex network operating system environment. The heart of the Active Directory service is the database and its related transactional log files, which include the following:

- **Ntds.dit** This file is the primary Active Directory database file (sometimes referred to as the *data store*) that resides on each domain controller (DC). It stores all of the objects, attributes, and properties for the local domain, as well as the configuration and schema portions of the database. By default, this file is installed into the %SYSTEMROOT%\NTDS folder. Although not required, it is recommended that you store this file on an NTFS partition for security purposes.

- **Edb*.log** This file format identifies transaction logs. Transaction log names can take one of several forms, including `edb.log`, `edb00001.log`, `edb00002.log`, and so forth. Each log file is a fixed 10MB in size, regardless of the amount of actual data stored in it. The current log file that is receiving updates to Active Directory is named `edb.log`. When this file is full, it is renamed to `edb00001.log` (or whatever the next number is in the sequence, if 00001 is taken), and a new empty `edb.log` is created. However, these logs don't keep piling up forever; they are regularly purged through a process called *garbage collection*, discussed later in the chapter.



NOTE

If you've worked with Exchange server, you might recognize the `edb.log` name as the name of the Exchange DataBase transaction logs, which work in much the same way (as each log fills, it is renamed and a new `edb.log` is created).

- **Res1.log** and **Res2.log** These files are known as the reserved (Res) log files. Their primary purpose is to ensure that Active Directory does not run out of disk space to use when logging transactions. If there is not enough free space to create a new transaction log, the reserved log is used. Because of this role, these log files are often referred to as *placeholders*. Like the `edb.log` files mentioned previously, these files are 10MB each.
- **Edb.chk** The “checkpoint” file is used to track the updates that have been written to the Active Directory database. You can think of this file as a list that is checked off as updates are flushed to disk from the Active Directory log files. If you shut down the system before all transactions have been written to the database, the checkpoint file will be consulted when you reboot the system so that any remaining transactions can be written to Active Directory.



NOTE

Microsoft recommends that you place the database and the log files on different physical disks, for performance purposes.

Let's take a deeper look at how Active Directory works, and the roles these files play in the process of updating and storing data.

Data Modification to the Active Directory Database

The Extensible Storage Engine (ESE) lies at the heart of the Active Directory database system. Changes to the Active Directory database on a DC occur through two primary means:

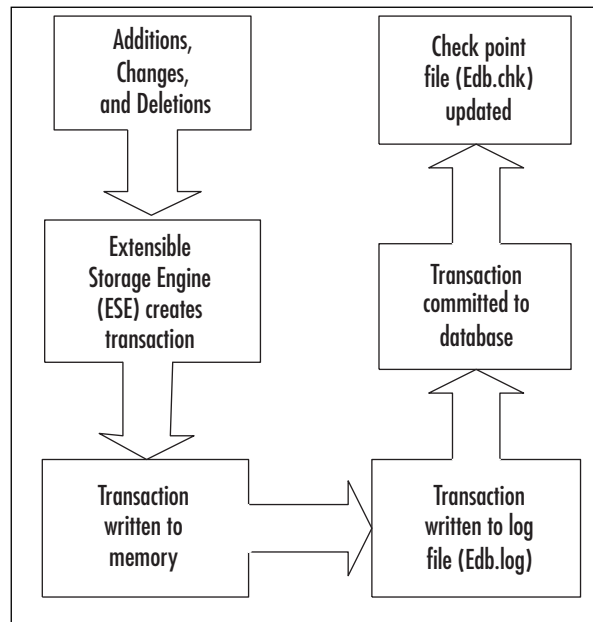
- An administrator creates, deletes, or updates objects in the database.
- Replication information, which contains new objects, deletion requests, or changes to existing objects is received from other DCs.

When changes to the database occur, the ESE captures each change as a single unit known as a *transaction*. A transaction contains the changed data and a set of metadata. This metadata can include the Globally Unique Identifier (GUID) assigned to the object, a timestamp, version, and other information. It's important to note that this update procedure applies to all changes in Active Directory, including objects, properties, and attributes.

A write request occurs when a change is made to the Active Directory. This initiates a transaction that consists of the changes, as well as the metadata described previously. ESE writes the transaction to the transaction buffer in memory, and then writes the transaction to the Edb.log file. After it has been successfully written in the log file, it is written to the Active Directory database file.

If a failure occurs, when Active Directory recovers, it examines the Edb.chk file to determine which transactions have not been written to the database. Transactions are not marked as written in this file until they have been fully committed to the database. This ensures that a failure that occurs partially through the process of writing data will not be marked as completed and leave inconsistent data in the Active Directory database. When a transaction has been committed, Active Directory compares the information written to the database with the information contained in the log file(s). When the two have been verified as identical, the Edb.chk file is updated and the transaction is marked as committed to the database.

Windows Server 2003 uses *circular* transaction logging. This means that, with the exception of the Edb.log, Res1.log, and Res2.log files, the log files are deleted after all of the transactions they contain have been committed to the database. Another important note about logging is that when you back up Active Directory by backing up the system state data (a process we discuss in the next section of this chapter), all events currently waiting to be written in your transaction logs are committed. The logs are fully committed when you shut down or reboot your server. Figure 11.1 illustrates this process.

Figure 11.1 The Active Directory Data Commitment Process**EXAM WARNING**

Be sure to remember that the ESE is the heart of Active Directory, and coordinates transactions between the log files, checkpoint file, and the database.

Head of the Class...**Extensible Storage Engine History**

The ESE is actually an Indexed Sequential Access Method table manager, and is used by the Active Directory, Exchange, Certificate Services, WINS, and the File Replication Service (FRS). It was previously called the Joint Engine Technology (JET); this evolved into ESE in subsequent versions of Microsoft products.

ESE uses a balanced tree (b-tree) structure to store data. This is a common database structure especially designed for fast storage and retrieval of semi-structured data. ESE is different from SQL, Oracle, and similar databases in that it was not designed to deal with direct queries from clients. It is a highly reliable database, designed to be able to recover from system crashes. This makes it an important component of Active Directory's high-availability goals.

The Tombstone and Garbage Collection Processes

The use of transaction logs is designed to ensure the integrity of data that changes on a single DC. Other mechanisms assure the integrity of changed information in relation to the Active Directory replication process. One of these is the *tombstone* process, which is used to replicate deletions from one DC to another.

When an object is deleted within Active Directory, the ESE engine creates a transaction for it (as described in the previous section) and it is moved to the deleted object's container within the database. It is important to note that the object is not immediately purged from the database. There is a delay between when the item is deleted and when it is finally removed from the database altogether. This delay is known as the *tombstone interval* and is set at 60 days by default.



TEST DAY TIP

The tombstone interval is configured in days, whereas the garbage collection interval is configured in hours. Both can be changed in Active Directory using ADSI Edit, LDP, or an ADSI script. However, Microsoft recommends that it is generally best not to change the intervals. The tombstone interval should always be at least as long as the longest replication interval in the forest.

The tombstone process exists to support the multimaster replication strategy of Windows Server 2003's Active Directory service. To understand this better, let's suppose that instead of using tombstoning, the object is immediately purged from Active Directory on the original DC when you delete it. At the same time, the DC's replication partners are notified to delete the object. Most receive the replication request, but one does not. In later replication, this DC might reintroduce the object into the databases of the other DCs. Because the other DCs have fully deleted the object, it might appear as a new object to them.

The tombstone process prevents this from occurring. Each DC holds the object in its deleted items container for the length of the tombstone interval. The default of 60 days allows for plenty of time to pass and ensures that all DCs on the network have sufficient time to receive the delete request. When this interval is reached, the object is marked as *expired*.

You should ensure that backups are performed during the tombstone interval. Restores of directory service data older than the tombstone interval should not be performed to prevent the reintroduction of objects that were deleted during this period but have since been purged from the database.

The garbage collection process works in conjunction with the tombstone process. It runs every 12 hours on DCs by default, and one of its primary functions is to purge expired objects from the database. After the expired objects are purged, any remaining unnecessary log files are deleted and an online defragmentation of the database occurs. This consolidates the free space that was generated by the deletions and increases the performance of the database.



NOTE

An online defragmentation doesn't rearrange the way the data is written to disk, although this can be done through an offline defragmentation. The online defragmentation rearranges how the data is written to the *database*, and compacts the data.

EXERCISE 11.01

MODIFYING THE GARBAGE COLLECTION INTERVAL

To modify the garbage collection interval, you must use either the LDP.exe or ADSIedit.msc (ADSI Edit) utility, or an ADSI script. We will use ADSI Edit, which must be installed from the support pack on the Windows Server 2003 CD-ROM.

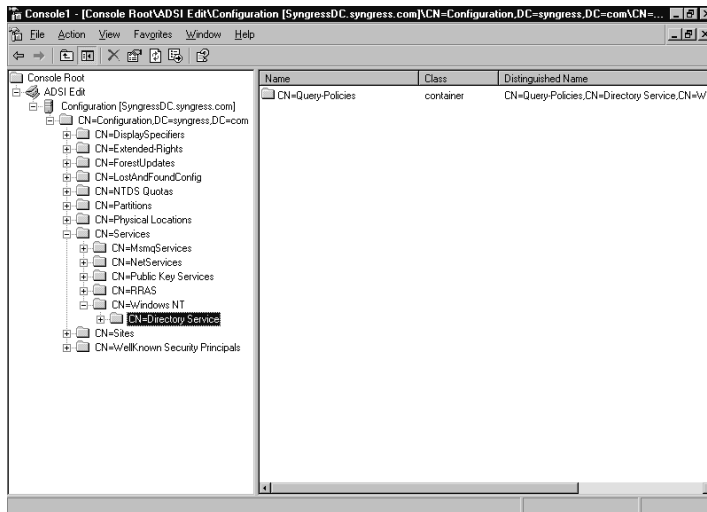
1. Insert your Windows Server 2003 CD-ROM into your CD-ROM drive.
2. On the CD-ROM, run **SUPPORT | TOOLS | SUPTOOLS.MSI**.
3. Click the **Next** button in the Windows Support Tools Setup Wizard that appears.
4. Select the **I Agree** option button on the End User License Agreement page, and click **Next**.
5. Enter your name and organization for licensing purposes, and click **Next**.
6. Choose the path of the location to which you want the utilities to install, and click the **Install Now** button.
7. When setup is complete, click the **Finish** button to exit the wizard.
8. Open **Start | Run**.
9. In the **Open:** text box, type **mmc** to open a blank MMC console.
10. Open **File | Add/Remove Snap-in...**
11. In the Add/Remove Snap-in dialog box, click the **Add** button.
12. Select **ADSI Edit** in the Add Standalone Snap-in dialog box, click the **Add** button, and then click **Close**.
13. Click the **OK** button in the Add/Remove Snap-in dialog box.
14. In the left pane, right-click **ADSI Edit** and select **Connect to...** from the context menu.
15. In the drop-down box for **Select a well known Naming Context:**, select **Configuration** and click the **OK** button, as shown in Figure 11.2.

Figure 11.2 The ADSI Edit Connection Settings Dialog Box



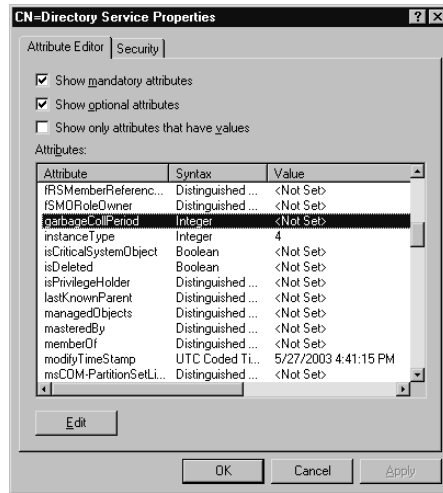
16. Expand the plus sign next to the **CN=Configuration,DC=<YOUR DOMAIN NAME>,DC=<YOUR DOMAIN NAME>** node in the left pane.
17. Expand the plus sign next to the **CN=Services** node in the left pane.
18. Expand the plus sign next to the **CN=Windows NT** node in the left pane. The **CN=Directory Service** node should appear as shown in Figure 11.3.

Figure 11.3 The ADSI Edit Utility with the Directory Service Node Selected



19. Right-click the node for CN=Directory Service in the left pane and select **Properties** from the context menu that appears.
20. In the CN=Directory Service Properties dialog box, scroll down in the **Attributes:** list and click **garbageCollPeriod**, as shown in Figure 11.4. If the current Value is <Not Set>, the default setting of 12 hours is in effect.

Figure 11.4 The Directory Service Properties Dialog Box



21. Click the **Edit** button.
22. In the Integer Attribute Editor dialog box, enter a value in hours for the new garbage collection interval. The minimum value is 1, and the maximum value is one-third of the tombstone lifetime in hours. Thus, the default tombstone time of 60 days equals 3600 hours. The maximum value (or one-third of 3600) would be 1200 hours. A higher value (such as 2000) in this box will be accepted, but will not be effective. The effective value will be 1200.
23. Click the **OK** button in the Integer Attribute Editor. Verify the new setting by checking the **Value** column for the **garbageCollPeriod** entry in the **Attributes:** list.
24. Click the **OK** button in the CN=Directory Service Properties dialog box.
25. Close the MMC console.

The tombstone lifetime can also be set using this tool. To set the tombstone lifetime, substitute `tombstoneLifetime` for `garbageCollPeriod` in the preceding steps. The tombstone lifetime value is set in days; the minimum value is 2.

System State Data

System state data is a term Microsoft uses to refer to a set of core configuration information in Windows 2000, XP, and 2003. The actual information included in the system state depends on the underlying configuration of the operating system, and which components are installed. System state data always includes the following:

- The Windows Registry
- The COM+ Class Registration database
- Boot and system files needed to start the operating system, including Ntldr and Ntdetect.com

Several additional components are included, depending on the configuration of the operating system:

- The Active Directory database and supporting files, if the computer is a DC
- The SYSVOL directory, if the computer is a DC
- The Certificate Services database, if the computer is functioning as a certificate authority (CA)
- The Internet Information Server (IIS) metabase, if IIS is installed on the computer
- Core cluster service configuration information, if the computer is part of a cluster

System state data is most commonly associated with backup and restore operations. The ability to back up the system state was originally created to deal with limitations in the Windows 2000 and XP backup utilities that did not allow open files to be backed up. Many of the most important system and configuration files, including most of those listed as system state, remain open and locked during system operation. Therefore, Microsoft needed a way to ensure that these files could be backed up. The backup tool released with Windows 2000 (and all subsequent versions) included the option to back up all of these critical open resources by selecting the system state backup option.

Backing up open files is no longer a problem when using the Windows Server 2003 Backup utility because of its ability to use *volume shadow copying*, as we will discuss later in this chapter. However, you should still back up these critical files using the system state option in the Backup utility because of the optimizations it provides. Restoring a system state backup restores all key system databases and configuration information. When restored over an existing default installation of Windows Server 2003, the system state completely restores the identity of the computer and all core configuration information.



EXAM WARNING

Remember that the system state data is used to back up all critical system and service information. The information it actually contains depends on the Windows Server 2003 components installed on the machine being backed up. It is important to note that you can only back up and restore system state data on the local computer; you cannot back up or restore system state data remotely.

Fault Tolerance and Performance

You can take several key actions to ensure fault tolerance and maximize performance for the Active Directory database. Maintaining proper backups is, by far, the most important action you can take to provide fault tolerance. We discuss this in more detail later in the chapter. Defragmentation of the Active Directory database is also a key performance component, and we cover this later in the chapter as well.

The very nature of Active Directory replication ensures fault tolerance for the Active Directory database in a multiple DC environment. This is because multiple copies of the Active Directory database remain on the network when any single DC becomes unavailable. Site topology has a strong impact on performance by ensuring that users are directed to domain components located near them, rather than those located across (slow) WAN links.

At the server level, Microsoft recommends hardware redundancy to promote fault tolerance on your DCs. This can include sophisticated systems with hot swap memory, additional processors, hot spare redundant drives, and so forth, or even server clustering, which is the ultimate in server redundancy. However, Microsoft's recommendations primarily refer to redundancy in regard to where the AD database and log files reside in the disk subsystem. For performance reasons, Microsoft recommends that the Active Directory database and log files be on separate physical disk drives attached to separate hard drive controllers or channels. In other words, they recommend that each disk be in its own data path so that there is no contention between these components in the file system. This means that both drives can be read from and/or written to at the same time.

SCSI technology can provide this same effect using a single controller. IDE technology can only provide this ability when a different controller is dedicated to each disk drive. In addition to separate data paths, Microsoft also recommends the use of RAID-1 or RAID-5 technology to improve performance and fault tolerance. RAID stands for Redundant Array of Independent (or Inexpensive) Disks and can be implemented via software using Windows Server 2003 or via hardware by purchasing one or more RAID capable disk controllers. Hardware RAID offers superior performance because it typically comes with its own onboard processor and memory, and thus does not tax the general resources of the computer system. In the following sections, we discuss each of the recommended RAID levels in more detail.

RAID-1

RAID-1 is more commonly referred to as *disk mirroring* or *disk duplexing*, and involves the use of two disk drives. It can be implemented through hardware, or through software using the disk management features built into Windows Server 2003. A RAID-1 volume created using the Windows software is referred to as a *mirrored volume*. With RAID-1, each drive is a “mirror image” of the other. *Disk duplexing* refers to mirroring when each drive is attached to its own separate controller or communication channel.



NOTE

Fault tolerant software RAID (RAID-1 and RAID-5) can only be implemented in Windows Server 2003 on disks that have been converted from basic to dynamic status. Dynamic disks offer more flexibility in management options, but cannot be accessed locally by operating systems prior to Windows 2000. Because it is recommended that production servers not be installed in a dual boot configuration, this should not be an issue with your DCs.

Duplexing can improve both read and write performance. In a duplexed environment, both disks can be read from or written to simultaneously. In a mirrored environment, only one of the drives can be read from or written to at a time. Because data must be written to both drives, write performance can be slightly degraded in a mirrored environment. For the reasons stated previously, this is less true in a duplexed environment. Data can be read from either disk in a mirrored environment and from both disks simultaneously in a duplexed environment, thereby increasing read performance. Fault tolerance is enhanced because if one of the drives malfunctions in a mirrored environment, or one of the drives, controllers, or communications channels malfunctions in a duplex environment, the server remains functional using the other hardware in mirrored configuration.

A duplexed volume appears to the operating system as a mirrored volume. The difference is in the hardware configuration; the software does not recognize any difference between mirrored and duplexed volumes. If a disk in a RAID-1 volume fails, you should break the mirror, add a new disk, and recreate the mirror to restore fault tolerance.



TEST DAY TIP

When using software RAID, the Windows Server 2003 boot and system files can only exist on a RAID-1 (mirrored or duplexed) volume. You cannot use a Windows RAID-5 volume for disks that hold the boot or system volumes. This limitation does not exist when using hardware RAID. In addition, hardware RAID offers higher performance and is preferred if the extra cost is acceptable to the organization.

RAID-5

RAID-5, like RAID-1, can be configured using special hardware or by using the Windows Server 2003 software. This level of RAID is commonly referred to as *disk striping with parity*. When Windows Server 2003 software is used to create a RAID-5 solution, Microsoft refers to it as a RAID-5 volume. RAID-5 technology can dramatically enhance both disk read performance and fault tolerance. RAID-5 involves the use of three or more disk drives that form an array. For optimum performance, each disk should be on its own controller or communications channel. When this is the case, all disks can be written to or read from simultaneously. Data is written to each of the disks in the array in 64KB chunks. This means that 128KB of data will be written in two 64KB chunks, one on the first hard drive in the array and one on the second. Because these writes can happen simultaneously, disk write time can be reduced using RAID-5 technology.

However, when data is written to a RAID-5 disk array, it also has fault tolerant information generated and written. This fault tolerance information is known as *parity*. In our previous example, the third disk in the array would have parity information written to it. Depending on the RAID-5 technology employed, the time it takes to calculate and write this parity information can negate the performance gains discussed previously when writing to the array. However, performance is still dramatically improved for disk reads. Parity is not needed during normal array read operations, which means that in our example, only the two 64K chunks need to be read back from the disk. The performance gain comes because they can be read at the same time.

As with mirroring, a RAID-5 volume can continue to function when a disk fails. The full information contained on the disk can be reconstructed mathematically using the information on the remaining disks in the array, including the parity information they contain. You should be aware, however, that read times will dramatically slow when disk failure has occurred. This is because of the parity calculations that must occur to reconstruct the data that is being requested.

If a disk in the RAID-5 volume fails, replace it as soon as possible to restore fault tolerance and increase performance.



TEST DAY TIP

If more than one drive malfunctions simultaneously in either of these RAID systems, data must be restored from tape.

Performing Active Directory Maintenance Tasks

In this section, we'll look at some of the day-to-day and less routine tasks that you can expect to perform in the process of managing your Active Directory environment. We'll

start with an overview of defragmenting the Active Directory database to improve performance and recover disk space. Next, we'll examine how you can move the Active Directory database and log files, which is sometimes necessary to increase performance or when free disk space gets tight. Finally, we'll look at key Event Viewer entries and performance counters that you should monitor to spot trouble early and keep your DC at maximum performance. In the following sections, we address each of these topics individually.

Defragmenting the Database

As mentioned previously, by default, Windows Server 2003 begins a maintenance cycle every 12 hours, known as the garbage collection process. The final portion of the Garbage Collection process is the performance of an online defragmentation and re-index of the Active Directory database. This is done to improve the performance of the database. To understand how it works, let's look at what fragmentation is and how it can occur.

Understanding Active Directory Database Fragmentation

When first created, the Active Directory database is a blank slate. Data is written to and read from it *contiguously*. This means that it is written to and read from in chunks of data that are next to one another on the physical disk. Over time, as you use the database, information is deleted from it. This causes gaps to occur in the data as it is stored on the physical disk. These gaps reduce the performance of the read and write operations.

There are two parts to any read or write operation:

- Data seeking
- Data reading or writing

In a read operation, *data seeking* refers to the process of locating the next piece of information on the drive that needs to be read into memory. *Data access* refers to actually reading the information off the drive. In a write operation, *seeking* refers to locating the next portion of free space on the drive to which data can be written, and *data writing* refers to actually committing the information to storage on the disk by creating tiny magnetized and unmagnetized areas that represent the ones and zeroes that make up binary data.

For each of these operations, an increase in gaps and a reduction in contiguous data or free space means that seek time is increased. If a particular piece of data must be written in small parts in different places across the disk, it takes longer to find all those parts and read them than if the data is in a row. The greater the seek time, the lower the performance of the database. The primary goal of defragmentation is to eliminate as much extra seek time as possible, thus restoring the database to maximum read and write performance. It is important to note that the online defragmentation process rearranges the data and free space so that they are contiguous and re-indexes the data to facilitate faster access. It does not, however, reduce the size of the database file and recover available free space from it.

The Offline Defragmentation Process

Although Windows Server 2003 runs an online defragmentation twice per day by default, there might be times when you need to actually recover available free space from the Active Directory database file. This can only be done by performing an offline defragmentation. This type of defragmentation process is much more invasive than its online counterpart, and should be done only when absolutely necessary. For example, you might have recently demoted the DC from acting as a Global Catalog (GC) server. In multidomain environments, GCs typically have much larger database files than the other DCs in their domain. This is because they must contain a subset of attributes from every object in the forest. When the GC role is removed from a server, this information is deleted from its Active Directory database file, leaving a large amount of free space that can be recovered.



EXAM WARNING

Be sure to remember that online defragmentation runs automatically but does not free disk space, and offline defragmentation, which must be performed explicitly, frees up disk space.

As with many invasive Active Directory database operations, you must be booted into a special mode known as the *Directory Services Restore Mode* to perform an offline defragmentation. This mode allows a DC to boot without initializing its copy of the Active Directory database. For most critical maintenance operations, the database must be offline. Because you must always log on to a Windows Server 2003 computer, a small version of a local directory service database (called the Security Accounts Manager, or SAM, database) remains on the computer after it has been promoted to a DC. This database has a single account, which is a local administrator account. This account is used when performing database maintenance operations in Directory Services Restore Mode.

You can access the Directory Services Restore Mode by booting or rebooting the computer, pressing the **F8** key when prompted, and selecting **Directory Services Restore Mode** from the Windows Advanced Options menu. You will be prompted to log on, and you must use the administrator account mentioned in the previous paragraph. Note that this account is not the domain administrator account; it is a special local account. The password for this local account is set during the installation of directory services on the local computer.

After you are authenticated by the local SAM, you can perform advanced directory services maintenance functions. Many of these are performed with the Ntdsutil utility, which can be used for many tasks, including the following:

- Offline defragmentation
- Database repair
- Moving the database

- Special types of directory services restore operations
- Changing operations' master roles
- Deleting DC entries from Active Directory



NOTE

The Ntdsutil tool is installed in the <systemroot>\System32 folder by default on Windows Server 2003 computers. In Windows 2000, Ntdsutil was not installed by default; you had to install it from the support\tools folder on the Windows 2000 installation CD-ROM.

To perform an offline defragmentation of the Active Directory database, follow the steps listed in Exercise 11.02.

EXERCISE 11.02

PERFORMING AN OFFLINE DEFRAGMENTATION OF THE ACTIVE DIRECTORY DATABASE

1. Back up the system state data for fault tolerance purposes. See the *Backing Up Active Directory* section later in this chapter for more information.
2. Boot or reboot the computer.
3. When prompted, press **F8** during Windows Server 2003 startup.
4. Select **Directory Services Restore Mode (Windows DCs only)** on the Windows Advanced Options menu that appears, and press the **Enter** key.
5. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.
6. You will see a number of checks performed while the system is booting, and you eventually will receive the Safe Mode logon prompt.
7. Log on by providing the password for the local administrator account and clicking the **OK** button.
8. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
9. Open a command prompt.

10. Type **ntdsutil** to enter the Ntdsutil utility. Note that this is a command-line utility, so the command prompt will change to ntdsutil:.
11. Type **files**. The command prompt should change to display file maintenance.
12. Type **compact to <drive>:\<directory>** to create a defragmented and compacted copy of the Active Directory database in the specified new location. For example, **compact to C:\ADTemp** creates a defragmented, re-indexed, and re-sized database file in the C:\ADTemp directory, as shown in Figure 11.5. The location specified can be on a local disk or on a mapped network drive. If there are spaces in the path where the file needs to be placed, it must be surrounded in quotes; for example, **"compact to c:\ad\july defrag"**.

Figure 11.5 The Ntdsutil Compact To Command

```

C:\WINDOWS\system32\cmd.exe - ntdsutil
C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: files
file maintenance: compact to C:\ADTemp
Opening database [Current].
Executing Command: C:\WINDOWS\system32\esentutl.exe /d"C:\WINDOWS\NTDS\ntds.dit"
/t"C:\ADTemp\ntds.dit" /p /o
Initiating DEFRAGMENTATION node...
Database: C:\WINDOWS\NTDS\ntds.dit
Temp. Database: C:\ADTemp\ntds.dit
Defragmentation Status (% complete)
0 10 20 30 40 50 60 70 80 90 100
|-----|-----|-----|-----|-----|-----|-----|-----|-----|
.....
Note:
It is recommended that you immediately perform a full backup
of this database. If you restore a backup made before the
defragmentation, the database will be rolled back to the state
it was in at the time of that backup.
Operation completed successfully in 7.922 seconds.
Spawned Process Exit code 0x0<0>
If compaction was successful you need to:
copy "C:\ADTemp\ntds.dit" "C:\WINDOWS\NTDS\ntds.dit"
and delete the old log files:
del C:\WINDOWS\NTDS\*.log
file maintenance: _

```

13. Type **quit** to return to the ntdsutil: prompt.
14. Type **quit** again to exit the utility.
15. Open Windows Explorer and rename the previously used **ntds.dit** file to **ntds.old.dit**.



NOTE

Step 15 is not specified in Microsoft's instructions, but we recommend it for fault tolerance purposes. As mentioned, an offline defragmentation is very invasive. It is possible that the compacted file will be corrupt and that Active Directory will not start after the procedure. If you don't take this step, you will be forced to do a system state restore to recover the previous database file. By simply renaming the file, you can boot back into Directory Services Restore Mode, delete the corrupt file, and rename `ntds.old.dit` back to `ntds.dit` to recover the system.

16. In Windows Explorer, copy the new **ntds.dit** file from the location you specified, using the **compact** command to specify the location of the primary **ntds.dit** file location.
17. In Windows Explorer, delete all files that end with the `.LOG` extension in your Active Directory log files folder.
18. Close the command prompt window and reboot the server normally.

Determining the Amount of Free Space to Be Recovered

As mentioned, you should perform an offline defragmentation only when you are confident that you can recover significant free space. Offline defragmentation is a very serious process that requires taking the DC offline for a period of time and performing a highly invasive database procedure. You might be wondering how you can tell if there is enough free space available for recovery to justify the operation.

Fortunately, Microsoft provides you with a way to determine this that doesn't require actually performing the procedure. By modifying the Registry, you can have the Directory Service generate an event log entry that estimates how much space would be freed up by an offline defragmentation. The following steps outline the process:

1. Click **Start** | **Run**.
2. In the **Open:** text box, type **regedt32.exe** and click the **OK** button.
3. In the Registry editor, navigate to `HKEY_LOCAL_MACHINE \SYSTEM \CurrentControlSet \Services \NTDS \Diagnostics`.
4. Double-click the **6 Garbage Collection** entry.
5. In the Edit DWORD Value dialog box, type **1** in the **Value Data** text box.

Continued

6. Click **OK**.
7. Verify that the value in the Data column for **6 Garbage Collection** is **0x00000001 (1)**. If it is not, repeat steps 4, 5, and 6.
8. Close the Registry editor.

Moving the Database or Log Files

At some point, it might become necessary to move the Active Directory database or log files. Most often, this occurs because you need to move the files to a new hard drive or array of hard drives. Performance might dictate this decision. New faster drives and controllers can be used to replace slower ones as usage increases. The decision might also be dictated by a lack of free space as the Active Directory database and components grow.

Moving the database or log files is relatively simple. It is done from the command line using the Ntdsutil utility discussed previously. Because the database and log files cannot be open when they are moved, the operation must be carried out while in Directory Services Restore Mode. The following steps outline how to move the Active Directory database and log files.

1. Back up the system state data for fault tolerance purposes. See the *Backing Up Active Directory* section later in this chapter for more information.
2. Boot or reboot the computer.
3. When prompted, press **F8** during Windows Server 2003 startup.
4. Select **Directory Services Restore Mode (Windows DCs only)** on the Windows Advanced Options menu that appears, and press the **Enter** key.
5. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.
6. You will see a number of checks performed while the system is booting, and eventually you will receive the Safe Mode logon prompt.
7. Log on by providing the password for the local administrator account and clicking the **OK** button.
8. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
9. Open a command prompt.
10. Type **ntdsutil** to enter the Ntdsutil utility. This is a command-line utility so the command prompt will change to ntdsutil:.
11. Type **files**. The command prompt should change to display file maintenance:.

12. Use one of the following commands to move the Active Directory database or log files, or update their paths.

- Type **move DB to <drive>:\<directory>** to move the ntds.dit database file to the new location specified. For example, **move DB to C:\AD** moves the database file to the C:\AD directory and updates the Registry to point to this new location, as shown in Figures 11.6 and 11.7.

Figure 11.6 Moving the Active Directory Database with Ntdsutil, First Screen Portion

```
C:\WINDOWS\system32\cmd.exe - ntdsutil
C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: files
file maintenance: move DB to C:\AD
Opening database [Current].

C:\Documents and Settings\Administrator>REM - *****
*****
C:\Documents and Settings\Administrator>REM - Script to move DS DB file
*****
C:\Documents and Settings\Administrator>REM - *****
*****

C:\Documents and Settings\Administrator>C:
C:\Documents and Settings\Administrator>cd \
C:\>mkdir "AD"
A subdirectory or file AD already exists.
C:\>cd "AD"
C:\AD>move "C:\WINDOWS\NTDS\ntds.dit" "C:\AD\ntds.dit"
1 file(s) moved.
C:\AD>C:\WINDOWS\system32\ntdsutil.exe files "set path DB \"C:\AD\ntds.dit\""
quit
C:\WINDOWS\system32\ntdsutil.exe: files
file maintenance: set path DB "C:\AD\ntds.dit"
Copying NTFS security from C:\WINDOWS\NTDS to C:\AD...
file maintenance: quit
C:\WINDOWS\system32\ntdsutil.exe: quit

C:\AD>C:\WINDOWS\system32\ntdsutil.exe files "set path backup \"C:\AD\DSADATA.BAK\""
quit quit
C:\WINDOWS\system32\ntdsutil.exe: files
file maintenance: set path backup "C:\AD\DSADATA.BAK"
The previous NTDS database location C:\WINDOWS\NTDS\dsadata.bak is unavailable.
Default NTFS security on NTDS folders will be set on reboot.
file maintenance: quit
C:\WINDOWS\system32\ntdsutil.exe: quit

C:\AD>C:\WINDOWS\system32\ntdsutil.exe files "set path working dir \"C:\AD\""
quit
```

Figure 11.7 Moving the Active Directory Database with Ntdsutil, Second Portion of the Screen

```
C:\WINDOWS\system32\cmd.exe - ntdsutil
C:\AD>C:\WINDOWS\system32\ntdsutil.exe files "set path working dir \"C:\AD\""
quit
C:\WINDOWS\system32\ntdsutil.exe: files
file maintenance: set path working dir "C:\AD"
Copying NTFS security from C:\WINDOWS\NTDS to C:\AD...
file maintenance: quit
C:\WINDOWS\system32\ntdsutil.exe: quit

C:\AD>C:\WINDOWS\system32\ntdsutil.exe files info quit quit
C:\WINDOWS\system32\ntdsutil.exe: files
file maintenance: info

Drive Information:
C:\ NTFS (Fixed Drive ) free(15.7 Gb) total(18.6 Gb)

DS Path Information:
Database : C:\AD\ntds.dit - 10.1 Mb
Backup dir : C:\AD\DSADATA.BAK
Working dir : C:\AD
Log dir : C:\WINDOWS\NTDS - 40.8 Mb total
res2.log - 10.0 Mb
res1.log - 10.0 Mb
ntds.INTEG.BAK - 91.2 Kb
edb00002.log - 10.0 Mb
edb.log - 10.0 Mb

file maintenance: quit
C:\WINDOWS\system32\ntdsutil.exe: quit

C:\AD>REM - *****
*****
C:\AD>REM - Please make a backup immediately else restore
*****
C:\AD>REM - will not retain the new file location.
*****
C:\AD>REM - *****
*****
Opening database [Current].
If new database was successful,
please make a backup immediately else restore
will not retain the new file location.

file maintenance:
```

- Type **move logs to <drive>:\<directory>** to move the Active Directory log files to the new location specified. For example, **move logs to C:\AD** moves the log files to the C:\AD directory and updates the Registry to point to this new location, as shown in Figures 11.8 and 11.9.

Figure 11.8 Moving the Active Directory Log Files with Ntdsutil, First Portion of the Screen

```

C:\WINDOWS\system32\cmd.exe - ntdsutil
ntdsutil: files
file maintenance: move logs to C:\AD
Opening database [Current].

C:\Documents and Settings\Administrator>REM - *****
*****
C:\Documents and Settings\Administrator>REM - Script to move DS log files
C:\Documents and Settings\Administrator>REM - *****
*****
C:\Documents and Settings\Administrator>C:
C:\Documents and Settings\Administrator>cd \
C:\>mkdir "AD"
A subdirectory or file AD already exists.
C:\>cd "AD"
C:\AD>move "C:\WINDOWS\NTDS\res2.log" "C:\AD\res2.log"
1 file(s) moved.
C:\AD>move "C:\WINDOWS\NTDS\res1.log" "C:\AD\res1.log"
1 file(s) moved.
C:\AD>move "C:\WINDOWS\NTDS\ntds.INTEG.RAW" "C:\AD\ntds.INTEG.RAW"
1 file(s) moved.
C:\AD>move "C:\WINDOWS\NTDS\edb00002.log" "C:\AD\edb00002.log"
1 file(s) moved.
C:\AD>move "C:\WINDOWS\NTDS\edb.log" "C:\AD\edb.log"
1 file(s) moved.
C:\AD>C:\WINDOWS\system32\ntdsutil.exe files "set path logs \"C:\AD\""" quit quit
C:\WINDOWS\system32\ntdsutil.exe: files
file maintenance: set path logs "C:\AD"

```

Figure 11.9 Moving the Active Directory Log Files with Ntdsutil, Second Portion of the Screen

```

C:\WINDOWS\system32\cmd.exe - ntdsutil
file maintenance: set path logs "C:\AD"
Copying NTFS security from C:\WINDOWS\NTDS to C:\AD...
file maintenance: quit
C:\WINDOWS\system32\ntdsutil.exe: quit
C:\AD>C:\WINDOWS\system32\ntdsutil.exe files info quit quit
C:\WINDOWS\system32\ntdsutil.exe: files
file maintenance: info

Drive Information:
    C:\ NTFS <Fixed Drive > free<15.6 Gb> total<18.6 Gb>

DS Path Information:
    Database : C:\AD\ntds.dit - 10.1 Mb
    Backup dir : C:\AD\DSADATA.BAK
    Working dir : C:\AD
    Log dir : C:\AD - 40.8 Mb total
              res2.log - 10.0 Mb
              res1.log - 10.0 Mb
              ntds.INTEG.RAW - 91.2 Kb
              edb00002.log - 10.0 Mb
              edb.log - 10.0 Mb

file maintenance: quit
C:\WINDOWS\system32\ntdsutil.exe: quit
C:\AD>REM - *****
*****
C:\AD>REM - Please make a backup immediately else restore
C:\AD>REM - will not retain the new file location.
C:\AD>REM - *****
*****
Opening database [Current].
If move log files was successful,
please make a backup immediately else restore
will not retain the new file location.

file maintenance: _

```


13. Ensure that the database (or log files) is now referenced in the proper location by typing **info**.
14. Type **quit** to return to the ntdsutil: prompt.
15. Type **quit** again to exit the utility.
16. Close the command prompt window and reboot the server normally.

Using the Set Path Option in Ntdsutil

It's important to properly move the Active Directory database and log files using the Ntdsutil command-line utility. This updates the Registry entries that point to the correct locations in the file system, thus allowing the system to find and initialize them when booting. If you are forced to restore these files to another location, or simply copy them to a new location using Windows Explorer, Active Directory will not initialize when the system is rebooted. Fortunately, Microsoft provides a way to fix this, using the Ntdsutil utility. To do so, follow these steps:

1. Boot or reboot the computer.
2. When prompted, press **F8** during Windows Server 2003 startup.
3. Select **Directory Services Restore Mode (Windows DCs only)** on the Windows Advanced Options menu that appears, and press the **Enter** key.
4. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.
5. You will see a number of checks performed while the system is booting, and you eventually will receive the Safe Mode logon prompt.
6. Log on by providing the password for the local administrator account and clicking the **OK** button.
7. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
8. Open a command prompt.
9. Type **ntdsutil** to enter the Ntdsutil utility. This is a command-line utility, so the command prompt will change to ntdsutil:.
10. Type **files**. The command prompt should change to display file maintenance:.
11. Use one of the following commands to move the Active Directory database or log files, or to update their paths.
12. Type **set path DB <drive>:\<directory>** to update the Registry to point to the new location of the ntds.dit file.

Continued

13. Type **set path logs <drive>:\<directory>** to update the Registry to point to the new location of the Active Directory log files.
14. Type **quit** to return to the ntdsutil: prompt.
15. Type **quit** again to exit the utility.
16. Close the command prompt window and reboot the server normally.

EXAM
70-294
OBJECTIVE
2.5.3

Monitoring the Database

It's important to implement a consistent Active Directory monitoring strategy to ensure database integrity, reliability, and performance within your forest. Regular monitoring can also improve your knowledge of Active Directory and assist you in determining when a problem is in the early stages of unfolding. This will lead to performance and service issues being resolved in a much more timely manner.

After you have been monitoring your environment for a period of time, you will become aware of the performance, reliability, and service issues that most often occur in your forest. Preparing responses in advance and training your help desk and administrative personnel can go a long way toward ensuring that these issues can be dealt with routinely and with maximum effectiveness. Scripted responses can also be used to this end.

It is important to remember that Active Directory does not exist in a vacuum. Performance and service issues that seem to relate to Active Directory can be caused by other key infrastructure components, such as name resolution services. Because of this, any Active Directory monitoring solution should also consist of key metrics from all the network services that play a role in supporting the directory service. In a large environment, these supporting components might be under the authority of another administrative group and you will need to enlist the cooperation of that group's members when defining your monitoring strategy. In the following sections, we look at some of the primary tools that you can use to monitor Active Directory.

Using Event Viewer to Monitor Active Directory


You can use the Windows Server 2003 Event Viewer tool, shown in Figure 11.10 and accessed via **Start | Programs | Administrative Tools | Event Viewer**, to view a variety of event logs on the DC. For monitoring the directory service, the following event logs are of particular interest:

- **DNS Server** This event log displays information relating to the DNS server service if it is installed on the DC. It is common in small environments and remote offices to have a single server acting as both DC and DNS server (along with other roles). Because clients running Windows 2000 and later operating systems

use DNS to locate DCs and GC servers, problems with this service can severely impact Active Directory availability on the network.

- **System** This event log displays critical information concerning the state of the operating system as a whole. Examining the System log should be part of the review procedure because underlying system stability is critical to optimum functionality of a DC. The System event log is also used to display messages that notify you when the DC's DNS record was not registered or updated properly.
- **Application** This event log displays extensive information from Group Policy and other related Active Directory components.
- **Directory Service** This is the primary event log for Active Directory. It includes information related to when the directory service starts and stops, the Garbage Collection process, online defragmentation, and much more.
- **File Replication Service** This service controls the replication of SYSVOL, which contains critical data such as Group Policy and replication topology connection information.

Figure 11.10 The Windows Server 2003 Event Viewer Utility



In a large domain, event logs can grow quite rapidly, which makes it difficult to search through them for key events. Microsoft recommends using the filter or search functionality to specifically seek out events matching the following criteria:

- All records with an Error severity level in the Directory Service or FRS event logs.
- All LSASS records in the System event log with a severity level of Error. The Local Security Authority subsystem (LSASS) is the primary security subsystem for Active Directory.

- All Kerberos V5 Key Distribution Center (KDC) records in the System event log with a severity level of Error. The KDC is the primary logon service for Windows 2000 and later clients in Active Directory.
- All USERENV records in the Application event log with a severity level of Error. This setting can indicate problems with the application of Group Policy.



NOTE

Administrators can set permissions on event logs to restrict access via Group Policy.

Head of the Class...

Understanding Event Log Entries

Most event logs (except the Security log) display three different types or levels of events:

- **Error events** indicate that there is a relatively serious problem that can result in a loss of functionality or loss of data.
- **Warning events** indicate events that can be indicative of potential future problems.
- **Information events** describe operations that have occurred that do not involve problems, such as the successful loading of a driver or starting of a service.

The Error event is the most severe, and demands action on the part of the administrator to prevent further problems. The error type is represented by a symbol in the log; an Information event is represented by a balloon with a blue “i” in it, a Warning event is represented by a yellow triangle with a black exclamation point inside it, and an Error event is represented by a red circle with a white “X” inside it.

Using the Performance Console to Monitor Active Directory

The Performance console is another tool that comes preinstalled in Windows Server 2003 and can be very helpful in monitoring Active Directory. The Performance console, shown in Figure 11.11 and accessed via **Start | Programs | Administrative Tools | Performance**, is capable of monitoring the server on which it is installed, and other remote servers. Data from any number of Windows Server 2003 computers can be com-

bined for tracking or display purposes. Windows contains a variety of performance metrics that can be monitored with this utility. The Performance utility consists of the three following components:

- **System Monitor** This portion of the utility is used to graphically display performance metrics. When System Monitor is selected in the left pane of the utility, the right pane fills with a large graph (by default) that displays three recommended general system counters (Memory:Pages/sec, PhysicalDisk:Avg disk queue length, and Processor:%Processor time) by default. The lower portion of the graph pane allows you to select between the counters, and numerical statistics for the selected counter are displayed between the graph and the selection box (see Figure 11.11). You can also change the view from graph to Histogram (bar chart) or Report (text).
- **Counter and Trace logs** These options allow for detailed levels of logging over time. In most cases, you won't have time to sit around all day watching the System Monitor, which charts real-time information. Instead, you'll need data that you can review and work with when it is convenient. In addition, keeping an archive of consistently logged data is essential to becoming familiar with system norms and spotting trends over time.
- **Alerts** The Alerts option allows you to specify critical thresholds that, when exceeded, cause some type of action to take place. The default action is to have the alert generate a message in Event Viewer. You can also have it send a network message, start logging to a preconfigured counter log, or execute a script, batch file, or program. A good example of an alert that can be configured for Active Directory is one that monitors the amount of free space on the volume that contains the Active Directory database and log files. You can have the system alert you when the amount of free space goes under a specified limit, to ensure that you don't run out of disk space on this volume.

The metrics that are monitored using the Performance console are called *counters*. These counters are grouped according to the *objects* to which they pertain. For example, you can monitor many different aspects of what is happening to memory by using the counters that pertain to the Memory object. The Memory:Pages/sec counter mentioned earlier is the Pages/sec counter within the Memory object. The notation used is a common way to see these counters described in print (including on Microsoft exams)—it also makes them easy to find. The first step is to select the Memory object. When selected, it will display a number of different memory-related counters that can be selected. Scroll down the list and pick the one that you are looking for. Figure 11.12 shows this dialog box and some of the counters available for the NTDS (directory services) object.

Figure 11.11 The Windows Server 2003 Performance Console Utility

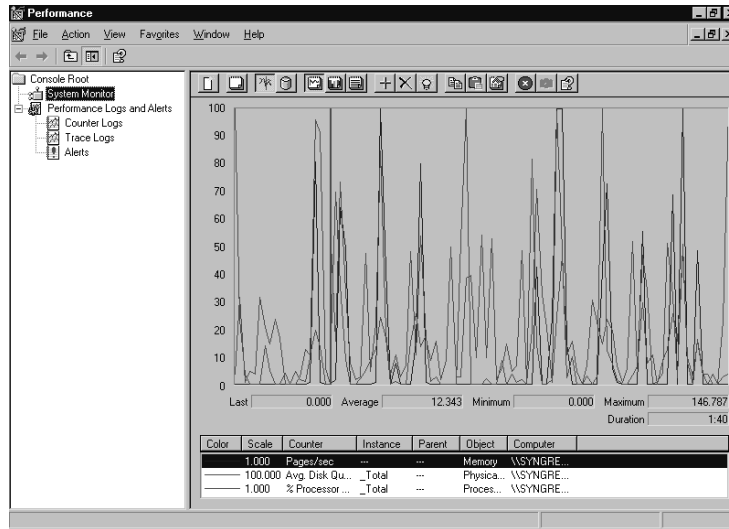
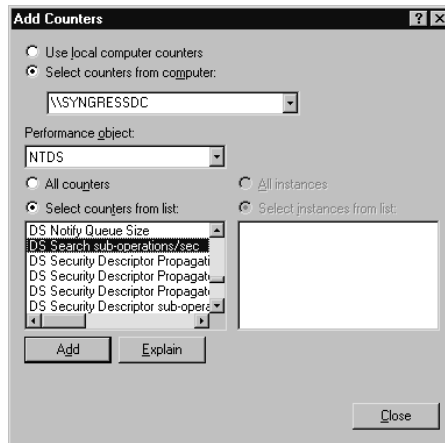


Figure 11.12 The Add Counters Dialog Box



Microsoft recommends that you use the following performance counters for monitoring Active Directory:

- **NTDS performance object counters:**
 - **DRA inbound and outbound counters** These counters are used to track the amount of replication information that flows into and out of a site. Significant changes can indicate a major increase in the amount of replication traffic or a shift in the site replication topology.

- **DS Search sub-operations/sec** Significant changes in this counter can indicate an application that is incorrectly targeting a DC, or performance problems involving the DC.
- **LDAP Searches/sec** This counter corresponds to the overall number of LDAP searches per second on the DC. It should be relatively consistent across all of your DCs in a well-planned and balanced environment. If it isn't, this counter can indicate that an application is incorrectly targeting a DC (rather than spreading its use out across several DCs). It can also indicate uneven client loads. This counter is also useful for tracking trends over time for capacity planning.
- **LDAP Client Sessions** This counter displays the number of clients that are connected to the LDAP services. It can also be used to track uneven client loads, which might be indicative of connection failures to other DCs in a well-planned and balanced environment. Like the LDAP Searches/sec counter, this counter is useful for tracking trends over time for capacity planning.
- **NTLM Authentications** This counter indicates the number of domain authentications taking place using the NTLM protocol. Windows 2000 and later clients should use Kerberos for authentication, but will fail back to NTLM when they are unable to authenticate using Kerberos. This counter can be used to indicate Kerberos authentication issues in these types of environments.
- **Kerberos Authentications** This counter indicates the number of domain authentications that take place using the Kerberos protocol. This counter is helpful in tracking authentication trends over time for capacity planning.
- **Processor object counters:**
 - **% Processor Time** This counter can be used to track the overall consumption of processor resources in the DC. Microsoft recommends that this counter not exceed 85 percent on a sustained basis.
 - **% DPC Time** This counter alerts you to delayed execution of processes resulting from the DC being too busy to execute them. Microsoft recommends a sustained threshold of 10 for this counter.
- **System object counters:**
 - **Processor Queue Length** This counter indicates that the system cannot keep up with processing requests. When you see the word *queue* in any counter, the counter tracks the number of things “waiting in line” to use the resource. Microsoft recommends that this counter not exceed a value of 2 on a sustained basis.
 - **Context Switches/sec** Most modern processors can only execute one thread at a time. Although it appears that the computer is running many pro-

grams at once, each program is actually sharing the processor with all others. Each *thread* (the smallest unit of executable code in a program) uses the processor for a short period of time and then passes it on to the next. This concept is referred to as *time slicing*. A context switch occurs when the processor switches between waiting processes. This counter can indicate too many applications (including operating system applications) for the processor to service, or applications that are too busy for the processor to keep up with. Microsoft recommends that this counter not exceed 70,000 on a sustained basis.

■ **Memory object counters:**

- **Page Faults/sec** This counter indicates when needed program code is not resident in memory and must be loaded from disk (from the page file). This is often an indication of a system in need of more physical RAM. Microsoft recommends a sustained threshold of 700 for this counter.
- **Available MBytes** This counter indicates the amount of available system memory. Microsoft recommends using this counter to configure an alert that will notify you when the DC is running low on memory resources.
- **PhysicalDisk:Current Disk Queue Length counter** This counter can be used to track the number of disk reads and writes that are waiting to be filled. This can be the result of a busy processor that is not able to keep up with IRQ requests, or a slow disk drive or subsystem. Microsoft recommends that this counter not exceed a value of 2 on a sustained base.



TEST DAY TIP

While the NTDS performance object is the primary one for Active Directory, be sure to remember that it's important to monitor other objects, too, as this can give you a larger picture of what is going on with the server and the services on which Active Directory depends.

In Exercise 11.03, you will learn how to configure System Monitor to display the key NTDS object counters.

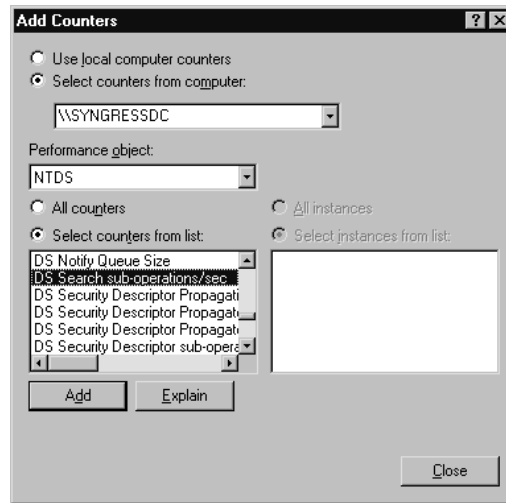
EXERCISE 11.03

USING SYSTEM MONITOR TO MONITOR ACTIVE DIRECTORY

1. Open the Windows Server 2003 Performance console from **Start | Programs | Administrative Tools | Performance**.
2. Select the **System Monitor** node in the left pane.

3. In the right pane, right-click on the graph and select **Add Counters...** in the context menu that appears. Note that you can also click the + button on the toolbar above the graph to add a new counter.
4. In the **Add Counters** dialog box, select **NTDS** from the **Performance object:** drop-down box.
5. In the **Select counters from list:** box, select the **DS Search sub-operations/sec** counter as shown in Figure 11.13.

Figure 11.13 The Add Counters Dialog Box



6. Click the **Add** button.
7. Repeating steps 5 and 6 after each addition, add the following counters: **LDAP Searches/sec**, **LDAP Client Sessions**, **NTLM Authentications**, and **Kerberos Authentications**.
8. Click the **Close** button in the Add Counters dialog box.
9. Your new counters should appear in the list in the right pane under the graph, as shown in Figure 11.14. Select one of these counters by clicking on it in this list.
10. Press **Ctrl + H** to highlight the counter in the graph as shown in Figure 11.15. This tool is often used to display a large number of counters, making it very difficult to tell them apart in the graph. Using the highlight feature makes this much easier.
11. Close the Performance console.

Figure 11.14 The Performance Console without Highlighting Enabled

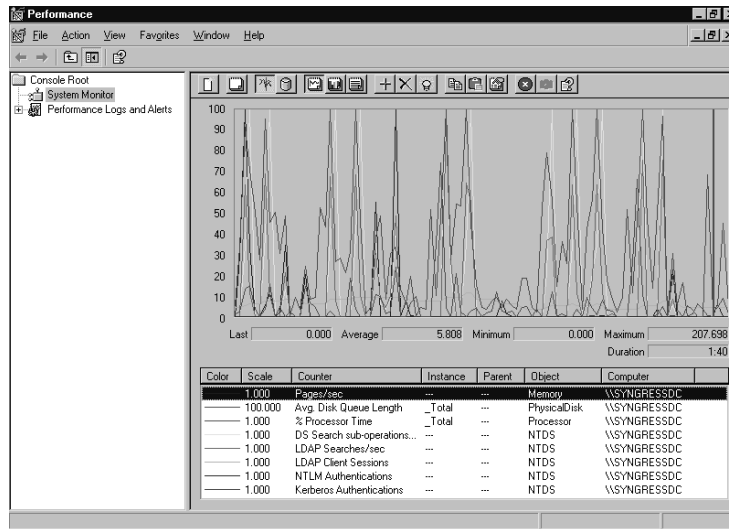
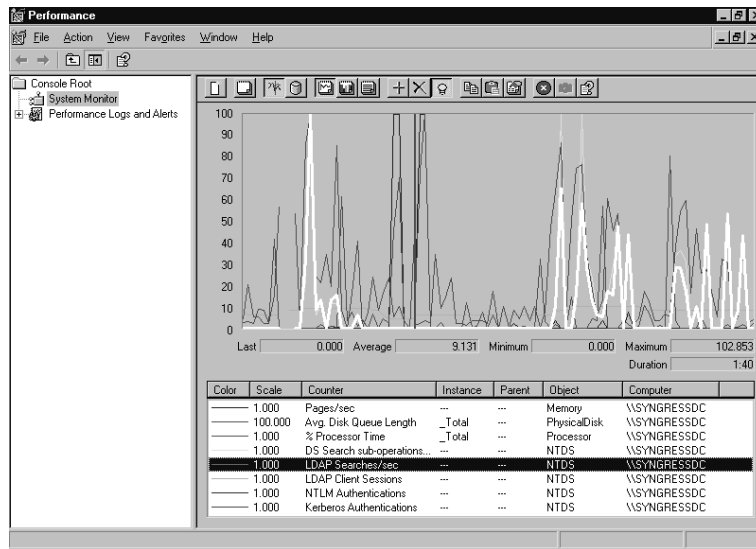


Figure 11.15 The Performance Console with Highlighting Enabled



**NOTE**

Running the Performance utility can have a negative impact on the performance of the server. You can decrease this impact by running the System Monitor in Report view instead of the default graph view, sample at longer intervals (at least three seconds apart), and select as few objects and counters to monitor as possible.

Backing Up and Restoring Active Directory

Although it's technically just a collection of files, Active Directory has its own unique backup and restore methods. That's because it is a *very important* collection of files, vital to the operation of your DC and network operations. In this section, you'll learn how to back up Active Directory from both the Windows Server 2003 Backup utility and from the command line. In Windows Server 2003, Microsoft gives you the ability to back up open files using the Volume Shadow Copy feature. In previous versions of Windows, you had to purchase a third-party backup utility to back up open files. After you learn how to back up Active Directory, you'll learn how to restore it to the DC. In the following subsections, we discuss the three different types of restore operations available, and when to use each.

Backing Up Active Directory

Several different methods can be used to back up Active Directory:

- As part of a full system backup
- As part of a partial system backup
- Back up the system state data only

In the past, the Active Directory database had to be backed up as system state data. As previously mentioned, this was because when the operating system was in use, the Active Directory files were open, locked and inaccessible to the backup process. Microsoft Backup and some other third-party backup programs are now able to use a new Windows Server 2003 feature known as the Volume Shadow Copy service to work around this open file issue.

Volume Shadow Copy makes a read-only copy of the information in these open files, which can be used for backup purposes. The original files continue to be accessed without any interference from the backup operation. When the backup is complete, the Volume Shadow Copy is deleted. The amount of disk space required by the Volume Shadow Copy will vary, based on the amount of data that changes on the disk during the backup procedure. If the underlying disk does not have enough free space to support Volume Shadow Copy, open files are not backed up.



TEST DAY TIP

Volume Shadow Copy is new in Windows Server 2003, and is used to allow open files to be backed up. Another form of shadow copying supported in Windows Server 2003 is Shadow Copies of Shared Folders, which is designed to provide point-in-time file copies for user recovery purposes. Both use the same technology; however, their uses differ.

When preparing a backup job, rather than specifying the individual files for Active Directory to be backed up, it is best to always use the *system state data* selection in the utility. System state will be backed up automatically when a full system backup is selected, and can be specified manually when a partial backup is selected. Using the system state backup feature ensures that all necessary files are backed up. Relying on your own knowledge of which files would be needed in the case of a restore is not recommended. Remember that Active Directory consists not only of the core Active Directory database and log files, but also Group Policy, replication topology, and other information contained in the SYSVOL folder.

In the next section, we examine how to back up Active Directory using the Windows Server 2003 Backup utility.



EXAM WARNING

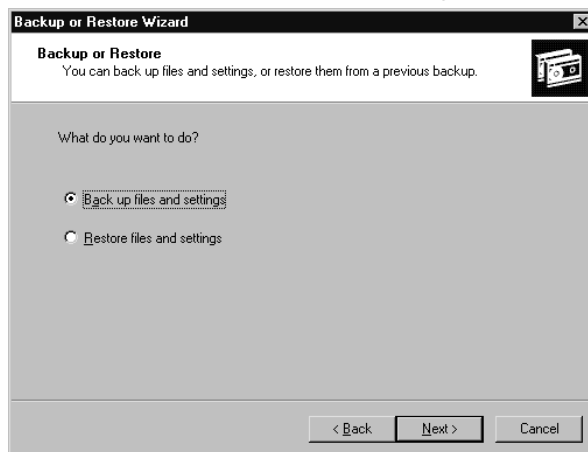
Be sure to remember that you can only back up and restore system state data on the local machine using the Backup utility.

Using the Windows Server 2003 Backup Utility

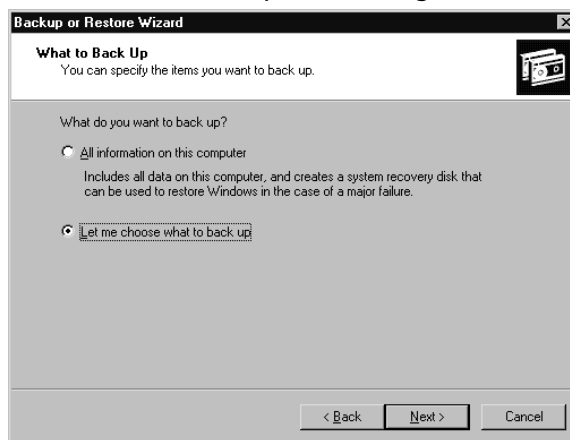
To perform a backup, your user account must be a member of the Administrators or Backup Operators group on the local computer, or have the **Back up files and directories** user right. The Windows Server 2003 Backup utility can be opened from **Start | All Programs | Accessories | System Tools | Backup** or by clicking **Start | Run** and typing **ntbackup**.

When opened, the utility displays a brief message stating that it is looking for local backup devices, and then opens the Backup or Restore Wizard. If you prefer not to use the wizard, you can click the **Advanced Mode** link. If you want to ensure that the utility always starts in advanced mode, you can clear the check box next to **Always start in wizard mode**. To use the wizard, follow these steps:

1. On the initial page of the wizard, click the **Next** button.
2. The next page in the wizard, shown in Figure 11.16, asks you if you want to perform a backup or a restore. Leave the default option button selection next to **Back up files and settings**, and click the **Next** button.

Figure 11.16 The Backup or Restore Wizard Page

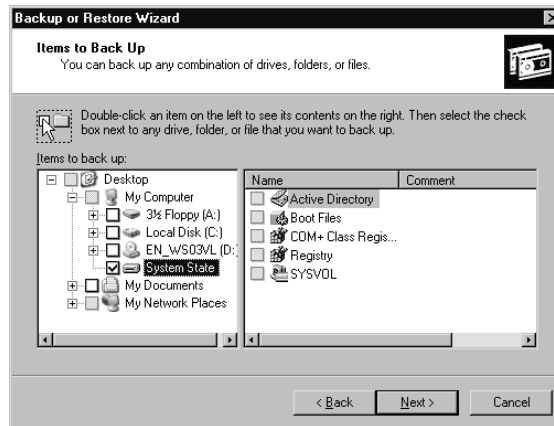
3. When asked what you want to back up, as shown in Figure 11.17, you have two choices:
 - **All information on this computer**, which backs up all volumes and includes system state data.
 - **Let me choose what to backup**, which allows you to select individual files and folders to back up, in addition to selecting or deselecting the system state data. Because this is the more complex setting that allows you to see more of the wizard, select this option and click the **Next** button.

Figure 11.17 The What to Back Up Wizard Page

4. The **Items to Back Up** screen, shown in Figure 11.18, presents an Explorer type interface that allows you to select which files, folders, and other components (such

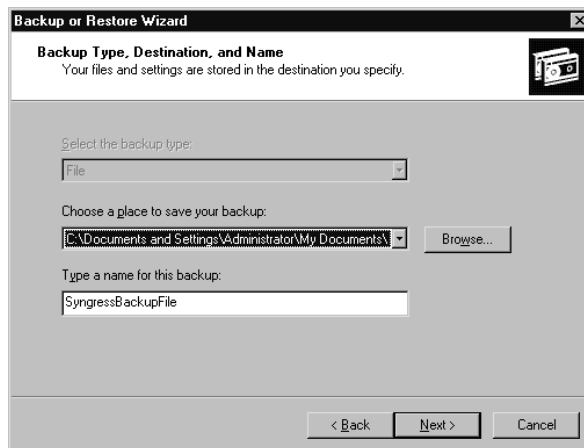
as system state data) you want to back up. In the left pane are selections for **My Computer**, **My Documents**, and **My Network Places**. These can be expanded by clicking on the plus sign next to each to see additional items that can be backed up. Often, these items also have plus signs next to them and can be expanded to view even more files and folders that can be selected for individual backup. These levels and check boxes can be used to determine the granularity of the backup. As an example, a selection box next to **My Documents** will back up all subfolders and files contained in the My Documents folder. If you only want to back up system state data, you should expand the plus sign next to **My Computer** and check the box next to **System State**. When you have made your selections, click the **Next** button.

Figure 11.18 The Items to Back Up Wizard Page



5. Next, choose the location to which the information will be backed up, as shown in Figure 11.19. If a valid backup device (a tape device) was detected while the Backup utility was opening, you will be able to select it in the drop-down list under **Select the backup type:**. If no device was detected, this box will be grayed out and you can only back up to a file on a local or mapped network drive. When you back up to a file on the disk, you can later copy the backup to other media, such as CD-ROM or DVD, if you choose to do so. In the **Choose a place to save your backup:** drop-down box, you can do one of the following:
 - Type a path into the box
 - Select a previously used path from the drop-down list
 - Click the **Browse** button and select a path and filename using a standard **Save As** Windows dialog box

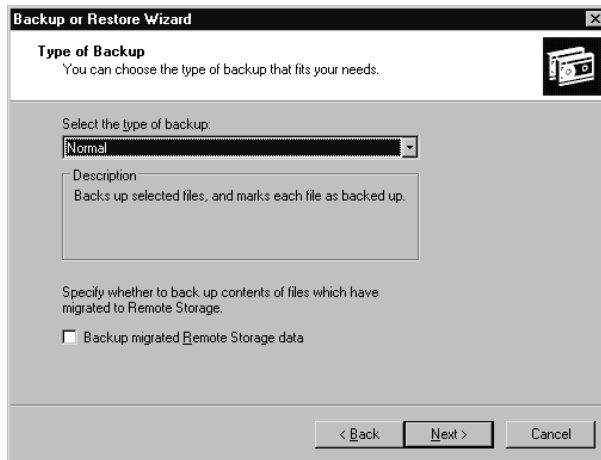
Figure 11.19 The Backup Type, Destination, and Name Wizard Page



6. At the bottom of the screen, enter the name of the backup file you want to use in the **Type a name for this backup:** text box. When you have completed the information on this screen of the wizard, click the **Next** button.
7. At this point, you can click the **Finish** button to start the backup immediately, using the default advanced settings. We want to make sure that you are familiar with all of the settings that this wizard offers, so instead click the **Advanced...** button to proceed. Note that if you chose a full backup, you will not be able to configure the advanced options; you will simply be prompted to finish the wizard and run the backup job.
8. The **Type of Backup** page, shown in Figure 11.20, allows you to specify the type of backup you want to use. The following options are available from the **Select the type of backup:** drop-down box:
 - **Normal** Specifies that all selected files, folders, and components will be backed up. This option resets the archive attribute. The archive attribute is set when a file, folder, or component has changed. If not cleared upon backup, the system cannot tell when the file has changed and needs to be backed up again. The archive attribute is read when Incremental and Differential backup operations are performed.
 - **Copy** Specifies that all selected files, folders, and components will be backed up. This option does *not* reset the archive attribute.
 - **Incremental** Uses the archive attribute to determine which files have changed since the last backup. Only the changed files are backed up. An Incremental backup clears the archive attribute so that subsequent backup operations can determine which files have changed. Clearing this attribute enables sequential backups to back up only those files that have changed since the last backup that ran.

- **Differential** Uses the archive attribute to determine which files have changed since the last backup. Only the changed files are backed up. A Differential backup does *not* clear the archive attribute. This means that subsequent backup operations back up files that have changed since the last backup ran, and other files that changed but were backed up by earlier backup operations.
- **Daily** This option reads the timestamps on files and only backs up files that were created or modified on the day of the backup. This option does *not* clear the archive attribute.

Figure 11.20 The Type of Backup Wizard Page



9. The **Backup migrated Remote Storage data** check box is located at the bottom of the **Type of Backup** page. Infrequently used files can be migrated to a near-time access point, using Remote Storage. When this occurs, they still show up for users as local files on the system in Explorer type interfaces, although they are stored remotely. Users are actually viewing reparse points, not the actual files. Because of this, it is possible that you selected files in step 4 that are not actually located on the local system disks, but have been migrated to Remote Storage. Files migrated to Remote Storage can be recalled seamlessly when the reparse point is clicked on by a user in an Explorer type interface. Checking this box ensures that these reparse points will be backed up. After completing your selections on this page of the wizard, click the **Next** button.
10. The next page in the wizard, shown in Figure 11.21, contains the following three check boxes:
 - **Verify data after backup** This reads the data back off the storage medium used and compares it to the original information backed up. You should be

aware that this will greatly extend the amount of time required to finish the backup job. However, when the data is critical and you need to be assured that it was backed up correctly, you might want to select this option.

- **Use hardware compression, if available** If the Backup utility detects a tape drive or other storage mechanism that is capable of hardware compression, this box will be available for selection. Typically, these types of compression are very advanced and it is recommended that you make use of them. This box will be grayed out if backup does not detect a device that supports this setting.
- **Disable Volume Shadow Copy** As mentioned earlier, this feature is used to back up open files. This option is enabled by default. If you select to back up the system state data, the option to disable it will be grayed out because it is required for backing up the system state information.

Figure 11.21 The How to Back Up Wizard Page



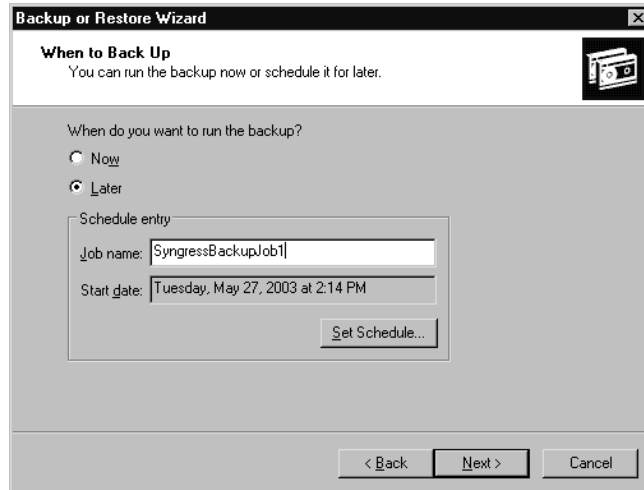
11. After making your selections on the **How to Back Up** wizard page, click the **Next** button.
12. The **Backup Options** page, shown in Figure 11.22, allows you to choose to append this backup to an existing backup by selecting the **Append this backup to the existing backups** option, or replace any existing backups on the media selected by choosing the **Replace the existing backups** option. If **replace the existing backups** is selected, the **Allow only the owner or the Administrator access to the backup data and to any backups appended to this medium** check box becomes available for selection. When checked, this allows only the user who created the backup file or an administrator to restore the backed up information. Click the **Next** button to continue with the wizard.

Figure 11.22 The Backup Options Wizard Page

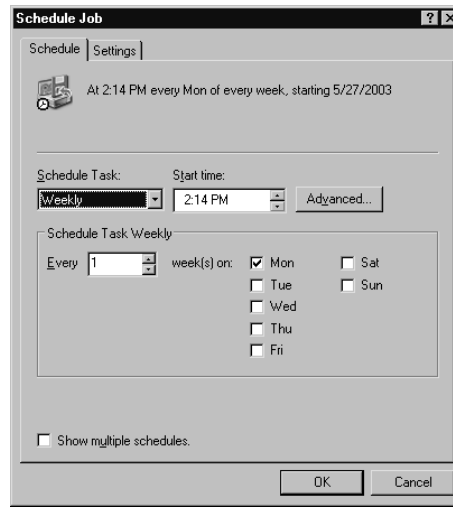


13. The Backup utility allows you to begin the backup immediately by selecting the **Now** option on the **When to Back Up** page, shown in Figure 11.23. However, you can also schedule a backup job to run at another time by selecting the **Later** option. When this option is selected, the **Job name:** text box becomes available, as does the **Start date:** option. Enter a descriptive name for the backup operation in the **Job name:** box.

Figure 11.23 The When to Back Up Wizard Page



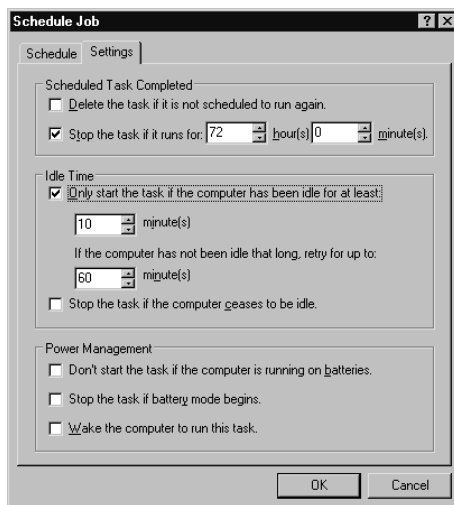
14. By default, the **Start date:** option is set to the current date and time when you click the **Later** option. To change this to another date and time, or use the more advanced schedule features, click the **Set Schedule...** button.
15. This displays the **Schedule Job** dialog box with the **Schedule** tab in the foreground, as shown in Figure 11.24. Several options can be selected in the **Schedule Task:** drop-down box, including:
 - **Daily** This setting allows you to specify a start time, and the number of consecutive days on which you would like the task to run. It also allows you to click the **Advanced...** button to bring up the **Advanced Schedule Options** dialog box. This box allows you to specify start and end dates, how often the task will repeat, and the maximum duration or time past which the backup job cannot run.
 - **Weekly** This option allows you to specify a start time, and you can click the **Advanced...** button to configure all of the advanced options listed previously. In addition, it allows you to specify the number of consecutive weeks the backup should run, and has selection boxes for each day of the week so that you can determine on which days the backup job should run. Figure 11.24 shows this option.
 - **Monthly** As with the **Daily** and **Weekly** options, this option allows you to specify a start time and click the **Advanced...** button to configure all of the advanced options listed previously. You can select the day of the month on which you want to have the job run. This can include patterns such as the **first Tuesday** of the month. Clicking the **Select Months** button brings up the **Select Months** dialog box with check boxes for each month of the year, all of which are selected by default.
 - **Once** This option will run the backup job one time, and allows you to specify a start time and access the **Advanced...** button options. It provides a **Run on:** drop-down box that enables you to select a date from a calendar. This is the default setting.
 - **At System Startup** This option starts the backup job when the computer is booted.
 - **At Logon** This option starts the backup job when a user logs on to the computer.
 - **When Idle** This option allows you to specify an idle setting for the computer in the **When the computer has been idle for:** entry box. This refers to the amount of time the system is not in use. The default is 10 minutes.

Figure 11.24 The Schedule Tab in the Schedule Job Dialog Box

16. At the bottom of the **Schedule** tab is the **Show multiple schedules** check box. When selected, it adds a new section to the top of the tab, which consists of a drop-down box, a **New** button, and a **Delete** button. The current schedule becomes the first entry in the drop-down box. Additional schedules can be created by clicking the **New** button and modifying the options on the **Schedule** tab. Changes to existing schedule entries can be made by selecting the schedule from the drop-down box and changing the settings on the tab. Any schedule can be deleted by selecting it in the drop-down box and clicking the **Delete** button.
17. The **Schedule Job** dialog box contains a second tab labeled **Settings**, which is shown in Figure 11.25. When selected, this tab displays a number of additional scheduling options, including the following:
 - **Delete the task if it is not scheduled to run again.** This removes the task from the list of scheduled tasks if it is not scheduled to run in the future.
 - **Stop the task if it runs for:** This allows you to specify the number of hours and minutes that a backup job can run before it is terminated. Failed backup jobs often just keep running and consuming resources. They can even cause subsequent jobs to fail because they still have exclusive use of the required system resources, such as tape drives. The default value is 72 hours.
 - **Only start the task if the computer has been idle for:** This allows you to specify how much time must pass since the computer has been used by a user before a backup must begin.

- **If the computer has not been idle that long, retry for up to:** This works in conjunction with the previous setting and allows you to specify how long the scheduler will continue to check to see if the required amount of idle time has been accumulated before giving up.
- **Stop the task if the computer ceases to be idle.** This terminates the backup job if a user begins to use the computer again.
- **Don't start the task if the computer is running on batteries.** Because backups require the use of system resources such as the hard drive, they can be very power intensive. This setting allows you to specify that you do not want to have the backup start if the computer is running on batteries. This is primarily a setting for laptops.
- **Stop the task if battery mode begins.** This setting terminates the backup job if the computer on which it is running switches over to battery power after the job has started. This setting is also primarily for laptops.
- **Wake the computer to run this task.** If a power-saving mode is in use on the computer, this selection can be used to wake the system up so that the backup job can be run at the scheduled time.

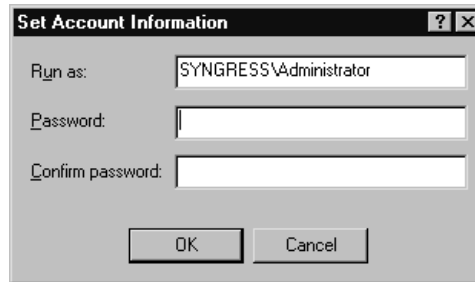
Figure 11.25 The Settings Tab in the Schedule Job Dialog Box



18. When you have finished configuring the options on each of the tabs in the **Schedule Jobs** dialog box, click the **OK** button.
19. Click the **Next** button in the wizard.

20. This brings up the **Set Account Information** dialog box, shown in Figure 11.26. The backup job must be set to run with the user rights of a member of the local administrators or backup operators group. Alternatively, it can be run by a user who has been granted the right to **Back up files and directories**. In the **Run as:** text box, specify a user account that meets these requirements. Provide the password associated with the account in the **Password:** and **Confirm password:** text boxes. Click the **OK** button to proceed.

Figure 11.26 The Set Account Information Dialog Box



21. Click the **Next** button in the wizard.
22. Review the summary information for the backup job and click the **Finish** button to close the wizard.
23. The backup will take at least a few minutes. The **Backup Progress** screen is displayed during this time, as shown in Figure 11.27. Even on the most basic Windows Server 2003 DC, the system state data will average approximately 500MB in size. Note that the file size can be even larger. The actual backup file can take up to twice as much disk space as the amount listed in the **Backup Progress** dialog box. As an example, the file shown in these images actually consumed 857,175KB of disk space.
24. When the backup has completed, click the **Close** button to close the **Backup Progress** dialog box, shown in Figure 11.28, or click the **Report...** button to view the backup log associated with the job. Clicking the **Report...** button will open the Notepad application with the log file displayed, as shown in Figure 11.29. You should review the log for any error messages, such as those pertaining to files that were skipped. After reviewing the log, close the Notepad application.

Figure 11.27 The Backup Progress Dialog Box During Backup

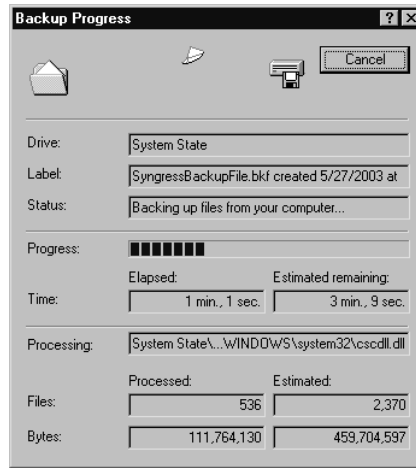


Figure 11.28 The Backup Progress Dialog After the Backup Has Completed

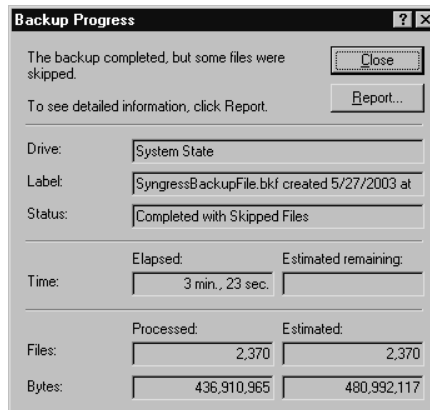
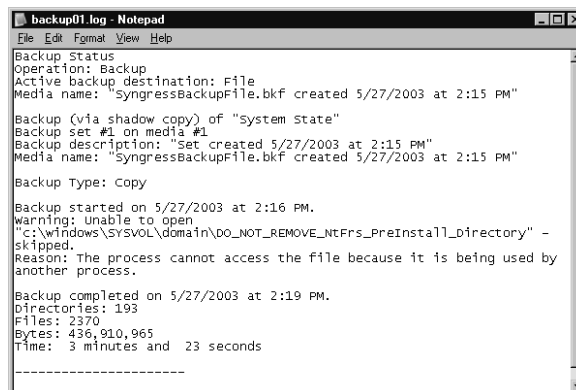


Figure 11.29 The Backup Log



Backing Up at the Command Line

Instead of using the graphical Backup utility, you can back up the system state data by using the command-line version of the Backup utility. This might be desirable for use with administrative scripts. The command-line utility is a full-featured backup program that can specify many of the same options covered in the previous section. To back up the system state data, open a command prompt (**Start | Run** and type **cmd**) and use the following command and options: **ntbackup backup systemstate /J “Syngress Backup Job” /F “C:\backupfile.bkf”**.

- **Ntbackup** is the name of the command-line backup utility.
- **Backup** is the option to specify a backup operation.
- **Systemstate** is the option used to specify that the system state data should be backed up.
- **/J** specifies the backup job name, which should be surrounded in quotes if it contains spaces.
- **/F** specifies the name of the backup file.

Note that when you run this command, the graphical utility appears to show you the progress of the job.

There are many more switches that you can use with the Ntbackup command-line utility; those described here are the ones you will most commonly use to back up the system state data.



TEST DAY TIP

Although you can back up the system state data from the command line using the Ntbackup utility, you cannot perform a restore with this utility. Restores must be done from the graphical Backup utility.

EXAM 70-294
OBJECTIVE 2.4 2.4.1 2.4.2

Restoring Active Directory

Windows Server 2003 includes three types of directory services restore methods:

- Primary
- Normal
- Authoritative

Microsoft has designed each of these restoration types to address a complex need that arises when restoring Active Directory or one of its related components. In addition to these three modes, specialty restore functionality is also provided within the Ntdsutil command-line

utility and the Directory Services Restore Mode. It is very important for you to know which modes, features, and utilities to use to restore your server in a given recovery scenario. An improper restore can destabilize your entire Active Directory forest.

New Restore Options in Windows 2003

The Active Directory restore options have seen some significant changes since Windows 2000. In Windows 2000, there were only two methods of restoration: Authoritative and Non-Authoritative. With Windows Server 2003, Authoritative restores remain unchanged; however, Non-Authoritative restores are now referred to as Normal restores. Despite the name change, they function exactly as they always have.

A new type of restore is added, the Primary restore. This is designed to be used when all DCs for a given domain have been wiped out and need to be restored. Under Windows 2000, this could be an exhaustive Authoritative restore process involving many hours of labor and double-checking. With the new Primary restore type, it is as simple as selecting a check box.

Directory Services Restore Mode

Before we discuss the three different restore methods that can be used, it is important to discuss the Directory Services Restore Mode. We mentioned this mode earlier in the chapter when discussing maintenance operations, such as moving the Active Directory database. Remember that the special feature of this mode is that it allows a DC to boot without initializing its copy of the Active Directory database. Because you must always log on to a Windows Server 2003 computer before you can use the operating system, a small version of a local directory service database (called a SAM database) remains on the computer after it has been promoted to a DC. This database has a single account, the local administrator account.

When you have booted to the Directory Services Restore Mode using the directions given earlier in the chapter, you must log on with this account. After you are authenticated, you can perform certain limited maintenance functions, such as running the Ntdsutil utility mentioned earlier. You can also run the Backup utility to perform restores of the Active Directory database. It is necessary to perform all restores while running in this mode, because the Active Directory database must be offline to be restored. In this mode, you are logged on to a local account and the Active Directory database is not in use.

Normal Restore

The simplest of all restore methods is the normal restore. This method can be used in the following circumstances:

- When a domain only has one DC, and the DC needs to be restored. You can also opt to use the primary restore method (covered later) for this scenario.
- If there are multiple DCs on the network for the domain, and at least one remains functional, a normal restore can be used to bring the downed DCs back to life.

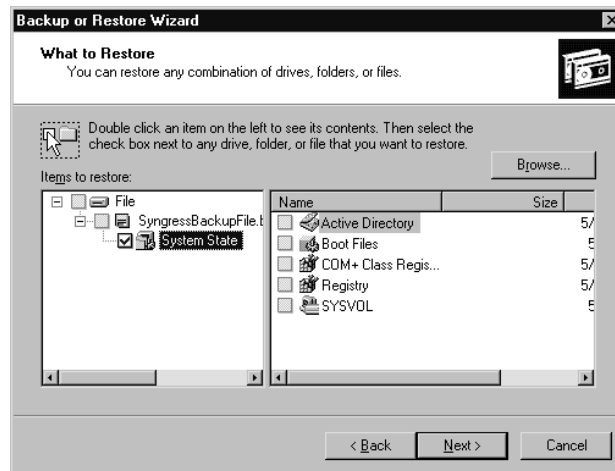
Like all Active Directory restores, a normal restore is performed by running the Backup utility while logged on to Directory Services Restore Mode. When the restore has completed, the DC is rebooted. When it comes back up, it begins normal replication with its replication partners. Because it was restored from a backup, some of its objects will have older version numbers than ones currently on the network. This will cause updates and deletions to be replicated to the DC and will bring its Active Directory database up to date. To perform a normal restore, follow these steps:

1. Boot or reboot the computer.
2. When prompted, press **F8** during Windows Server 2003 startup.
3. Select **Directory Services Restore Mode (Windows DCs only)** in the **Windows Advanced Options** menu that appears, and press the **Enter** key.
4. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.
5. You will see a number of checks performed while the system is booting, and eventually you will receive the Safe Mode logon prompt.
6. Log on by providing the password for the local administrator account and clicking the **OK** button.
7. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
8. Open the Windows Server 2003 Backup utility from **Start | All Programs | Accessories | System Tools | Backup**.
9. On the initial page of the wizard, click the **Next** button.
10. Select the option button next to **Restore files and settings**, as shown in Figure 11.30, and click the **Next** button.
11. The **What to Restore** page, shown in Figure 11.31, contains an Explorer style interface similar to the one you encountered while configuring your backup job. Click the plus sign next to **File** in the left pane. This should reveal the file to which you backed up the system state data earlier. If it doesn't, you can click the **Browse...** button and select the file from the **Open Backup File** dialog box. Click the plus sign next to the file to which you backed up and select the check box next to the backup you want to restore that appears beneath it. Click the **Next** button after making your selection.

Figure 11.30 The Backup or Restore Wizard Page



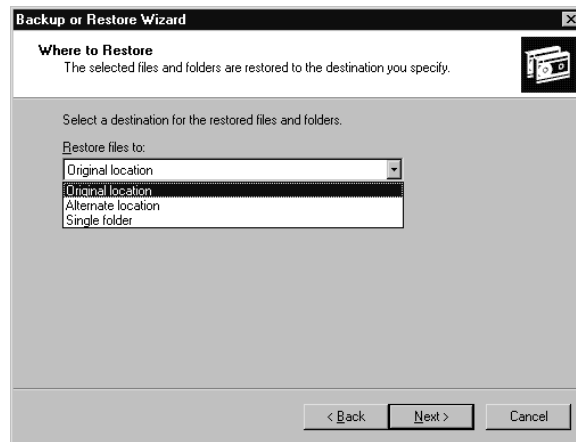
Figure 11.31 The What to Restore Wizard Page



12. At this point in the wizard, you can click the **Finish** button and allow the restore to proceed with the default advanced settings. However, we want you to see more of the settings that are available within the wizard, so click the **Advanced...** button.
13. The **Where to Restore** page, shown in Figure 11.32, appears with three options that can be selected from the **Restore files to:** drop-down box.
 - **Original location** This option restores all files to their original locations and is the default. When you select this option and click the **Next** button, a dialog box appears, informing you that restoring system state will always overwrite the current system state information unless you restore to an alternate location. Click the **OK** button to proceed to the next screen.


- **Alternate location** Selecting this option reveals the **Alternate location:** text box and a **Browse...** button that opens the **Restore Path** dialog box. You can use this option to restore the files to a different location. This can be helpful for verification and file comparison purposes.
- **Single folder** This option reveals the **Alternate location:** text box and **Browse...** button, which opens the **Restore Path** dialog box. As with the **Alternate location** setting, you can use this option to restore the files to an alternate location. When this option is selected, all restored files are placed in a single directory, rather than having their directory structures restored.

Figure 11.32 The Where to Restore Wizard Page



14. Click the **Next** button after making your selection.
15. Depending on your selection, a **Warning** dialog box (shown in Figure 11.33) might appear to inform you that a restore of system state data will always overwrite the current system state data unless you choose to restore it to an alternate location. Click the **OK** button if you receive this dialog box.

Figure 11.33 The System State Restore Warning Dialog Box



16. The **How to Restore** page, shown in Figure 11.34, contains the following three options:

- **Leave existing files (Recommended)** This option ensures that the restore process does not overwrite any files that currently exist on the DC.
- **Replace existing files if they are older than the backup files** This option permits the files on the disk to be overwritten, but only if the backup file is newer than the one currently on the DC.
- **Replace existing files** Always copies the files from the backup media to the DC and replaces all files existing on the DC, regardless of whether they are newer.

Figure 11.34 The How to Restore Wizard Page

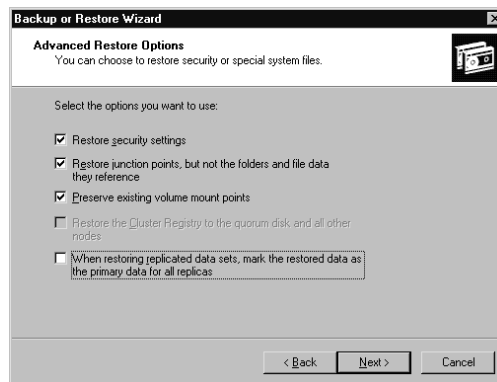


17. After making your selection, click the **Next** button to proceed.
18. The **Advanced Restore Options** page, shown in Figure 11.35, contains the following five check boxes:
 - **Restore security settings** This option is selected by default, and should remain selected. It shows the power that a user with restore rights has, because any such user can, by deselecting this check box, restore the files without their associated permissions. In some circumstances, difficulties can arise when restoring data that was on a disk formatted in the NTFS file system, which supports file level permissions, to one using the FAT file system, which does not support file level permissions. In circumstances like these, clearing this check box has been known to resolve some of the issues. This is because selecting this box restores a wide range of extended data (permissions, auditing information, and ownership information) that is not supported by the FAT file system.
 - **Restore junction points, but not the folders and file data they reference** Among other things, junction points are used to reference mounted drives. In Windows Server 2003, volumes can be mounted in folders of another volume, instead of being accessed through a drive letter. If you do not

restore junction points, you will not be able to restore the information on mounted drives unless you recreate the junction points manually.

- **Preserve existing volume mount points** This option relates to the preceding point. When using mounted drives, it is necessary to create mount points, which are the empty folders to which the volume is mounted (thus creating the mounted drive). When selected, this box protects existing mount points on the volume being restored. This is helpful if you have already formatted the disk to which you are restoring and added these mount points prior to beginning the restore. However, if you have formatted the disk to which you are restoring and have not added these mount points back manually, clearing this check box will restore your old mount points from tape. This option is selected by default.
- **Restore the Cluster Registry to the quorum disk and all other nodes** This option restores the cluster quorum database and replicates it to all of the nodes in the server cluster. This option will be grayed out if the DC is not part of a server cluster.
- **When restoring replicated data sets, mark the restored data as the primary data for all replicas** This option is used to perform a primary restore and is covered in detail later in the chapter.

Figure 11.35 The Advanced Restore Options Wizard Page



19. Click the **Next** button after making your selections.
20. Click the **Finish** button to begin the restore.
21. The restore will take at least a few minutes and display its progress as shown in Figure 11.36. When it is finished, click the **Close** button (shown in Figure 11.37) to close the **Restore Progress** dialog box, or click the **Report...** button to view the backup log associated with the job. Clicking the **Report...** button will display the Notepad application with the log file displayed, as shown in Figure 11.38. You

should review the log for any error messages, such as those pertaining to files that had to be skipped. When you have finished reviewing the log, close the Notepad application.

Figure 11.36 The Restore Progress Dialog Box During a Restore

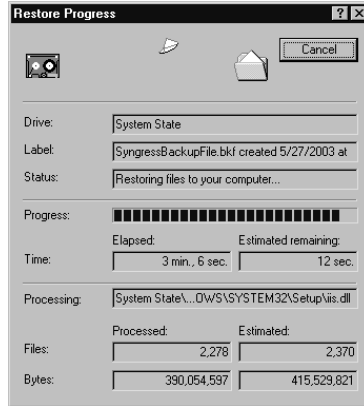


Figure 11.37 The Restore Progress Dialog Box After the Restore Has Completed


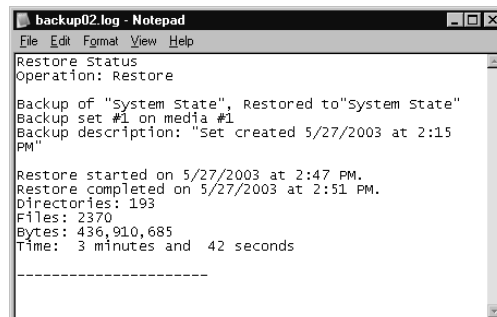


Figure 11.38 The Restore Log



22. Click the **Yes** button in the **Backup Utility** dialog box when prompted to restart and reboot the server normally.

Keep Them Separated

Although members of the administrators and backup operators local groups have the right to back up and restore files, in larger environments it is important to separate these roles for security purposes. Instead of making your backup administrators members of the backup operators group, you should consider creating two new groups instead. Grant one of the groups the right to back up files and folders, and give the other group restore privileges. Let's look at why this is a good idea.

The advanced restore option **Restore security settings** must be checked in order for NTFS permissions, auditing settings, and ownership information to be restored with a file. In other words, your company's files are a check box away from having no security permissions associated with them. This means that anyone with both backup and restore privileges can restore any file he or she wants to view with no permissions or ownership information. There is nothing standing between someone with these user rights and untraceable access to the information on your network.

You can limit the potential for abuse by making some administrators responsible for backups (and adding them to the backup group you created earlier) and other administrators responsible for restores (and adding them to the restore group you created earlier). Although those in the restore group can still restore any file without permissions, auditing, or ownership information, they cannot choose which files are backed up for them to view. Ensuring that the restore group is comprised of the most trusted employees who perform these roles will further enhance security.


Additional levels of security can also be realized through the use of encryption, including both EFS and third-party file encryption tools. Encrypted files are not decrypted before being backed up. They are stored on the tape in encrypted form, and are restored in encrypted form. As a result, although they can be restored without their associated permissions, they are still encrypted and will not be viewable by the restore administrator.

Authoritative Restore

There are times when a normal restore of Active Directory isn't sufficient; for example, when you accidentally delete an OU. Within a few minutes, the deletion will have replicated to the other DCs in the domain. If you perform a normal restore in an effort to repopulate the OU back into Active Directory, it will not work. When the DC reboots after the restore and replicates with its replication partners, they will have a higher version number for the deleted OU, and the restored DC will be told to delete the object all over again. To restore the object, you must use an authoritative restore.

An authoritative restore is exactly like a normal restore, up to a point. Once the system state data has been restored, rather than rebooting the server, the Ntdsutil command-line utility is used to mark one or more objects as *authoritative*. This gives them a very high version number so that when the server is rebooted and the replication process takes place, the other servers in the domain will see the high version number and replicate the object to their own Active Directory databases. To restore a database authoritatively, follow the steps from the preceding section up to number 18, and then proceed to these steps:

1. Click the **No** button in the **Backup Utility** dialog box when asked to restart.
2. Close the Backup utility, if it does not close by itself.
3. Open a command prompt (click **Start | Run** and type **cmd**).
4. Type **ntdsutil** to enter the Ntdsutil utility. Note that this is a command-line utility so the command prompt will change to ntdsutil:.
5. Type **authoritative restore**. The command prompt should change to display authoritative restore:.
6. Use one of the following commands to mark Active Directory or a portion of it as authoritative.
 - Type **restore database** to mark the domain and configuration containers of the database as authoritative. The schema container cannot be marked as authoritative; consequently, an authoritative restore can not be performed for the schema. Because you cannot delete objects from the schema, this is not an issue.
 - Type **restore subtree** followed by the distinguished name of the object in Active Directory that you want to restore; for example, **restore subtree OU=student,DC=syngress,DC=com** to restore the OU named “student” in the syngress.com domain.
 - The **verinc** option can be used with either the **restore database** or **restore subtree** command. Remember, when an object or the database is restored authoritatively, a large version number is applied to it. The **verinc** option is designed to be used when you need to perform another authoritative restore, on top of an existing authoritative restore. It allows you to choose your own version number, thus ensuring that it will be higher than the one used previously by the utility. The proper syntax is **restore database verinc %d** or **restore subtree <distinguished name of object to mark authoritative> verinc %d**, with **%d** being the desired increment for the version number.
7. Click **Yes** in the **Authoritative Restore Confirmation** dialog box, as shown in Figure 11.39.

Figure 11.39 The Authoritative Restore Confirmation Dialog Box

8. Review the screen output while the command completes. Figure 11.40 shows the completed operation.

Figure 11.40 The Completed Authoritative Restore Process

```

C:\WINDOWS\system32\cmd.exe - ntdsutl
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: authoritative restore
authoritative restore: restore subtree OU=student,DC=syngress,DC=com

Opening DIT database... Done.

The current time is 05-27-03 17:56.31.
Most recent database update occurred at 05-27-03 15:13.43.
Increasing attribute version numbers by 100000.

Counting records that need updating...
Records found: 000000001
Done.

Found 1 records to update.

Updating records...
Records remaining: 000000000
Done.

Successfully updated 1 records.
Authoritative Restore completed successfully.
authoritative restore: _

```

9. Type **quit** to return to the ntdsutl: prompt.
10. Type **quit** again to exit the utility.
11. Close the command prompt and reboot the server normally.

Primary Restore

The primary restore method is new in Windows Server 2003, and is designed for situations where all DCs for a given domain have gone down and you need to rebuild the domain from backup. The first server that is restored in this situation should be restored using this method. Additional DCs should be restored using the normal restore method. A primary restore is also the new preferred method to use when restoring what Microsoft refers to as a *standalone DC*, which means the DC in a domain with only one DC. If you have a domain with only one DC and that server goes down, use this method to restore it.

Performing a primary restore is similar to performing a normal restore. The only difference is that you select the check box next to **When restoring replicated data sets, mark the restored data as the primary data for all replicas** in the Advanced portion of the Restore wizard, as shown in Figure 11.35. Refer to step 14 in the *Normal Restore* section of this chapter, or complete Exercise 11.04, which walks you through the entire process of performing a primary restore.

EXERCISE 11.04

PERFORMING A PRIMARY RESTORE

1. Reboot or boot your DC.
2. When prompted, press **F8** during Windows Server 2003 startup.
3. On the Advanced Startup Options menu that appears, select **Directory Services Restore Mode**.
4. Log on by providing the password for the local administrator account and clicking the **OK** button.
5. Open the Windows Server 2003 Backup utility from **Start | All Programs | Accessories | System Tools | Backup**.
6. On the initial page of the wizard, click the **Next** button.
7. Select the option button next to **Restore files and settings**, and click the **Next** button.
8. Click the plus sign next to **File** in the left pane. If your backup file does not appear, click the **Browse...** button and select the file from the **Open Backup File** dialog box.
9. Click the plus sign next to the file to which you backed up the system state data and select the check mark next to the backup you want to restore that appears beneath it.
10. Click the **Next** button after making your selection.
11. Click the **Advanced...** button.
12. Accept the default restore location, **Original location**, and click the **Next** button.
13. Select the **Replace existing files** option and click the **Next** button to proceed.
14. On the **Advanced Restore Options** page, select the check box next to **When restoring replicated data sets, mark the restored data as the primary data for all replicas** and accept all other defaults.

15. Click the **Next** button.
 16. Click the **Finish** button to begin the restore.
 17. The restore will take at least a few minutes. When it is finished, click the **Report...** button to view the restore log associated with the job. Review it for any error messages, such as those pertaining to files that had to be skipped. After reviewing the log, close the Notepad application.
 18. Close the Backup utility and reboot the server normally.
-

EXAM
70-294
OBJECTIVE
2.5.3

Troubleshooting Active Directory Availability

Microsoft recommends checking the Event Viewer logs and careful monitoring of performance counters as initial steps when troubleshooting Active Directory availability. As mentioned previously, each of these tools can provide you with detailed and extensive information regarding where to begin your efforts. Another very important factor to consider when troubleshooting Active Directory is name resolution. Windows 2000 and later computers use the DNS service to locate Active Directory components, including GC servers and DCs.

Setting Logging Levels for Additional Detail

The default level of logging for all aspects of Active Directory is 0. This is the lowest level of logging, and while it guarantees that fatal and critical errors will be logged, it omits substantial amounts of information that can be beneficial when troubleshooting. The possible range is from 0 (which logs the least amount of information) to 5 (which logs the most). Most of the information is logged to the application log in Event Viewer.



WARNING

Setting the logging value above 3 for any aspect of Active Directory can fill the application log very quickly and substantially degrade system performance. In general, the level should be elevated temporarily only in instances when you need more information for troubleshooting purposes.

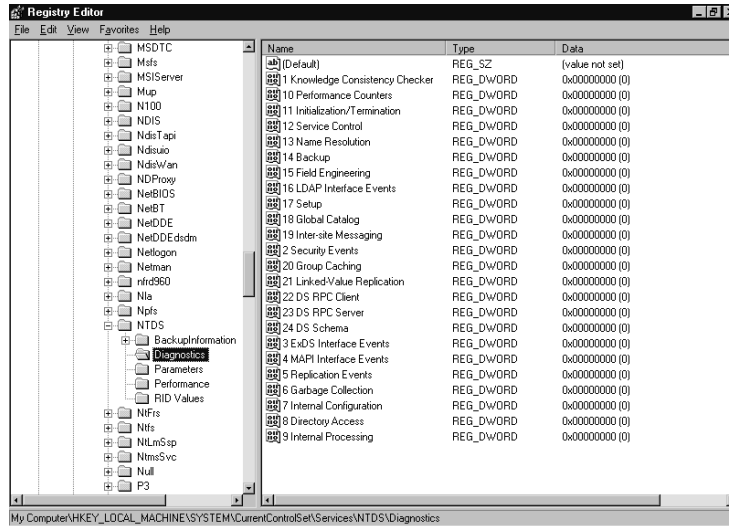
There is a wide range of individual aspects of Active Directory for which you can specify individual logging levels by editing the Registry. All of the pertinent values are located in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTDS\Diagnostics Registry subkey, shown in Figure 11.41. The available settings include:

- Knowledge Consistency Checker (KCC)
- Security Events
- ExDS Interface Events
- MAPI Interface Events
- Replication Events
- Garbage Collection
- Internal Configuration
- Directory Access
- Internal Processing
- Performance Counters
- Initialization/Termination
- Service Control
- Name Resolution
- Backup
- Field Engineering
- LDAP Interface Events
- Setup
- Global Catalog
- Inter-site Messaging
- Group Caching
- Linked-Value Replication
- DS RPC Client
- DS RPC Server
- DS Schema

In addition to the additional detail that can be specified for logging to the Event Viewer, Active Directory provides log sources for tracking and troubleshooting purposes. These are located in the %SYSTEMROOT%\Debug folder. Included are logs that were created during the installation of AD that provide significant information about the configuration of Active Directory and its related services. Other key logs in this directory relate to the FFRS. For example, the NtFrs_XXXX.log files contain detailed information about the function of FRS on the DC. There might be several of these on your system. Generally, another is added with each system reboot, and they will appear as NtFrs_0001.log to NtFrs_0005.log. When the maximum number of 5 is reached, the oldest is deleted and a

new one is created in its place, and all existing log file names will be decremented by 1. New logs will also be created when existing logs get full. By default, these logs generally hold between 1.5 to 2.5MB of information before Active Directory considers them full.

Figure 11.41 Configuring Additional Logging Detail Using the Registry



Using Ntdsutil Command Options

A number of repair options within the Ntdsutil command-line utility provide assistance in ensuring the consistency of the database. In the following subsections, we'll examine the use of these options in troubleshooting and maintaining Active Directory health and availability.

Using the Integrity Command

The *integrity* command is used to detect low-level corruption of the database. It performs its work at the binary level, which means that it reads every byte of the ESE database structure looking for corruption. Note that although the ESE structure forms the basis of Active Directory, this command might not parse all Active Directory database information. Some critical Active Directory information is additional to and outside the knowledge of the *esentutil* command that this option uses. Because of the detailed checking it performs, this tool often takes a while to complete its operations.

Where Did Esentutl Come From?

Esentutl.exe was included in Windows NT 4.0. A similarly named utility, eseutil.exe, comes with Exchange. As the name implies, these were designed for repairing ESE databases. In Windows 2000/2003, Ntldsutil acts as a “front end” for esentutl. You can also run the Esentutl utility itself, by typing **esentutl** at the command line with one of the following parameters to denote mode of operation:

- **/d <database name>** Defragmentation mode
- **/r <logfile base name>** Recovery mode
- **/g <database name>** Integrity mode
- **/k <database name>** Checksum mode
- **/p <database name>** Repair mode
- **/m[mode-modifier] <filename>** File dump

We address the use of *esentutl* by itself later in this chapter.

In addition to the byte-level corruption check mentioned previously, the *Ntldsutil integrity* command also performs a full check on the integrity of the directory service files. After successfully running the command, Microsoft suggests that you perform a semantic database analysis (covered in a later section). The *Ntldsutil integrity* command must be performed when the database is offline, so you have to run it from Directory Services Restore Mode. To use the command, follow these steps:

1. Boot or reboot the computer.
2. When prompted, press **F8** during Windows Server 2003 startup.
3. Select **Directory Services Restore Mode (Windows DCs only)** in the **Windows Advanced Options** menu that appears, and press the **Enter** key.
4. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.
5. You will see a number of checks performed while the system is booting, and eventually you will receive the Safe Mode logon prompt.
6. Log on by providing the password for the local administrator account and clicking the **OK** button.
7. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
8. Open a command prompt.

9. Type **ntdsutil** to enter the Ntdsutil utility. This is a command-line utility so the command prompt will change to **ntdsutil:**.
10. Type **files**. The command prompt should change to display **file maintenance:**.
11. Type **integrity**.
12. View and evaluate the information displayed on the screen as the process runs. Figure 11.42 shows an error-free display, and Figure 11.43 shows a display showing errors.

Figure 11.42 A Successful Integrity Check Showing No Errors

```

C:\WINDOWS\system32\cmd.exe - ntdsutil
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: files
File maintenance: integrity
Opening database [Current].
Executing Command: C:\WINDOWS\system32\esentutl.exe /g"C:\WINDOWS\NTDS\ntds.dit"
/o

Initiating INTEGRITY mode...
Database: C:\WINDOWS\NTDS\ntds.dit
Temp. Database: TEMPINTEGI360.EDB

Checking database integrity.

          Scanning Status (% complete)
0    10   20   30   40   50   60   70   80   90  100
|-----|-----|-----|-----|-----|-----|-----|
.....

Integrity check successful.
Operation completed successfully in 4.847 seconds.

Spawned Process Exit code 0x0(0)

If integrity was successful, it is recommended
you run semantic database analysis to ensure
semantic database consistency as well.

File maintenance:

```

13. Type **quit** to return to the ntdsutil: prompt.
14. Type **quit** again to exit the utility.
15. Close the command prompt window and reboot the server normally.

Figure 11.43 An Integrity Check Showing Errors

```

C:\WINDOWS\system32\cmd.exe - ntdsutl
C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: Files
File maintenance: integrity
Opening database [Current].*** Error: DBInitializeJetDatabase failed with [ Jet
Error -528].
Error While Doing Soft Recovery
Executing Command: C:\WINDOWS\system32\esentutl.exe /g"C:\WINDOWS\NTDS\ntds.dit"
/o
Initiating INTEGRITY mode...
Database: C:\WINDOWS\NTDS\ntds.dit
Temp. Database: TEMPINTEG316.EDB
Checking database integrity.
The database is not up-to-date. This operation may find that
this database is corrupt because data from the log files has
yet to be placed in the database.
To ensure the database is up-to-date please use the 'Recovery' operation.

Scanning Status (% complete)
0 10 20 30 40 50 60 70 80 90 100
|-----|-----|-----|-----|-----|-----|-----|-----|
.....

Integrity check completed. Database is CORRUPTED!

Operation terminated with error -1206 (JET_errDatabaseCorrupted, Non database fi
le or corrupted db) after 4.467 seconds.

Spawned Process Exit code 0xffffb4a(-1206)
If integrity was successful, it is recommended
you run semantic database analysis to ensure
semantic database consistency as well.

File maintenance:

```

Using the *recover* Command

Remember that transactions are written to log files before being committed to the Active Directory database file. In the event of power failure or other system problems, not all transactions will be written to the database. When the system is booted, ESE should use the checkpoint, log, and database files to determine what was committed properly to the database and what still needs to be written. Although this process works in most cases, occasionally inconsistencies result and it is necessary to run the process again manually. The *recover* command performs a “soft” recovery of the database log files, which means that it writes transactions from the log files to the directory service database. This process is sometimes also referred to as “re-running” the log files manually.

Like the other commands used in conjunction with the *Ntdsutl* command, the *recover* command must be run from Directory Services Restore Mode. As with the other maintenance commands covered in this section, Microsoft recommends running a semantic database analysis after the *recover* command has completed successfully. To run the *recover* command, follow these steps:

1. Boot or reboot the computer.
2. When prompted, press **F8** during Windows Server 2003 startup.
3. Select **Directory Services Restore Mode (Windows DCs only)** in the Windows Advanced Options menu that appears, and press the **Enter** key.
4. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.

5. You will see a number of checks performed while the system is booting, and eventually you will receive the Safe Mode logon prompt.
6. Log on by providing the password for the local administrator account and clicking the **OK** button.
7. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
8. Open a command prompt.
9. Type **ntdsutil** to enter the Ntdsutil utility. This is a command-line utility, so the command prompt will change to **ntdsutil:**.
10. Type **files**. The command prompt should change to display **file maintenance:**.
11. Type **recover**.
12. View and evaluate the information displayed on the screen as the process runs. Figure 11.44 shows a successful recover operation, and Figure 11.45 shows a failed recover operation.

Figure 11.44 A Successful Recover Operation

```

C:\WINDOWS\system32\cmd.exe - ntdsutil
C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: files
file maintenance: recover
Executing Command: C:\WINDOWS\system32\esentutl.exe /redb /1"C:\WINDOWS\NTDS" /s
"C:\WINDOWS\NTDS" /8 /o

Initiating RECOVERY mode...
  Logfile base name: edb
  Log files: C:\WINDOWS\NTDS
  System files: C:\WINDOWS\NTDS

Performing soft recovery...

Operation completed successfully in 2.303 seconds.

Spawned Process Exit code 0x0<0>

If recovery was successful, it is recommended
you run semantic database analysis to ensure
semantic database consistency as well.

file maintenance: _

```

13. Type **quit** to return to the **ntdsutil:** prompt.
14. Type **quit** again to exit the utility.
15. Close the command prompt window.

Figure 11.45 A Failed Recover Operation

```

C:\WINDOWS\system32\cmd.exe - ntdsutil
C:\Documents and Settings\Administrator>set SAFEBOOT_OPTION=DSREPAIR
SAFEBOOT_OPTION=DSREPAIR
C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: files
file maintenance: recover
Executing Command: C:\WINDOWS\system32\esentutl.exe /redb /1"C:\WINDOWS\NTDS" /s
"C:\WINDOWS\NTDS" /8 /o

Initiating RECOVERY mode...
  Logfile base name: edb
  Log files: C:\WINDOWS\NTDS
  System files: C:\WINDOWS\NTDS

Performing soft recovery...

Operation terminated with error -1032 (JET_errFileAccessDenied, Cannot access fi
le, the file is locked or in use) after 20.18 seconds.

Spawned Process Exit code 0xfffffb8(-1032)

If recovery was successful, it is recommended
you run semantic database analysis to ensure
semantic database consistency as well.

file maintenance:

```

Using the Semantic Database Analysis Command

The *semantic database analysis* command is the primary command that is used to verify the full integrity of the Active Directory database. You might be wondering what the difference is between this command and the *integrity* command from the **files:** prompt. Recall that the *integrity* command works by calling the **Esentutl** utility, which has full knowledge of the ESE database system but not necessarily all portions of the Active Directory database. The *semantic database analysis* command is specific to Active Directory and does not use the *Esentutl* command. As its name implies, it analyzes the Active Directory database, based on Active Directory semantics (whereas the *integrity* command bases its check on ESENT database semantics). Running *semantic database analysis* includes checks for the following:

- **Reference counts**
 - Counts references from the data table and the link table to ensure that they match the listed counts for the record.
 - Ensures that each object has a full distinguished name, GUID, and nonzero reference count.
 - For each deleted object, the utility verifies that it does not have a distinguished name or GUID and makes sure that it has a deleted time and date.
- **Deleted objects**
 - Verifies that the object has a deleted time and date.
 - Ensures that the object has a special relative distinguished name.
- **Ancestor checks** Determines if the Distinguished Name Tag is equal to:
 - The ancestor list of the parent

- The current Distinguished Name Tag
- **Security descriptor checks**
 - Verifies a valid descriptor.
 - Ensures that it has a control field.
 - Verifies that the discretionary access control list is not empty.
 - A warning is generated if deleted objects without a discretionary control access list are located.
- **Replication checks.**
 - Checks the up-to-dateness vector in the directory partition head to ensure that the correct number of cursors exist.
 - Checks to ensure that every object has a property metadata vector.

Errors generated by the *semantic database analysis* command are written to `dsdit.dmp.xx` log files, which are located in the profile directory of the user running the utility (for example, `C:\Documents and Settings\Administrator`). As with most low-level database tools, this command must be run when the database is not initialized (in other words, in Directory Services Restore Mode). Microsoft recommends that you perform a full backup of the system state data prior to running this command. Follow these steps to perform a semantic database check:

1. Boot or reboot the computer.
2. When prompted, press **F8** during Windows Server 2003 startup.
3. Select **Directory Services Restore Mode (Windows DCs only)** in the **Windows Advanced Options** menu that appears, and press the **Enter** key.
4. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.
5. You will see a number of checks performed while the system is booting, and eventually you will receive the Safe Mode logon prompt.
6. Log on by providing the password for the local administrator account and clicking the **OK** button.
7. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
8. Open a command prompt.
9. Type **ntdsutil** to enter the Ntdsutil utility. This is a command-line utility, so the command prompt will change to **ntdsutil.**
10. Type **Semantic database analysis**, and press the **Enter** key.

11. At the **semantic checker:** prompt, type **Verbose on**, and press **Enter**. This option displays the Semantic Checker.
12. Choose one of the following options:
 - To start the Semantic Checker and not have it repair any of the errors it encounters, type **Go**, and press the **Enter** key.
 - To start the Semantic Checker and have it repair the errors it encounters, type **Go Fixup**, and press the **Enter** key.
13. View and evaluate the information displayed on the screen as the process runs. There is very little difference visually between the two modes. Figure 11.46 shows the **go** mode.

Figure 11.46 Semantic Database Analysis Using Go Mode

```

C:\WINDOWS\system32\cmd.exe - ntdsutil
C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: Semantic database analysis
semantic checker: Verbose on
Verbose mode enabled.
semantic checker: go
Fixup mode is turned off
Opening database [Current].....Done.

Getting record count...2971 records
Getting security descriptor count...76 security descriptors

Writing summary into log file dsdit.dmp.0
SDs scanned:          76
Records scanned:     2971
Processing records..Done.

semantic checker: _
  
```

14. Type **quit** or **q** to return to the ntdsutil: prompt.
15. Type **quit** or **q** again to exit the utility.
16. Close the command prompt window.
17. Navigate to your profile directory and use Notepad to open the log file (shown in Figure 11.47) that you saw in step 15 (for example, dsdit.dmp.0).
18. View the contents of the log, paying careful attention to any warning messages, and then close Notepad.
19. Reboot the server normally.

In Exercise 11.05, you will practice using the Ntdsutil tool to perform a semantic database analysis of Active Directory.

Figure 11.47 The Log File Produced by Semantic Database Analysis Go Mode

```

Property Metadata vector missing for 2($ROOT_OBJECTS)
INFO: Uptodate vector found for NC head 1458(syngress)
INFO: Uptodate vector found for NC head 1459(Configuration)
WARNING: Deleted object 1471 has timestamp [12/29/9999] later than now
WARNING: Deleted object 1477 has timestamp [12/29/9999] later than now
INFO: Uptodate vector found for NC head 1480(Schema)
warning SE_DACL_PROTECTED for 1687(Volumetable)
warning SE_DACL_PROTECTED for 1693({31B2F340-016D-11D2-945F-00C04FB984F9})
warning SE_DACL_PROTECTED for 1696({6AC1786C-016F-11D2-945F-00C04FB984F9})
warning SE_DACL_PROTECTED for 1743(AdminSDHolder)
warning SE_DACL_PROTECTED for 1746(WMI Policy)
warning SE_DACL_PROTECTED for 1748(SOM)
warning SE_DACL_PROTECTED for 1814(Administrator)
warning SE_DACL_PROTECTED for 1820(Administrators)
warning SE_DACL_PROTECTED for 1825(Print Operators)
warning SE_DACL_PROTECTED for 1826(Backup Operators)
warning SE_DACL_PROTECTED for 1827(Replicator)
warning SE_DACL_PROTECTED for 1859(krbtgt)
warning SE_DACL_PROTECTED for 1861(Domain Controllers)
warning SE_DACL_PROTECTED for 1862(Schema Admins)
warning SE_DACL_PROTECTED for 1863(Enterprise Admins)
warning SE_DACL_PROTECTED for 1865(Domain Admins)
warning SE_DACL_PROTECTED for 1870(Server Operators)
warning SE_DACL_PROTECTED for 1871(Account Operators)
INFO: Uptodate vector found for NC head 1885(domainDnsZones)
WARNING: Deleted object 1888 has timestamp [12/29/9999] later than now
INFO: Uptodate vector found for NC head 1895(ForestDnsZones)
WARNING: Deleted object 1898 has timestamp [12/29/9999] later than now
12971 total records walked.
Summary:
Active objects      2960
Phantoms            1
Deleted             10
Security descriptor summary:
SD count:           76
Total SD size before single-instancing:      815 kb
Total SD size after single-instancing:       69 kb

```

EXERCISE 11.05

PERFORMING A SEMANTIC DATABASE ANALYSIS

1. Boot or reboot your computer.
2. When prompted, press **F8** during Windows Server 2003 startup.
3. Select **Directory Services Restore Mode (Windows DCs only)** in the **Windows Advanced Options** menu that appears, and press the **Enter** key.
4. Select your operating system (for example, **Windows Server 2003, Enterprise**), and press the **Enter** key.
5. You will see a number of checks performed while the system is booting, and eventually you will receive the Safe Mode logon prompt.
6. Log on by providing the password for the local administrator account and clicking the **OK** button.
7. Click the **OK** button in the dialog box that notifies you that Windows is running in safe mode.
8. Open a command prompt.
9. Type **ntdsutil** to enter the Ntdsutil utility. This is a command-line utility so the command prompt will change to **ntdsutil:**

10. Type **semantic database analysis**, and press the **Enter** key.
11. At the **semantic checker:** prompt, type **Verbose on**, and press **Enter**. This option displays the Semantic Checker.
12. To start the Semantic Checker and not have it repair any of the errors it encounters, type **Go**, and press the **Enter** key.
13. View and evaluate the information displayed on the screen as the process runs. Note the name of the log file that information is being written to (for example, dsdit.dmp.0).
14. Type **quit** or **q** to return to the ntdsutil: prompt.
15. Type **quit** or **q** again to exit the utility.
16. Close the command prompt window.
17. Navigate to your profile directory and use Notepad to open the log file that you saw displayed in step 13 (for example, dsdit.dmp.0).
18. View the contents of the log, paying careful attention to any warning messages, and then close Notepad.
19. Reboot the server normally.



EXAM WARNING

Remember, only the *semantic database analysis* command is specific to Active Directory (checks the database with respect to Active Directory semantics). Although the *integrity* command can also be used from within the Ntdsutil utility, it is more generic to the ESE database engine and does not see all information that pertains to Active Directory.

Using the *esentutil* Command

Although all of the operations covered previously in this section used the Ntdsutil command-line utility, most actually performed their work by calling the *Esentutil* command. ESENT (Extensible Storage Engine for NT) is one of the acronyms used to refer to the ESE database system that Active Directory uses. The *Esentutil* command is the maintenance command that is associated with this database system. Because Microsoft prefers that you use the *Ntdsutil* command for all low-level database maintenance operations, they built calls to most of the major *Esentutil* operations into it. However, you do not have to use Ntdsutil to perform these operations. The following are two of the commands from earlier in the chapter with their associated *Esentutil* command-line arguments:

- **Integrity** %SYSTEMROOT% \System32\esentutil.exe /g "C:\Windows\NTDS\ntds.dit" /o
- **Recover** %SYSTEMROOT%\System32\esentutil.exe /redb /I"C:\Windows\NTDS" /s "C:\WINNT\NTDS" /8 /o

The *esentutl.exe* command used in conjunction with the */p* switch, shown in Figure 11.48, is considered the most dangerous of all the low-level database commands. In Windows 2000, this command was available as the *repair* option in Ntdsutil, and has been removed in the version of Ntdsutil that ships with Windows Server 2003. This option performs a very low-level and highly invasive binary database repair operation. It is very likely that you will lose some data when using this option, and it is highly possible that it will be data essential to your Active Directory database.

You should use this command with the */p* switch only when you have been advised to do so by Microsoft support personnel, or when you feel that you have tried everything else to get Active Directory to initialize. *Always make a backup of your database file before you run this utility.* In most cases, you will be resorting to this option when Active Directory can no longer initialize, and you will be booted to Directory Services Restore Mode. The simplest way to back up the database and related components in this scenario is to copy them to a second location in the file system, using Windows Explorer.

If Active Directory can initialize and you still feel you should (or Microsoft tech support asks you to) run this command, you must boot into Directory Services Restore Mode first. The database must be offline for low-level operations such as this. Microsoft recommends running a semantic database analysis after this command has completed successfully. To use the *repair* command, enter the following at a command prompt: `%SYSTEMROOT%\system32\esentutl.exe /p "C:\Windows\NTDS\ntds.dit" /!10240 /8 /o`

Figure 11.48 The *esentutl* Repair Process

```

C:\WINDOWS\system32\cmd.exe
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

C:\Documents and Settings\Administrator>%SYSTEMROOT%\system32\esentutl.exe /p "C:\Windows\NTDS\ntds.dit" /!10240 /8 /o

Initiating REPAIR mode...
    Database: C:\Windows\NTDS\ntds.dit
    Temp. Database: TEMPREPAIR204.EDB

Checking database integrity.

        Scanning Status (% complete)

0    10  20  30  40  50  60  70  80  90 100
|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|
.....

Integrity check successful.

Note:
It is recommended that you immediately perform a full backup
of this database. If you restore a backup made before the
repair, the database will be rolled back to the state
it was in at the time of that backup.

Operation completed successfully in 4.576 seconds.

C:\Documents and Settings\Administrator>

```



TEST DAY TIP

Remember, there is no longer a Repair option on the files: menu in Ntdsutil as there was with Windows 2000. To perform a binary repair of the database in Windows Server 2003, you must run *esentutl /p*.

Changing the Directory Services Restore Mode Password

Because the Directory Services Restore Mode password is set during the installation of Active Directory, administrators often have difficulty remembering the password that was used when it is needed later. Fortunately, there is a way to change this password without having to remember what it was originally: by using the Ntdsutil command-line utility. To use this feature, the server on which you want to change the password cannot be running in Directory Services Restore Mode. Ntdsutil can be used to change the password on the DC locally, or another DC within the forest. To change the Directory Services Restore Mode password, follow these steps:

1. Open a command prompt.
2. Type **ntdsutil** to enter the **Ntdsutil** utility. This is a command-line utility, so the command prompt will change to **ntdsutil:**.
3. Type **Set DSRM Password**.
4. At the **Reset DSRM Administrator Password:** prompt, type **Reset Password on server <SERVER NAME>**.
5. At the **Please type password for DS Restore Mode Administrator Account:** prompt, type the new password that you want to use.
6. At the **Please confirm new password:** prompt, re-type the new password that you want to use.
7. Review the feedback on the screen to ensure that the operation was successful. Figure 11.49 shows the full procedure.
8. Type **quit** or **q** to return to the ntdsutil: prompt.
9. Type **quit** or **q** again to exit the utility.
10. Close the command prompt window.

Figure 11.49 Using *Ntdsutil* to Reset the DSRM Password on a Server



```

C:\WINDOWS\system32\cmd.exe - ntdsutil
C:\Documents and Settings\Administrator>ntdsutil
ntdsutil: Set DSRM Password
Reset DSRM Administrator Password: Reset Password on server syngressdc
Please type password for DS Restore Mode Administrator Account: *****
Please confirm new password: *****
Password has been set successfully.
Reset DSRM Administrator Password:
  
```

Summary of Exam Objectives

In this chapter, we examined the many factors involved in ensuring Active Directory availability. In the first section, we began with a thorough examination of the Active Directory database and its related files: Ntds.dit, Edb*.log, Res1.log, Res2.log, and Edb.chk. Next, we discussed the ESE database engine and how data is updated in Active Directory. You learned that ESE captured changed information and its related metadata, and packaged it into *transactions*. These transactions are initially written to memory, and then to the Edb.log file. Finally, they are written to the Active Directory database, and the checkpoint file is updated so that the system knows what has and has not been fully committed to the database.

We also examined the Garbage Collection process. You learned that it runs every 12 hours by default and is responsible for deleting expired objects (those that had already been tombstoned), removing any fully committed log files and running an online defragmentation of the database. You also became familiar with the importance of the tombstone process and how it ensures that deleted Active Directory objects do not accidentally get reintroduced into the database.

Next, we discussed the significance of system state data. You learned that the system state data contains the most critical system configuration information. You also learned that the data it contains depends on the Windows Server 2003 components installed in the system. Finally, you learned the critical role that hardware plays in fault tolerance and performance.

In the second section of the chapter, we examined Active Directory maintenance tasks. We began with a look at the defragmentation process and you learned that there are two different types of defragmentation for the Active Directory database: online and offline. Both types rearrange the contents of the database so that the stored data and free space are contiguous, and re-index the database. These measures provide greater database performance. However, the online defragmentation process only rearranges the way the data is written to the *database*, whereas the offline defragmentation process rearranges the way the data is written to the *disk*. You also learned that the online defragmentation process runs automatically, but does not shrink the database file as free space becomes available in it. The offline defragmentation process must be run manually, but does recover free space from the database and shrink its file.

Next, you learned how to move the Active Directory database and log files, using the Ntdsutil command-line utility. We then looked at monitoring Active Directory with the Event Viewer utility and Performance console. In this final portion of the section, we looked at Microsoft's specific recommendations as to which event log messages to look for, and which performance counters to track for optimum performance.

The third section of the chapter covered backing up and restoring Active Directory. You learned that the best way to back up Active Directory was with system state data. You also learned that there are three ways to restore Active Directory: primary, normal, and authoritative. You learned that all restore operations must be performed while the database is offline in Directory Services Restore Mode. You discovered that the primary restore type is new in Windows Server 2003 and is designed to be used when all DCs in a given domain

must be restored. In this situation, a primary restore should be performed on the first DC that is restored, followed by normal restores on the others.

You learned that a normal restore is used when other DCs still exist on the network with current versions of Active Directory on them. In this scenario, the downed DC is restored and uses the Active Directory replication process to bring its database up to date upon reboot. Finally, you learned that the authoritative restore method is used when an object, such as an OU, has been deleted from Active Directory and must be reintroduced back into the directory. As part of the restore process, the object is marked as authoritative, using the `Ntdsutil` utility. This gives it a higher version number than the tombstoned object that exists on the other active DCs and ensures that it will propagate back out to them when the restore is completed.

The final section of the chapter examined troubleshooting Active Directory availability. In this section, you saw that a number of additional Active Directory logging levels can be enabled using the Registry. You also learned about some of the additional `Ntdsutil` commands that can be used to check the integrity of and repair the database. We discussed how the *integrity* command uses the `Esentutil` utility to perform a binary integrity check of the database. The *recover* command uses the `Esentutil` utility to perform a “soft” recovery of the Active Directory database by re-running its log files. The *semantic database analysis* command is the only native Active Directory command and thus provides the most thorough examination of Active Directory database integrity. We examined some of the available switches that can be used with the *esentutil* command, including the `/p` switch that is used to perform a full, binary-level database repair. Finally, we showed you how to use the `Ntdsutil` command to change the Directory Services Restore Mode password.

For mission-critical computers such as DCs that hold information in Active Directory that is vital to the operation of the network, high availability is not a luxury, it’s a necessity. We hope that the information in this chapter will be useful to you when you take Exam 70-294, and on the job as well.

Exam Objectives Fast Track

Understanding Active Directory Availability Issues

- ☑ The major files associated with the Active Directory database are `Ntds.dit`, `Edb.log`, `Edb*.log`, `Edb.chk`, `Res1.log`, and `Res2.log`.
- ☑ The Extensible Storage Engine (ESE) governs the process of creating, deleting, and modifying objects, properties, and attributes in Active Directory.
- ☑ The Garbage Collection process runs every 12 hours by default and purges expired objects from the Active Directory database.

- ☑ *System state data* is the term Microsoft uses to refer to the core set of configuration information on a Windows Server 2003 computer. The configuration information it contains depends on the components that are installed.

Performing Active Directory Maintenance Tasks

- ☑ An online defragmentation of the Active Directory database runs every 12 hours by default, following the Garbage Collection process. Online defragmentations do not shrink the size of the Active Directory database file.
- ☑ Offline defragmentation of the Active Directory database must be performed manually and does shrink the size of the database file. The Ntdsutil command-line utility is used to perform an offline defragmentation of the database.
- ☑ You can use the Ntdsutil utility to move the Active Directory database or log files to another location in the file system.
- ☑ The Performance console contains special counters that enable you to monitor Active Directory performance and availability.

Backing Up and Restoring Active Directory

- ☑ The primary tools for backing up and restoring Active Directory are the graphical Windows Server 2003 Backup tool and the **Ntbackup** command-line utility.
- ☑ Backing up the system state data is the recommended method to use in backing up Active Directory and its related components.
- ☑ A *normal* restore is used to restore Active Directory on a DC when changes made since the backup can be replicated from other DCs within the domain.
- ☑ An *authoritative* restore is used when an object is accidentally deleted from Active Directory and needs to be restored and propagated back to all DCs in the domain.
- ☑ A *primary* restore is used to restore the first server in a domain that has lost all its active DCs.

Troubleshooting Active Directory Availability

- ☑ Windows 2000 and later clients use the DNS service to locate Active Directory, so before you suspect directory services problems, ensure that DNS is functioning properly.
- ☑ Additional levels of Active Directory logging detail can be specified in the Registry.

- ☑ The *Ntdsutil* command contains a number of useful troubleshooting commands, including *integrity*, *recover*, and *semantic database analysis*.
- ☑ The *esentutil.exe /p* command can be used to perform a binary-level repair of the AD database.
- ☑ *Ntdsutil* can be used to change the Directory Services Restore Mode password on any DC in the forest.

Exam Objectives

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the Exam Objectives presented in this chapter, and to assist you with real-life implementation of these concepts. You will also gain access to thousands of other FAQs at ITFAQnet.com.

Q: What does ESE provide in an Active Directory environment?

A: The Extensible Storage Engine (ESE) is the heart of the Active Directory update process. It listens for changes to Active Directory, creates transactions, writes the transactions to the log file and Active Directory database, and updates the checkpoint file. It is also responsible for creating new transaction logs and deleting them when the current log file has had all of its transactions committed.

Q: What purpose does the system state data serve?

A: *System state data* is a term Microsoft uses to refer to core configuration information on a Windows 2000 or later computer. The actual information contained in the system state data depends on the components that are installed in the underlying operating system. All Windows Server 2003 systems include the Registry, COM+ database, and boot and system files in system state data. DCs also include the Active Directory database and related components such as SYSVOL as part of the system state data. The system state data is used to back up the most important configuration information for the operating system and its most critical services.

Q: What utility is used to move the Active Directory database and log files and perform an offline defragmentation of the Active Directory database?

A: The *Ntdsutil* command-line utility is a low-level database maintenance tool that provides this functionality.

- Q:** Into what mode must the server be booted in order to perform an offline defragmentation of the Active Directory database?
- A:** The Directory Services Restore Mode is a special boot mode that does not initialize the Active Directory database. Because it is not loaded, the database can have low-level operations performed on it; for example, move and offline defragmentation operations.
- Q:** What is the primary object that I should monitor in the Performance console for Active Directory directory services?
- A:** The NTDS object is the primary object that Microsoft provides for monitoring Active Directory. It contains individual counters that can be used to monitor directory services.
- Q:** What can I use to back up Active Directory from the command line?
- A:** The *Ntbackup* command can be used in conjunction with the *systemstate* option to back up Active Directory from the command line.
- Q:** What are the three types of restores that are possible with Active Directory in Windows Server 2003?
- A:** Microsoft has added a new restore type in Windows Server 2003, the primary restore. The normal (formerly non-authoritative) and authoritative restore types supported in Windows 2000 are also still available.
- Q:** Which commands can I use within *Ntdsutil* to verify the consistency of the Active Directory database?
- A:** You can use the *integrity* and *semantic database analysis* commands to verify the consistency of the database.
- Q:** I suspect a problem with my Active Directory database. Should I run the repair option in *Ntdsutil*?
- A:** This option no longer exists within the version of *Ntdsutil* that is included with Windows Server 2003. It must be run manually, using the *esentutil.exe* command with the */p* switch. You should only run this command under the guidance of senior support personnel, or Microsoft tech support. It is meant for use when everything else has been tried and you are using it as a last resort. Active Directory data loss is a likely result of using this tool.

Self Test

A Quick Answer Key follows the Self Test questions. For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

Understanding Active Directory Availability Issues

1. Your supervisor asks you how Active Directory knows which transactions have been committed to the database. You explain that this is tracked in a file known as:
 - A. Edb.log
 - B. Ntds.dit
 - C. Edb.chk
 - D. Edb00001.log
2. You are studying the Windows Server 2003 Resource Kit in preparation for deployment of your company's new infrastructure. You learn that updates to the Active Directory database are transaction based. Which of the following are components of a transaction? (Choose all that apply.)
 - A. Existing Active Directory data that has been altered by an administrator or by the replication process
 - B. New Active Directory data that has been altered by an administrator or by the replication process
 - C. All unchanged information that describes the object
 - D. Metadata
3. You are explaining the importance of system state data to your supervisor. Among other things, you tell her that the information it contains depends on the components installed in the underlying operating system. Which of the following do you tell her are always part of system state data on a Windows Server 2003 DC? (Choose all that apply.)
 - A. The certificate database
 - B. The Registry
 - C. The Active Directory database
 - D. The metabase
 - E. SYSVOL

4. You are trying to explain the Active Directory deletion process to one of your coworkers. Despite your best efforts, he can't seem to grasp the concept of the Garbage Collection process. Which of the following key points do you reiterate? (Choose all that apply.)
- A. The Garbage Collection process runs every 12 hours by default.
 - B. An offline defragmentation of the database runs automatically at the end of the Garbage Collection process.
 - C. The Garbage Collection process deletes all expired objects from the Active Directory database.
 - D. The Garbage Collection process deletes all tombstoned objects from the Active Directory database.

Performing Active Directory Maintenance Tasks

5. Your Windows Server 2003 DC is running out of disk space on the partition containing the database and log files. Which of the following steps will best rectify the situation before critical failure and additional performance degradation occurs? (Choose all that apply.)
- A. You can compress the database file using, NTFS compression.
 - B. You can install an additional hard drive.
 - C. You can move the Active Directory database and log files, using Windows Explorer, and reboot the server.
 - D. You can move the Active Directory database and log files, using the Ntdstutil utility
6. You recently started a new job as a network administrator. Upon reviewing the servers for which you are responsible, you notice that Active Directory is failing to start on one of your DCs. Further investigation reveals that Active Directory cannot find the location of the log files. You verify that the log files are located in the D:\ADLogs directory. Which of the following actions is the best to use in resolving this problem?
- A. You can delete the log files so that they are automatically recreated.
 - B. You can boot into Directory Services Restore Mode and use the *set path logs* command within the Ntdsutil utility.
 - C. You can boot into Directory Services Restore Mode and use the *move logs to* command within the Ntdsutil utility.
 - D. You can directly modify the Registry so that it points to the new location.

7. You have been promoted to network engineer and are training someone from the help desk to move up to your previous network admin role. You stress the importance of monitoring the log files each day to spot Active Directory problems early. Which of the following logs do you specify as important for the new admin to monitor? (Choose all that apply.)
- A. The System event log
 - B. The Application event log
 - C. The Directory Service event log
 - D. The File Replication Service event log
 - E. The DNS Server event log
8. The network engineer for whom you work is concerned about the impact of monitoring too many performance counters. She asks that you limit your monitoring to the core directory service performance object and its counters only. Which of the following will this allow you to monitor?
- A. DS Search sub-operations/sec
 - B. DRA inbound and outbound counters
 - C. Context Switches/sec
 - D. % DPC Time

Backing Up and Restoring Active Directory

9. You've just taken over a new job and have been reviewing the backup logs. On one of your new DCs, the log shows that the Active Directory database is failing to back up. It also shows that several key components of the operating system, such as the Registry, are not even selected for backup. Which of the following is the best option for ensuring that these critical system components get backed up?
- A. Verify that Volume Shadow Copy is enabled
 - B. Select to back up system state data
 - C. Enable full system backup
 - D. Verify that there is sufficient free space on the drive to hold a shadow copy
10. Your coworker has been working to script an automated restore of system state data but has not been successful. He asks for your help. What do you tell him?

- A. A macro is required to script a restore using Ntbackup.
 - B. A CMD file must be used for scripting a restore.
 - C. A BAT file must be used when scripting a restore.
 - D. It is not possible to script a restore using the Ntbackup utility.
11. Over the weekend, a tornado ripped through your company's Kansas office. You arrived to find no functioning DCs. The office has its own domain within your company's forest. Duplicate hardware is due to arrive in two days, and offsite backup tapes should arrive tomorrow. Which of the following restore methods will you use for the three DCs at the facility? (Choose all that apply.)
- A. Primary
 - B. Normal
 - C. Non-authoritative
 - D. Authoritative
12. One of the help desk employees has been assigned to assist with backups. She just sent you an e-mail stating that although she can back up her own files and some of the files in her department, she is unable to back up any other files on the network. Which of the following can enable her to back up any file? (Choose all that apply.)
- A. Her user account can be added to the local administrators group.
 - B. Her user account can be granted the *Back up files and directories* right.
 - C. Her user account can be added to the local backup operators group.
 - D. Her user account can be added to a group that has been granted the *Back up files and directories* right.

Troubleshooting Active Directory Availability

13. You've been asked to train a new junior administrator on the basics of Active Directory troubleshooting. Which of the following do you tell him you will be covering? (Choose all that apply.)
- A. The Performance console
 - B. Event Viewer
 - C. DNS
 - D. DCPROMO

14. Your well-honed administrative skills are telling you that something is not quite right with your Active Directory database. You've reviewed the log files and performance monitor data. No users are complaining about connectivity problems. What can you do next to attempt to verify your hunch?
- A. You can run *Ntdsutil* with the *integrity* option.
 - B. You can run *esentutil* with the */g* switch.
 - C. You can run *esentutil* with the */p* switch.
 - D. You can run *Ntdsutil* with the *Semantic database analysis* option.
15. The administrator with whom you work believes that the Active Directory problem she's troubleshooting could be solved if she could manually re-run the Active Directory log files. Which of the following commands does exactly that? (Choose all that apply.)
- A. The *Ntdsutil integrity* command
 - B. The *Ntdsutil recover* command
 - C. The *esentutil /r* command
 - D. The *esentutil /g* command

Self Test Quick Answer Key

For complete questions, answers, and explanations to the Self Test questions in this chapter as well as the other chapters in this book, see the Self Test Appendix.

- | | |
|-------------------------|-----------------------|
| 1. C | 9. B |
| 2. A, B, D | 10. D |
| 3. B, C, E | 11. A, B |
| 4. A, C | 12. A, B, C, D |
| 5. B, D | 13. A, B, C, D |
| 6. B | 14. A, B, D |
| 7. A, B, C, D, E | 15. B, C |
| 8. A, B | |

MCSA/MCSE 70-294

Self Test Questions, Answers, and Explanations

This appendix provides complete Self Test Questions, Answers, and Explanations for each chapter.

Chapter 1: Active Directory Infrastructure Overview

Introducing Directory Services

1. An employee has retired from the company, and you have just disabled his account so no one can log on to the domain as this user. When this change is made, where will it be stored in the directory?
 - A. Domain partition
 - B. Configuration partition
 - C. Schema partition
 - D. Application partition

A. The domain partition contains information about the domain, inclusive to information on users, computers, resources, and attributes associated with each.

Answer B is incorrect, because the configuration partition stores information on the topology of Active Directory, and includes information on how the domains, domain trees, and forests within a network are configured. The configuration partition also includes information about the locations of DCs and the GC, which is a subset of the data contained in Active Directory. **Answer C** is incorrect, because the schema partition contains information on the object classes that exist in Active Directory, and which attributes each has. The schema partition contains information that defines object classes and attributes used within the domain. It determines what types of objects can exist within Active Directory, and what attributes each object can have. **Answer D** is incorrect, because the application partition is used to store data that is specific to different applications running on the network
2. Your company's employees are represented by two unions. Management has a union that represents the managers' interests, while others in the company belong to another union. Each union requires that dues be deducted from paychecks to pay for their representation. The Finance department has requested that a field be added to each user account, so that a code can be entered on the account to show which union each employee belongs to. They have asked you to create this field. When this new attribute has been added to user objects, where will it be stored in the directory?
 - A. Domain partition
 - B. Configuration partition
 - C. Schema partition
 - D. Application partition

- C.** The schema partition stores information on object classes and the attributes for each class. Each of the objects in Active Directory has specific attributes that determine the information that is stored about them. The objects that exist in Active Directory, and which attributes each has, is determined by the schema. The schema partition contains information that defines object classes and attributes used within the domain.
- Answer **A** is incorrect, because the domain partition contains information about the domain. It stores information on objects, but isn't used to store rules on what information can be stored in objects. Answer **B** is incorrect, because the configuration partition contains data on the topology of Active Directory, and includes information on how the domains, domain trees, and forests within a network are configured, as well as information on the locations of DCs and the GC. Answer **D** is incorrect, because the application partition stores data that is specific to different applications running on the network. Programs can use this partition to store settings that are needed while a particular application is running on a server.
3. You perform a search of Active Directory over the network, in search of an object stored in the directory. In performing this search, what protocol will be used?
- A. IPX/SPX
B. Directory Access Protocol (DAP)
C. Lightweight Directory Access Protocol (LDAP)
D. X.500
- C.** Active Directory uses LDAP for communications between clients and directory servers.
- Answer **A** is incorrect, because IPX/SPX is a protocol that isn't used for searching Active Directory. It is a protocol that is largely used on older versions of Novell NetWare, or Windows networks that need to communicate with NetWare servers. Answer **B** is incorrect, because DAP is the full version of LDAP. DAP contains more code than LDAP, and has too many features to be efficient with personal computers, thin clients, and communication over the Internet. Answer **D** is incorrect, because X.500 is a standard and not a protocol. X.500 was developed by the International Telecommunication Union and published by the International Organization for Standardization, and defines directory services. DAP is one of the protocols defined by X.500.
4. A user has the username JohnB. He wants to access a Microsoft Access database called db.mdb that's located on a DC called syngress.com, in a directory called DB. Using the URL, what will this user enter into a browser to access the database?
- A. JohnB@syngress.com
B. syngress.com
C. http://syngress.com/DB/db.mdb

D. \\syngress.com\DB\db.mdb

C. This choice uses the URL to access the file. The URL allows you to access files, applications, and other resources over networks.

Answer **A** is incorrect, because this is a UPN. In Active Directory, each user account in Active Directory has a logon name and UPN suffix. The logon name is the account name, and the UPN suffix is the domain that the user will log on to. Answer **B** is incorrect, because this is a DNS domain name. It identifies the domain, and doesn't provide enough information to access the database. Answer **D** is also incorrect, because this uses the UNC path to identify the location of the file. While the UNC name can be used to access the file, it does not use the URL that's specifically asked for in the question.

5. A user with the username of JaneD works in the Sales department. Her account is located in the syngress.com domain. Based on this information, what canonical name would be used to identify this object in Active Directory?

A. CN=JaneD

B. /CN=JaneD /OU=Sales /DC=syngress /DC=com

C. OU=Sales

D. /syngress.com/Sales/JaneD

D. A canonical name is another way of showing the distinguished name (DN) of an object, but in a format that is easier to read. Information used in the DN is reversed, so that it begins with the highest level and works its way down to the object's name.

Answer **A** is incorrect, because this is the relative distinguished name (RDN). An RDN is a portion of the DN, and is used to uniquely identify an object with a parent container. As each object must have a unique name with the directory structure, the RDN identifies an object within a particular OU. Answer **B** is incorrect because it is the DN. It is comprised of the object's common name (CN), organizational unit (OU), and domain component (DC). Answer **C** is incorrect, because this is the organizational component of the DN.

Understanding How Active Directory Works

6. You are making changes to object classes and attributes used in Active Directory. On which of the following DCs will you make these changes?

A. Schema Master

B. RID Master

C. Infrastructure Master

D. PDC Emulator

- A.** The Schema Master is a DC that is in charge of all changes to the schema, which defines what object classes and attributes are used in the forest. Updates to the schema must be performed on the DC acting in the role of Schema Master. This DC is used to write to the directory's schema, which is then replicated to other DCs in the forest.
- Answer **B** is incorrect, because the RID Master is responsible for creating a unique number for every object in a domain. These numbers are issued to other DCs in the domain, so that when an object is created, a sequence of numbers that uniquely identifies the object can be applied to it. Answer **C** is incorrect, because the Infrastructure Master is responsible for updating changes that are made to group memberships. Answer **D** is incorrect, because the PDC Emulator is designed to act as a Windows NT primary DC and receives preferred password changes, which it replicates to other DCs in the domain.
7. Your network consists of two forests, with two domains in one forest and three domains in the other. Based on this information, how many of the following master roles will be in the forests and domains?
- A. There will be five Schema Masters, Domain Naming Masters, RID Masters, PDC Emulators, and Infrastructure Masters.
- B. There will be two Schema Masters, Domain Naming Masters, RID Masters, PDC Emulators, and Infrastructure Masters.
- C. There will be five Schema Masters and Domain Naming Masters, and two RID Masters, PDC Emulators, and Infrastructure Masters.
- D. There will be two Schema Masters and Domain Naming Masters, and five RID Masters, PDC Emulators, and Infrastructure Masters.
- D.** The Schema Master and Domain Naming Master are forestwide master roles, and are unique to one DC in every forest. The RID Master, PDC Emulator, and Infrastructure Master are domainwide master roles, and there is only one per DC in each domain.
- Answer **A** is incorrect because there can only be one Schema Master and Domain Naming Master in each forest. Since there are two forests, this means there would be two of each on the network. Answer **B** is incorrect, because there must be one RID Master, PDC Emulator, and Infrastructure Master per domain. Since there are five domains, this means that there are five of each on the network. Answer **C** is incorrect, because the number of domainwide master roles and forestwide master roles needed for the network are inverted. There must be two Schema Masters and Domain Naming Masters, and five RID Masters, PDC Emulators, and Infrastructure Masters.
8. A user recently changed her last name, and you make changes to the user object in the directory to reflect this. Just before the change, inter-site replication has taken place using the default schedule. Just after the change, a link between the DC on which the changes

were made and the DC in the other site fails. It will be another hour until the link is back up again. There are four DCs in each site. Which of the following will occur?

- A. Replication between the DCs will occur normally, because at least two connections to each DC are created by the Knowledge Consistency Checker (KCC). Because one has failed, the other connection will be used.
 - B. Replication between the DCs won't occur. After 15 seconds, a notification of the change will be sent out, and replication partners will then request updated data.
 - C. Replication will occur normally, because the information won't be replicated until three hours after the last replication.
 - D. Another link will be used to replicate the data, based on the information gathered by the topology generator.
- C.** The scenario states that replication occurs by using the default schedule, and the link was back up one hour after the failure. By default, inter-site replication occurs every 180 minutes (three hours), and will use the site link to meet this schedule 24 hours a day, 7 days a week.
- Answer A** is incorrect, because this choice describes intra-site replication. With intra-site replication (where replication occurs within a site), at least two connections to each DC are created by the KCC, so that if one connection fails, the other connection can be used. **Answer B** is incorrect, because this also describes intra-site replication. By default, when a change is made on a DC, it will wait 15 seconds and then send notification to its closest replication partner. If it has more than one replication partner, it will send out notifications in three-second intervals to each additional partner. When the partner receives this notification, it will send out a request for updated directory information to the original DC, which then responds by sending the updated data. **Answer D** is incorrect, because replication occurs every three hours by default. Although the inter-site topology generator at each site will build a topology to send the data to other DCs, intra-site replication will occur at scheduled times. Information on site link objects is used to determine the best link to use for passing this data between sites.

Using Active Directory Administrative Tools

9. You are using the Microsoft Management Console (MMC) to administer objects in Active Directory. You decide to view information about a DC. Which of the following snap-ins will you use to view this information?
 - A. Active Directory Users and Computers
 - B. Active Directory Domains and Trusts
 - C. Active Directory Sites and Services
 - D. Dcgpofix

- A.** The Active Directory Users and Computers console is an MMC snap-in that allows you to administer user and computer accounts, domain controller objects, groups, printers, OUs, contacts, and other objects stored in Active Directory. In this snap-in, the Domain Controllers container contains objects representing Windows 2000 and Windows Server 2003 DCs that reside in the domain.
- Answer **B** is incorrect, because the Active Directory Domains and Trusts console is used to manage domains and the trust relationships between them. Answer **C** is incorrect, because the Active Directory Sites and Services console is used to create and manage sites, and control how the directory is replicated within a site and between sites. Answer **D** is incorrect, because *Dcgpofix* is a command-line utility that's used to restore the default domain policy and default Domain Controllers policy to the way they were when initially created.
10. Your company has merged with another company that uses UNIX machines as their servers. Users in your Windows Server 2003 domain need to access information on these UNIX machines, but you don't want to have information accessed by clients outside your domain. Which of the following types of trusts will you create to make it possible to share information in this way?
- A. One-way transitive forest trust
- B. Two-way transitive realm trust
- C. One-way nontransitive realm trust
- D. External trust
- C.** A realm trust is used when a relationship needs to be created between a Windows Server 2003 domain and a non-Windows realm that uses Kerberos version 5, such as one running UNIX. A realm trust can be one or two ways, and can be either transitive or nontransitive.
- Answer **A** is incorrect, because a forest trust is used to connect two different forests, so users in each forest can use resources in the other. Because there is only one Windows Server 2003 forest, this type of trust can't be used. Answer **B** is incorrect, because although a realm trust is used to connect these networks together, the scenario states that a one-way trust is needed. Answer **D** is incorrect, because an external trust is used to create a relationship between a Windows Server 2003 domain and one running Windows NT 4.0, or it can be used to connect two domains that are in different forests and don't have a forest trust connecting them.

Implementing Active Directory Security and Access Control

11. There is a concern about someone accessing objects in Active Directory using someone else's account. You want to audit an object to view how users and groups are accessing it. In setting auditing on this object, where will information on what to track be stored?
- A. Discretionary access control list
 - B. Security access control list
 - C. X.509
 - D. Auditing isn't provided on objects
- B.** The security access control list (SACL) is used to track an object's security based on how a user or group accesses the object. Information on what to audit is kept in access control entries (ACEs), which are stored within the SACL. These entries control what is audited, and contain information on the events to be logged. With auditing, records can be kept on the security of objects, and whether specific users or groups are able to successfully access them.
- Answer **A** is incorrect, because a DACL is a listing of ACEs for users and groups, and includes information on the permissions that they have to a file. The DACL controls whether a user is granted or denied access to an object. Answer **C** is incorrect, because X.509 is a standard that specifies the syntax and format of digital certificates. Answer **D** is incorrect, because auditing can be performed on objects.
12. You are configuring permissions on Active Directory so that managers can modify the user objects in the OU representing the department each manager is in charge of. In configuring these permissions, you also want each manager to have the ability to create new OUs within the OU representing his or her department. You want to give the most restrictive permissions to achieve these tasks. What permissions will you give these managers? (Choose all that apply.)
- A. Read
 - B. Write
 - C. Create All Child Objects
 - D. Delete All Child Objects
- A, B, and C.** Read allows the user to view objects, attributes, ownership, and permissions on an object. Write allows the user to change attributes on an object. Create All Child Objects allows the user to add objects to an OU. In giving these permissions, the managers can achieve all of the tasks required in the scenario.
- Answer **D** is incorrect, because Delete All Child Objects allows the user to delete objects from an OU. The scenario doesn't state that you are to give managers this capability.

13. You have set permissions on a parent container, and want to prevent these permissions from being applied to a child container within it. How will you achieve this?
- A. In **Active Directory Users and Computers**, open the properties of the parent OU, and select the **Security** tab. Click the **Advanced** button, and when the dialog box appears, ensure that the **Allow Inheritable Permissions From Parent To Propagate To This Object** check box is checked.
 - B. In **Active Directory Users and Computers**, open the properties of the parent OU, and select the **Security** tab. Click the **Advanced** button, and when the dialog box appears, ensure that the **Allow Inheritable Permissions From Parent To Propagate To This Object** check box is cleared.
 - C. In **Active Directory Sites and Services**, open the properties of the parent OU, and select the **Security** tab. Click the **Advanced** button, and when the dialog box appears, ensure that the **Allow Inheritable Permissions From Parent To Propagate To This Object** check box is cleared.
 - D. The objective cannot be achieved. Permissions will always be inherited by child objects. You must move the OU so it is at the same level in the hierarchy as the parent container.
- B.** In **Active Directory Users and Computers**, open the properties of the parent organizational unit, and select the **Security** tab. Click the **Advanced** button, and when the dialog box appears, ensure the **Allow Inheritable Permissions From Parent To Propagate To This Object** checkbox is cleared. When this is done, any permissions that are modified on parent objects don't apply to the child. Permissions for the child object must be explicitly assigned.
- Answer **A** is incorrect, because when the Allow Inheritable Permissions From Parent To Propagate To This Object check box is checked, object inheritance is enabled. This means that the permissions from parent OUs will apply to child OUs. Answer **C** is incorrect, because Active Directory Sites and Services isn't used to set permissions on objects. It is used to create and manage sites, and control how the directory is replicated within a site and between sites. Answer **D** is incorrect, because permissions can be blocked on child OUs.

What's New in Windows Server 2003 Active Directory?

14. You are upgrading your existing network to use Windows Server 2003. The network has Windows NT 4.0 domain controllers and the Windows Server 2003 server you're adding to the domain. After adding the first Windows Server 2003 DC to the network, you want to raise the domain functional levels to the highest level available for your network. To which level will you raise the domain?

- A. Windows 2000 mixed
 - B. Windows 2000 native
 - C. Windows 2003 interim
 - D. Windows 2003
- C.** The Windows 2003 interim level is used when your domain consists of Windows NT BDCs and Windows Server 2003 DCs. It provides the same functionality as Windows 2000 mixed mode, but is used when you are upgrading Windows NT domains directly to Windows Server 2003.
- Answer **A** is incorrect, because this level is used to accommodate Windows 2000 and Windows NT servers. Answer **B** is incorrect, because this level is used when there are no Windows NT BDCs in the domain. Answer **D** is incorrect, because this level is used when there are only Windows Server 2003 DCs in the domain.
15. You are upgrading your domain to use a mix of Windows 2000 and Windows Server 2003 DCs. After installing the first Windows Server 2003 DC on the domain, you want to raise the domain functional level to the highest level possible. Which of the following will you choose?
- A. Windows 2000 mixed
 - B. Windows 2000 native
 - C. Windows 2003 interim
 - D. Windows 2003
- B.** Windows 2000 native is used for domains with Windows 2000 and Windows Server 2003 DCs. Windows 2000 native removes support for replication to Windows NT BDCs, so these older servers are unable to function as DCs. In this level, only Windows 2000 and Windows Server 2003 DCs can be used in the domain.
- Answer **A** is incorrect, because Windows 2000 mixed allows domains to contain Windows NT BDCs and interact with Windows 2000 and Windows Server 2003 servers. Because it only provides basic Active Directory features, this isn't the highest level you could choose. Answer **C** is incorrect, because Windows 2003 interim is used when your domain consists of Windows NT and Windows Server 2003 DCs. Answer **D** is incorrect, because the Windows 2003 level is used when there are only Windows Server 2003 DCs in the domain.

Chapter 2: Working with User, Group, and Computer Accounts

Understanding Active Directory Security Principal Accounts

1. You create a new user account and assign it permissions to resources. When this account is created, a SID is given to the account to uniquely identify it. When the user logs on and attempts to access one of these resources, which of the following will the SID be compared to when determining access?
 - A. Access token
 - B. SACL
 - C. DACL
 - D. SID

C. The discretionary access control list (DACL) determines whether a security principal is granted or denied access to a resource. When a user tries to access an object, the user's SID is compared to entries in the DACL. If the user's SID or the SID of a group he or she belongs to matches an entry in the DACL, the user is permitted to use the resource.

Answer **A** is incorrect, because an access token is created when the user logs on, and contains information about the user's name, group affiliation, SID, and SIDs for the groups of which he or she is a member. The SIDs contained in the access token are compared to the SIDs listed in the DACL. Answer **B** is incorrect, because the system access control list (SACL) is used for auditing access to a resource. The SACL contains information on whether logging should be generated on attempts to access a resource. Answer **D** is incorrect, because SID is an acronym for security identifier.
2. A user attempts to access a resource, and entries in the ACL are analyzed to match the SID. The system examines the entire ACL, but no match is found. Which of the following will occur?
 - A. The user will be denied access.
 - B. The user will be granted access.
 - C. The account will be disabled.
 - D. Each ACE in the ACL will be read until a match is eventually found.

A. Since no entry is found, the security principal is implicitly denied access. Each ACE in the DACL is analyzed in sequence to determine a security principal's access. The SID of the user and any groups he or she belongs to is compared to the ACEs in the DACL.

Windows Server 2003 will look at each ACE until an entry is found that explicitly denies access to the resource, one or more entries are found that explicitly grants access to the resource, or until the entire DACL is searched but no ACE is found that explicitly grants or denies access.

- Answer **B** is incorrect; if no entry is found in the ACL, the user is implicitly denied access. If the user was granted access when no entry was found, then anyone who hadn't been given permission to the resource would have access. Answer **C** is incorrect, because accounts aren't disabled when they are denied access to a resource. Answer **D** is incorrect, because if the ACL were reviewed until an ACE matching the SID was found, it would never stop looking if no match existed.
3. A RID server has temporarily gone offline. During this time, you seize the RID Master role on another DC. After the original RID server becomes available again, you are concerned that duplicate SIDs might now exist for objects in Active Directory. Which of the following tools would you use to find and delete duplicates?
- A. Active Directory Users and Computers
 - B. MOVETREE
 - C. WHOAMI
 - D. NTDSUTIL
- D**. NTDSUTIL is a tool that runs from the command line, and has the ability to manage SIDs. Using this tool, you can find security principals with duplicate SIDs and delete them.
- Answer **A** is incorrect, because although Active Directory Users and Computers can be used to create and manage security principals, it doesn't have the functionality to find and remove duplicate SIDs. Answer **B** is incorrect, because MOVETREE is used to move objects to other locations in Active Directory. Answer **C** is incorrect, because WHOAMI is used to view information about the user who is currently logged on. While it can view SIDs associated with the user, it can't be used to detect and delete duplicate SIDs.

Working with Active Directory User Accounts

4. You want to use Remote Assistance to help users with problems by connecting to their machine and taking control of it remotely. When this action is performed, which of the following accounts is automatically created and used?
- A. HelpAssistant
 - B. Support_388945a0

- C. Guest
- D. InetOrgPerson
- A.** The HelpAssistant account is automatically created in Active Directory when a Remote Assistance session is established. Remote Assistance allows a user or administrator to connect to the desktop on other users' machines and assist them by talking them through a problem or taking control of their computer remotely.
- Answer **B** is incorrect, because the Support_388945a0 account is used to control access to scripts that use the account. The account allows users who don't have administrator access on the computer to perform tasks that require this access through signed scripts in Help and Support Services. Answer **C** is incorrect, because the Guest account is used by users who don't have an account of their own, or want to connect to the network without using their own account. Answer **D** is incorrect, because InetOrgPerson is a class that's used to define a specific type of user account. It isn't an actual account that's used for authentication and accessing resources.
5. Your network consists of an Active Directory domain with DCs running Windows Server 2003 and another network running Novell NetWare. You are preparing to migrate the Novell NetWare network to Windows Server 2003, but want to use an account that will connect the directory services of these two networks together. Which of the following class of user account will you create?
- A. HelpAssistant
- B. Support_388945a0
- C. InetOrgPerson
- D. None. A regular user account should be created.
- C.** InetOrgPerson is an object class used to create user accounts. This type of account is used to represent users in non-Microsoft directory services. The InetOrgPerson is used to allow applications written for other directories, or when migrating from these directory services to Active Directory.
- Answer **A** is incorrect, because the HelpAssistant is a user account that's used for Remote Assistance sessions. Answer **B** is incorrect, because the Support_388945a0 user account is used to control access to signed scripts. Answer **D** is incorrect, because regular user accounts aren't used for this purpose.
6. You are configuring a user account to use Terminal Services. Which of the following tabs on the user's account would you use to configure this user?
- A. General, Address, Organization
- B. Terminal Services Profile, Profile, Account
- C. Environment, Sessions, Remote Control
- D. Published Certificates, Member Of, Object

- C.** Each of these tabs is used to configure Terminal Services settings for the user account. The Environment tab contains settings for configuring the startup environment, Sessions contains timeout and reconnection settings, and Remote Control configures remote control settings for Terminal Services.
- Answer **A** is incorrect, because these tabs are used to configure personal properties regarding the user. The General tab contains basic information about the account, the Address tab provides contact information (such as the user's postal address), and the Organization tab contains information on the user's organization, department, and manager. Answer **B** is incorrect, because not all of these tabs deal with Terminal Services. Terminal Services Profile contains configuration information for the user's Terminal Services profile and home directory, but the other two have nothing to do with Terminal Services. The Profile tab contains settings on the user's logon script, profile, and home directory, while the Account tab contains settings that control account expiration, logon hours, and other account options. Answer **D** is incorrect, because the Published Certificates tab lists X.509 certificates and what they're used for, the Member Of tab lists the groups the user is a member of, and the Object tab provides information about the user object.

Working with Active Directory Group Accounts

7. You are creating a new group in Active Directory. In creating this group, you want users to be able to send e-mail to the group so that all members receive a copy of the message. Which type of group could be used for this purpose?
 - A. Security
 - B. Distribution
 - C. Both security and distribution
 - D. Neither security nor distribution
- C.** Security and distribution groups are e-mail entities, and allow e-mail to be sent to the group so that all members receive a copy. Distribution groups are used by applications like Microsoft Exchange to send e-mails to collections of users. Security groups are primarily used for access control, but can be given an e-mail address that will allow members to receive e-mail messages as a unit.
 - Answer **A** is incorrect, because distribution groups can also be used to send e-mail to collections of users. Answer **B** is incorrect, because security groups also serve as an e-mail entity, and can be used to send e-mail to collections of users. Answer **D** is incorrect, because both security and distribution groups can be used as entities to send bulk e-mail to users.

8. You created a new domain using DCs that are all running Windows Server 2003. The domain is part of a forest consisting of the domain you belong to, and three other domains. Each of these three other domains uses a Windows 2000 native functional level. The domain you belong to is running at the default domain functional level, and Active Directory has been configured so that all users in the domain have their own account. When adding users and groups to the groups you created, you decide that you want to change the scope of the Accounting and Sales groups. Which of the following must be taken into account when changes are made to these groups? (Choose all that apply.)
- A. If the group has a domain local scope, it cannot contain universal groups.
 - B. Domain local groups can be converted to universal groups.
 - C. Global groups can be converted to universal groups.
 - D. None of the groups in the domain can be universal groups.
- A and D.** If the group has a domain local scope, it cannot contain universal groups, and none of the groups in the domain can be universal groups. When a domain has a functional level of Windows 2000 mixed, domain local groups cannot contain universal groups in its membership. They can only contain accounts and global groups. Universal groups can only be created in domains with a domain functional level of Windows 2000 native or higher.
- Answers **B** and **C** are incorrect, because universal groups can't be created in domains using a Windows 2000 mixed functional level. If the domain functional level was set to Windows 2000 or higher, then domain local and global groups could be converted to universal groups.
9. Your network consists of several domains in a forest that has been set to Windows Server 2003 forest functionality. You are preparing to create a group that will contain user accounts from this domain and other domains, and will be used to access resources located in several of these domains. What will be the scope of the group you create?
- A. Universal
 - B. Global
 - C. Domain local
 - D. Distribution
- A.** Universal security groups can contain accounts and groups from any domain in the forest, and can be assigned permissions to resources in any domain in the forest that has the correct functional level.
- Answer **B** is incorrect. Although global groups can be assigned permissions in any domain, they can only contain accounts and groups from the domain in which they were created. Answer **C** is incorrect, because although domain local groups can contain accounts and groups from any domain, they can only be assigned permissions for

resources within the domain in which they were created. Answer **D** is incorrect, because distribution groups aren't used for access control and can't be assigned permissions.

10. You are an administrator in the domain dev.knightware.ca, which is a child domain beneath the forest root domain knightware.ca. You want to provide a user in this domain with the ability to create a forest trust between the dev.knightware.ca domain and the domain bookworms.ca. Which of the following built-in groups would you add this user to so he can create such a trust?
- A. Incoming Forest Trust Builders
 - B. Administrator
 - C. Account Operators
 - D. None of the above
- D.** None of the above. Forest trusts are created between forest-level domains. Because the dev.knightware.ca isn't a forest-level domain, such a trust can't be created regardless of which group the user is added to.
- Answer **A** is incorrect, because the Incoming Forest Trust Builders group is only available in forest root domains. Because this domain isn't a forest root domain, the user can't create a one-way incoming trust. Answer **B** is incorrect, because forest trusts can only be made between forest root domains. As such, the Administrator can't create such a trust from this child domain. Answer **C** is incorrect, because the Account Operators group doesn't have the necessary rights to create a forest trust.

Working with Active Directory Computer Accounts

11. You want a new member of the IT staff to be able to create new computer accounts using Active Directory Users and Computers. For this person to be able to perform this task, which of the following groups has the appropriate rights to create a computer account? (Choose all that apply.)
- A. Backup Operators
 - B. Account Operators
 - C. Domain Admins
 - D. Domain Users
- B** and **C.** To create a new computer account, a person needs to be a member of either the Account Operators group, Domain Admins group, Enterprise Admins group, or a user or group that has been delegated authority to create a new account.

- Answer **A** is incorrect, because the Backup Operators group gives a person the right to back up and restore data. Answer **D** is incorrect, because the Domain Users group contains all users in a domain, and provides them with the rights commonly needed by basic users.
12. You have given a user the **Add workstations to a domain** right, so he can have his computer join the domain. In giving the user this right, how many computer accounts can the user create?
- A. 1
 - B. 10
 - C. Unlimited
 - D. None
- B.** If a user has been issued the **Add workstations to a domain** right, he or she has the ability to configure computers to join the domain. When adding computers to the domain, the user can create up to 10 computer accounts.
- Answer **A** is incorrect, because a user can create up to 10 computer accounts in a domain. Answer **C** is incorrect, because for a user to create an unlimited number of computer accounts, he or she would need to be a member of the Account Operators group, Domain Admins group, Enterprise Admins group, or be granted additional special permissions, specifically the ability to Create Computer objects in the AD computers containers. Answer **D** is incorrect, because a user can create up to 10 computer accounts in a domain.
13. A new computer account is created in Active Directory Users and Computers for a workstation running Windows 2000 Professional. When viewing its properties, you check the Member Of tab and see that it is already included in the membership of a group. Which of the following groups is this account a member of?
- A. Domain Users
 - B. Domain Computers
 - C. Domain Controllers
 - D. Enterprise Admins
- B.** When a computer account is created, it is automatically made a member of the Domain Computers group. It is not a member of other groups, although you can add it to the membership of additional groups through the Member Of tab in the computer account's properties.
- Answer **A** is incorrect. This is the default group of which all user accounts are members. Answer **C** is incorrect, because the Domain Controllers group consists of all DCs within the domain. Answer **D** is incorrect, because the Enterprise Admins group is an

administration group with great control over domains in a forest. For this reason, accounts are not automatically added to it upon creation.

Managing Multiple Accounts

14. Your company has an external DNS domain that is used for the company's Web site, and an internal DNS domain that is used for the network. The external DNS domain is hosted on a UNIX server that hosts the company's Web site. The Web site using this external domain name is well known, and due to its popularity, users are confused as to which domain to log on to. Which of the following can you do to allow users to log on to the internal network using the external domain name?
- A. Use DSADD to add users to the external DNS domain.
 - B. Use MOVETREE to add users to the external DNS domain.
 - C. Use Active Directory Users and Computers to create an alternate UPN suffix.
 - D. Use Active Directory Domains and Trusts to create an alternate UPN suffix.
- D.** Use Active Directory Domains and Trusts to create an alternate UPN suffix. A UPN suffix is part of the upn that is used for logging on to a domain. The logon name indicates the account that's being logged on with. Alternative UPN suffixes can be created using Active Directory Domains and Trusts, so users can log on to their domain using a different UPN suffix. In this case, the alternate UPN suffix could be the same name as the external DNS domain, and can then be used by users to log on to the internal Active Directory domain.
- Answer **A** is incorrect, because DSADD is used to add new objects to Active Directory. Because the external DNS address is hosted on a UNIX server that isn't connected to the internal network, this tool cannot be used. Answer **B** is incorrect, because MOVETREE is used to move objects between different domains within a forest. It cannot be used for managing domain names that are hosted by UNIX servers. Answer **C** is incorrect, because Active Directory Users and Computers isn't used to create alternate UPN suffixes.
15. You want to move a user account from an OU located in one domain, to an OU located in another domain. Which of the following tools will you use to perform this task?
- A. Active Directory Users and Computers
 - B. Active Directory Domains and Trusts
 - C. DSMOVE
 - D. MOVETREE
- D.** MOVETREE is the Active Directory Object Manager, which is a command-line tool that allows you to move objects from one location in Active Directory to another. Using this tool, you can move an object from an OU in one domain to an OU located in another domain.

- Answers **A** and **C** are incorrect because both of these tools can only move objects within a domain. Active Directory Users and Computers is a graphical tool, while DSMOVE is a command-line tool. With both of these, you can move an object from one location in the directory tree to another location, but you can't move them to locations in other domains. Answer **B** is also incorrect, because Active Directory Domains and Trusts isn't used to move accounts within Active Directory.

Chapter 3: Creating User and Group Strategies

Creating a Password Policy for Domain Users

1. What is a potential drawback of creating a password policy on your network that requires user passwords to be 25 characters long?
 - A. Users will be more likely to write down a password that is so difficult to remember.
 - B. User passwords should be at least 30 characters long to guard against brute-force password attacks.
 - C. There are no drawbacks; this solution creates network passwords that will be impossible for an unauthorized user to penetrate.
 - D. Windows Server 2003 will not allow a password of more than eight characters.

A. A 25-character password is unreasonably long for most environments, and could prompt your users to write them down. Passwords that are written down can be discovered, especially if they are stored near the computer. This can easily render the use of strong passwords meaningless.

Answer **B** is incorrect. A password length of 8 to 14 characters is usually sufficient to guard against most brute-force attacks. Answer **C** is also incorrect. A 25-character password will create the issues described in Answer **A**. Answer **D** is incorrect because Windows passwords can be up to 255 characters in length.
2. You have recently started a new position as a network administrator for a Windows Server 2003 network. Shortly before the previous administrator left the company, the *syskey* utility was used on one of your domain controllers to create a password that needs to be entered when the machine is booted. You reboot the domain controller, only to discover that the password the previous administrator documented is incorrect. You are unable to contact your predecessor to obtain the correct one. How can you return this DC to service as quickly as possible?
 - A. Reformat the system drive on the server and reinstall Windows Server 2003.
 - B. Boot the server into Directory Services Restore Mode and restore the DC from a point before the previous administrator ran the *syskey* utility.

- C. Boot the server into Safe Mode and run *syskey* again to change the password.
- D. Use *ntdsutil* to seize the PDC Emulator role and transfer it to another DC.
- B.** If you misplace the password or diskette that's created when you run the *syskey* utility, your only option is to restore the system from a point before the *syskey* utility was run.
- Answer **A** is incorrect. It is not the quickest way to restore the DC to service, because you will lose application and Registry data stored on the system drive. All applications will need to be reinstalled and any shares recreated. Answer **C** is also incorrect. You cannot change the *syskey* password without knowing the original one. This is by design so that an attacker cannot circumvent *syskey* security by simply rebooting the server. Answer **D** is incorrect because transferring the PDC Emulator role will do nothing to return this DC to service.
3. According to Microsoft, which of the following would be considered weak passwords for a user account named jronick? (Choose all that apply.)
- A. S#n\$IUsN7
- B. soprano
- C. ronickrj
- D. Oo!dlx2
- E. new
- B, C, and E.** Microsoft considers a password weak if it is all lowercase, contains any portion of the user's account name (in this case, *jronick*), or contains a word found in the English dictionary (such as *soprano* or *new*); therefore, Answers **B, C, and E** are correct.
- Answers **A and D** are incorrect because both of these passwords meet the criteria for strong passwords. They are at least seven characters long and contain a mix of uppercase, lowercase, numeric, and non-alphanumeric characters.
4. You have implemented a password policy that requires your users to change their passwords every 30 days and retains their last three passwords in memory. While sitting in the lunch room, you hear someone advise his coworker that all she needs to do to get around that rule is to change her password four times so that she can go back to using the password that she is used to. What is the best way to modify your domain password policy to avoid this potential security liability?
- A. Increase the maximum password age from 30 days to 60 days.
- B. Enforce password complexity requirements for your domain users' passwords.
- C. Increase the minimum password age to seven days.
- D. Increase the minimum password length of your users' passwords.

- C.** If your password policy retains three unique passwords in memory, this setting can help prevent your users from changing their passwords four times in rapid succession so that they can change them back to their initial passwords. A minimum password age of seven days will force users to wait at least seven days before they can change their passwords. Very few users will want to be forced to use a new password every week for three weeks, just to get back to their favorite.
- Answer A** is incorrect. Increasing the maximum password age will not circumvent the security issue that is allowing users to change back to their favorite password. **Answer B** is also incorrect. Password complexity has nothing to do with how often a password can be changed. **Answer D** is incorrect because the minimum password length setting also has nothing to do with how often a password can be changed.
5. You are a new network administrator for a Windows Server 2003 domain. In making user support calls, you have noticed that many users are relying on simplistic passwords such as their children's or pets' names. Passwords on the network are set to never expire, so some users have been using these weak passwords for years. You change the default Group Policy to require strong passwords. Several weeks later, you notice that the network users are still able to log on using their weak passwords. What is the most likely reason why the weak passwords are still in effect?
- A. You must force the users to change their passwords before the strong password settings will take effect.
- B. The Group Policy settings have not replicated throughout the network yet.
- C. Password policies need to be set at the organizational unit (OU) level, not the domain level.
- D. The users reverted back to their passwords the next time they were prompted to change them.
- A.** The password complexity policy only applies to new or changed passwords within the domain. It is not applied retroactively to existing passwords. If your users' passwords are set to never expire, they will not be forced to change to strong passwords.
- Answer B** is incorrect. Active Directory replication should not take several weeks to replicate, even on the largest networks. **Answer C** is also incorrect. Password policies can only be set at the domain level, not on individual OUs. **Answer D** is incorrect because Windows would reject passwords that did not meet the new complexity requirements.

Creating User Authentication Strategies

6. You have created an e-commerce Web application that allows your customers to purchase your company's products via the Internet. Management is concerned that customers will not feel comfortable providing their credit card information over the Internet. What is the

most important step to secure this application so that your customers will feel confident that they are transmitting their information securely and to the correct Web site?

- A. Use IP restrictions so that only your customers' specific IP addresses can connect to the e-commerce application.
 - B. Issue each of your customers a smart card that they can use to authenticate to your e-commerce Web site.
 - C. Place your company's Web server behind a firewall to prevent unauthorized access to customer information.
 - D. Install a Secure Sockets Layer (SSL) certificate on your Web server.
- D.** Installing an SSL certificate provides *mutual authentication* so that your customers will know that they are communicating with the correct Web site and not being redirected to another site that's being used to steal their information.
- Answer **A** is impractical because your customers (and their associated IP addresses) will change from day to day as you get new referrals. Answer **B** is also incorrect. The costs associated with supporting smart cards for your customers in this environment could be quite high, which makes Answer **D** more appropriate. Answer **C**, although a good security practice, is incorrect because it will not protect your customers' data while it is being transmitted to and from your Web site. Protecting data during transit requires the kind of encryption offered by an SSL certificate.

7. Your network environment consists of Windows 2000 Professional, Windows XP Professional, and Windows NT 4.0 Workstation computers. You have just upgraded all domain controllers to Windows Server 2003. The domain and forest functional levels are both set to Windows Server 2003. The company does not use any Web applications or services. Which of the following authentication protocols will be used on the network? (Choose all that apply.)

- A. Digest
- B. NTLM
- C. Kerberos
- D. SSL

- B, C.** Answer **B** is correct. NTLM will be used for authentication involving the pre-Windows 2000 computers in the domain. In the scenario, this means that the Windows NT 4.0 Workstation computers will use NTLM for authentication. Answer **C** is also correct. Kerberos is the default protocol for all Windows 2000 and later computers. Therefore, the network's Windows 2000 Professional and Windows XP Professional clients will use Kerberos for authentication.

- Answers **A** and **D** are incorrect. Both Digest authentication and SSL are used in a Web services environment. The question states that there are no Web applications or services in use.
8. You've decided to implement Web-based authentication. You have a wide range of domains, domain controllers, and domain functional levels in your enterprise Windows Server 2003 forest. Because you are a homogenous Windows environment, you decide to implement digest authentication. Which of the following requirements must you keep in mind when planning to implement digest authentication? (Choose all that apply.)
- A. Digest authentication requires IE 5 or later on the clients.
 - B. There must be at least one Windows Server 2003 DC in the IIS server's domain.
 - C. User passwords must be stored with reverse encryption.
 - D. There must be at least one Windows 2000 or later DC in the IIS server's domain.
- A, C, D.** Answer **A** is correct. Clients must be using Internet Explorer 5 or later to use digest authentication. Answer **C** is also correct. Digest authentication requires user passwords to be stored with reversible (cleartext) encryption. Answer **D** is correct. Digest authentication requires at least one Windows 2000 or later domain controller in the domain that the IIS server is a member of.
- Answer **B** is incorrect because digest authentication requires at least one Windows 2000 or later domain controller in the domain that the IIS server is a member of. The domain controller does not have to be a Windows Server 2003 DC.

Planning a Smart Card Authentication Strategy

9. Your network configuration includes a Terminal Server designed to allow users at remote branches to access network applications. The Terminal Server often becomes overloaded with client requests, and you have received several complaints regarding response times during peak hours. You have recently issued smart cards for the users located at your corporate headquarters and would like to prevent those users from using their smart cards to access the Terminal Server. How can you accomplish this goal in the most efficient manner possible?
- A. Enable auditing of logon/logoff events on your network to determine which smart card users are accessing the Terminal Server, and then speak to their supervisors individually.
 - B. Create a separate OU for your Terminal Server. Create a global group containing all smart card users, and restrict the logon hours of this group for the Terminal Server's OU.
 - C. Enable the "Do not allow smart card device redirection" setting within Group Policy.

- D. Create a global group containing all smart card users, and deny this group the “Log on locally” right to the computers on your network.
- C.** “Do not allow smart card device redirection” only allows the use of smart card credentials for local workstations and servers. Smart cards cannot be used to log on to a Terminal Services session which this policy is enabled.
- Answer **A** is incorrect. This solution requires too much administrative overhead and has no guarantee of being effective. Answer **B** is also incorrect. Account policies such as logon hours can only be set at the domain level, not at the OU level. Answer **D** is also incorrect, because this will prevent smart card users from logging on to *any* machine on your network, not just the Terminal Server.
10. You have attached a smart card reader to your Windows XP Professional workstation’s serial port. The reader is not detected when you plug it in and is not recognized when you scan for new hardware within Device Manager. The smart card reader is listed on the Microsoft Web site as a supported device, and you have verified that all cables are connected properly. Why is your workstation refusing to recognize the smart card reader?
- A. The manufacturer-specific installation routine is not compatible with Windows Server 2003.
- B. The workstation needs to be rebooted before it will recognize the card reader.
- C. Smart card readers are only supported on machines running Windows Server 2003.
- D. You are not logged on as a member of the Domain Admins group.
- B.** If the smart card reader attaches via a serial port, the workstation might need to be rebooted before Windows Server 2003 will recognize the new hardware.
- Answer **A** is incorrect. Even with a manufacturer installation routine, it will probably still be necessary to reboot the computer so that it will recognize the device. Answer **C** is also incorrect. Smart card readers are supported on a wide variety of Windows platforms using third-party software; and Windows 2000, XP Professional, and Server 2003 using software built into Windows. Answer **D** is incorrect because a reboot would still be necessary if the device wasn’t being recognized.
11. You have recently deployed smart cards to your users for network authentication. You configured the Smart card Logon certificates to expire every six months. One of your smart card users has left the company without returning her smart card. You have disabled this user’s logon account, but management is concerned that she will still be able to use the smart card to access network resources. How can you be sure that the information stored on the former employee’s smart card cannot be used to continue to access network resources?
- A. Monitor the security logs to ensure that the former employee is not attempting to access network resources.

- B. Use the smart card enrollment station to delete the user's Smartcard Logon certificate.
 - C. Deny the Autoenroll permission to the user's account on the Smartcard Logon Certificate template.
 - D. Add the user's certificate to the CRL on your company's CA, and publish the CRL.
- D.** Every CA maintains a CRL that lists certificates that have been revoked and are no longer valid. Once the certificate is listed on the CRL, the Windows Server 2003 domain will not consider the certificate on the user's smart card valid.
- Answer **A** is incorrect. Although a security best practice, this answer takes no proactive actions to prevent the former employee from accessing network resources. Answer **B** is also incorrect. The user did not return her smart card, so it cannot be inserted into a smart card device on an enrollment station for modification. Answer **C** is incorrect because this will not disable the existing certificate that is stored on the user's smart card.

Planning a Security Group Strategy

12. One of your coworkers is trying to grasp the concept of distribution and security group types. He asks you what the two primary benefits are for the security group type. What do you tell him? (Choose two.)
- A. You tell him that they can have permissions and user rights assigned to them.
 - B. You tell him that they can function for messaging just like a distribution group type.
 - C. You tell him that they allow for quick and efficient delegation of administrative responsibility in Active Directory.
 - D. You tell him that they can only be used for messaging and granting permissions to Active Directory, file system, Registry, and printer objects.
- A, C.** Answer **A** is correct. The primary function of a security group is for assigning permissions and user rights. Answer **C** is also correct. Because they can be assigned permissions and user rights, security groups allow an administrator to delegate administrative responsibility in Active Directory. This delegation might involve granting a group full control of an OU, or simply the ability to change the passwords for user accounts that reside within it.
- Answer **B** is incorrect. Although this is a true statement, it is not a primary benefit of the security group type. In fact, Microsoft does not recommend using security groups for messaging. Answer **D** is also incorrect, because security groups can also be used to assign user rights.

13. Your boss has been looking over marketing material from Microsoft. She asks you how you plan on using universal groups. You administer a single domain environment that is about to be upgraded to Windows Server 2003. What do you tell her?
- A. You tell her that because you will be using a Windows Server 2003 functional level domain, you will be using only universal groups.
 - B. You tell her that because you will be using a Windows 2000 native functional level domain, you will be using only universal groups.
 - C. You tell her that you will use universal groups to replace global groups, but will still be using domain local groups for resource access.
 - D. You tell her that you will not be using universal groups.
- D.** Because you are using a single domain, there is no need to use universal groups. Universal groups are designed to simplify administration in a multidomain environment.
- Answers **A** and **B** are incorrect. Although each of these functional levels activates universal security groups, because you are in a single domain environment they are unnecessary. Answer **C** is also incorrect. Microsoft's domain models always call for using domain local groups for resource access; however, universal groups should not be used in this environment for the reasons already stated.
14. Last night you finished configuring a complex set of groups for your new Windows Server 2003 Active Directory environment. You spent this morning adding users to their appropriate groups. Now that the Active Directory environment is configured, you are trying to add the groups into ACLs in the file system. For some reason, they aren't showing up in the list of groups to select from. You can see all the default groups that the operating system and Active Directory installed. Why can't you see the groups you created?
- A. You don't have permission.
 - B. You didn't activate the groups in Active Directory.
 - C. You created distribution groups.
 - D. You created security groups.
- C.** Distribution groups are designed to be used for messaging, not applying security in Windows Server 2003. Because of this, they are not available when configuring ACLs in the operating system.
- Answer **A** is incorrect. Because you can see and select the default system groups, you know that you have the appropriate permissions. Answer **B** is also incorrect. Groups are activated the minute you create them in Active Directory. Answer **D** is incorrect because if you had created security groups, they would be showing up with the default groups.

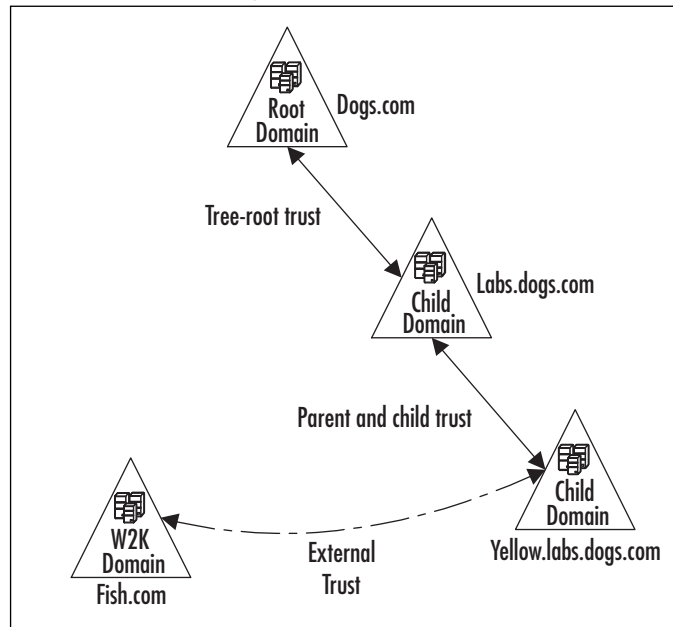
15. Your company has a single domain environment that will be upgraded to Windows Server 2003. One of the company's existing Windows NT 4.0 BDCs must remain in place because a custom application requires it. This application will not be migrated until sometime next year. The company has many departments, each of which has sub-departments and teams. The company would like to take advantage of Windows Server 2003's new group nesting capabilities. Which of the following group models is appropriate for this company?
- A. AGDLP
 - B. AGGDLP
 - C. AGGUDLP
 - D. AGUDLP
- A.** Because a Windows NT domain controller must remain on the network, only Answer **A** is correct. The company wants to take advantage of the additional group nesting capability in Windows Server 2003 but cannot until they upgrade all of their domain controllers to at least Windows 2000. Additional group nesting is only supported by the Windows 2000 native and Windows Server 2003 domain functional levels. These functional levels do not support Windows NT domain controllers.
- Answers **B**, **C**, and **D** are all incorrect. Although they allow for the use of group nesting, all require a higher domain functional level, which is not possible because of the Windows NT domain controller.

Chapter 4: Working with Forests and Domains

Understanding Forest and Domain Functionality

1. Your Yellow.labs.dogs.com Information Assurance department has just used John the Ripper, a password-cracking program, against your Active Directory user accounts. They report that 40 percent of your passwords were compromised within an eight-hour period. After some research, you determine that removing the LM authentication for down-level clients will make password guessing much more difficult. Further, you decide to require Kerberos authentication at all of your DCs. After some telephone calls, you assure your manager and CIO that all computers in all domains within your forest are loaded with Windows 2000, 2003, or XP. No legacy workstations are present that require backward LM or NTLM compatibility. With their permission, you apply the security setting using Group Policy. Almost immediately, authentication fails between Fish.com and the Yellow.labs.dogs.com domain, which are located in different AD forests and connected to each other via an external one-way trust. However, ping still works between hosts in the two domains. Previously, all users in both domains had complete access to all resources in the two domains, but now every authenticated access fails. Refer to Figure 4.56. Why is this problem happening?

Figure 4.56 Question #1 Diagram



- A. Network problems.
 - B. Selective authentication problems.
 - C. NT LAN Manager authentication problems.
 - D. Trust transitivity problems.
- C.** Answer C is correct because external trusts are limited in the same ways as an NT 4.0 trust. The NT LAN Manager establishes all *external* trusts. When Kerberos-only authentication is enabled, your two external trusts fail due to their reliance on NTLM.
- A** is incorrect, because basic network connectivity has not been compromised, as demonstrated by the successful *ping* command. Answer **B** is incorrect, because *selective authentication* requires explicit permissions between all users and all resources in the two domains. Previously full unrestricted access demonstrates that selective authentication was not configured. Answer **D** is incorrect, because as an external nontransitive trust, only Yellow.labs.dogs.com ever had access to Fish.com in the first place, and only Yellow.labs.dogs.com was affected by the outage. No transitivity issues were identified.
2. Over a weekend, you upgrade your Windows NT 4.0 domain through the Windows Server 2003 interim functional level up to the Windows Server 2003 functional level. Most client PCs are running Windows NT 4.0. Your test NT workstation continues to operate, but on Monday morning you start getting trouble tickets from several workstations that cannot log on to the new domain. As a temporary workaround, those users are able to log on from other workstations. What is the most likely cause?

- A. The malfunctioning workstations are running Windows NT 4.0 SP2.
 - B. The users' old passwords do not meet the new complexity requirements of Windows Server 2003.
 - C. Some workstations did not pick up the new trust relationships, and need to be removed and re-added to the domain.
 - D. Some user accounts have not replicated to all DCs yet. Choose the **Synchronize all domain controllers** option from the **Active Directory Domains and Trusts** MMC snap-in.
- A.** Windows NT 4.0 systems prior to SP3 were not compatible with NTLMV2 authentication, which is a minimum authentication requirement of Windows Server 2003.
- Answer **B** is incorrect because, although password complexity requirements are set by default in Windows Server 2003, the defined problem exists at the workstation, not the user account. Answer **C** is incorrect because trusts operate at the domain level, not at the workstation level. Answer **D** is incorrect, because replication is more robust and timely in Active Directory than it was in NT. This is unlikely to be the problem in a healthy LAN environment. Although it is possible that Active Directory replication has not yet completed after the installation, the stated problem lies with specific workstations, not user accounts.
3. A large company has just merged with yours. They have a corporate 24-hour Enterprise Administrators group running a forest at the Windows Server 2003 functional level. Your company has NT, 2000, and XP clients on a single NT domain. Your PDC and four of your five BDCs are still running Windows NT. Two of your permissions groups have more than 5000 members. The new conglomerate has given you a budget for upgrading the local infrastructure so that you can participate in the corporate forest structure through a cross-forest trust. What are the minimum requirements you must meet before joining the corporate forest? (Choose all that apply.)
- A. Ask the Enterprise Administrators to temporarily set their forest to interim level using manual LDAP administrative tools.
 - B. Install new servers and create a pristine local forest root using your namespace.
 - C. Upgrade your PDC to Windows Server 2003, and join the corporate forest during the upgrade.
 - D. Upgrade your four Windows NT BDCs to Windows Server 2003.
- A** and **C.** Answer **A** is correct because your best strategy is to use interim mode due to having large security groups. Microsoft generally recommends interim mode for NT-to-2003 direct upgrades. Because the corporate forest is already at the Windows Server 2003 functional level (federated forests can only be created at that level), they must temporarily drop to interim level for you to join during the upgrade. Answer **C** is cor-

rect because your local domain will inherit the necessary Windows Server 2003 interim functional level from the corporate forest at the time of joining. This is the step that installs Active directory in your domain.

- Answer **B** is incorrect because the requirement is to join an existing domain, not start up a new one. If you create your own forest, you will also have your own Enterprise Administrators group, which goes against the corporate centralized forest administration strategy. Answer **D** is incorrect because the Windows 2000 BDC must be either removed or upgraded for the use of interim level, not the NT BDCs. The NT BDCs must be upgraded before the project is complete so that the corporate forest can return to the Windows Server 2003 functional level (non-interim). It is not, however, a requirement for the joining process.
4. Two forests are joined with a cross-forest trust, linking the two for resource sharing. Both forests synchronize their time with the Internet, they both contain some Windows NT 4.0 workstations, and they reside in different time zones. One day, the cross-forest trust fails, and users cannot authenticate across their forest boundaries, although the network is healthy at layer 3. What might have happened?
- A. The two forests are using different time sources, or one forest-root PDC emulator has malfunctioned.
- B. The Windows NT systems need to reboot, or use the scheduler service to run the *net time* \\<timecomputer> /set /yes command on each NT machine.
- C. DCs in both forests should be set up to automatically run the *w32tm /config /syncfromflags:manual /manualpeerlist:Peerlist* command, listing only each other as their time source.
- D. The PDC Emulator needs its time reset, since Windows NT 4.0 workstations and servers automatically synchronize with their PDC.
- E. The scenario is impossible, since Windows NT cannot function in a Windows Server 2003 functional level domain, and cross-forest trusts can only be created in such a domain.
- A.** In a federated forest, individual enterprises can choose to attune with different external time sources. If these sources diverge for any reason, although each forest is chronologically homogenous, they might not agree with *each other*, resulting in a failure of all cross-forest authentication. Additionally, Windows 2003 and Windows XP systems use the W32Time service to synchronize with their authenticating DC. Each DC synchronizes with their domain's PDC Emulator, while each domain's PDC Emulator synchronizes with the forest-root PDC Emulator operations master, making that server the authoritative W32Time source for the entire enterprise. For that reason, this is often the server that is configured to connect to the external time source. The *net time* command will typically synchronize with some internal host, which in turn synchronizes with that

authority. If no additional fault tolerance is set up, the loss of the forest-root PDC Emulator for a long enough period will result in significant drift of the enterprise time.

- Answer **B** is incorrect because, although the *net time* command is a good idea in general, since NT will not otherwise synchronize, in this case it will only serve to more accurately synchronize the NT machines with the malfunctioning forest time. Answer **C** is incorrect because each enterprise should synchronize with an *external* authoritative time source, not with each other. Synchronizing with each other will solve the authentication problem, but it is a bad engineering practice since both forest times will tend to drift with each other and no longer be accurate with the real world. Answer **D** is incorrect because, even though the PDC Emulator operations master in each forest root functions as the authoritative time source for the W32Time service, Windows NT 4.0 systems do not use this service. With the *net time* command, you can use any server. Answer **E** is incorrect because, although all DCs in a Windows Server 2003 functional level forest must be running Windows Server 2003, member servers and workstations can continue to run Windows NT.
5. Your IT manager wants you to link four divisions of the company through a ring of eight unidirectional cross-forest trusts. He uses this reasoning: If multiple forest trusts are established, then authentication requests made in any domain of any forest can pass through multiple forest trusts, hence multiple Kerberos domains, on their way to their destination. Why is he wrong?
- A. While each cross-forest trust is transitive at the forest level, where all domains in both forests can authenticate, they are not transitive at the federated forest level as he suggests. The trust path cannot include more than one cross-forest trust.
- B. Cross-forest trusts are not transitive, and will not allow pass-through authentication.
- C. To create a ring of trust around four forests, you only need four cross-forest trusts.
- D. Cross-forest trusts are bidirectional, so only three trusts are needed to link all four forests. Completing the “ring” is not necessary.
- A.** A web of trusts to and from each forest would be needed to accomplish complete trust at the multiforest level.
- Answer **B** is incorrect because cross-forest trusts *are* transitive. Answers **C** and **D** are both incorrect because cross-forest trusts are not bidirectional.
6. You are the database administrator for a large enterprise, and you are trying to install an Active Directory-integrated application. The product literature says it implements new classes within the directory. The installation fails with an obscure message. What is the most likely reason for the failure?
- A. The application uses a service account that is not a member of the Enterprise Administrators group.


- B. You are not a member of the Domain Admins account in the domain where you are trying to install the application.
 - C. You are not in the Schema Administrators group.
 - D. You do not have permissions to create GPOs.
 - C.** By default, the ability to modify the schema is limited to those in the Schema Administrators group. Implementing new schema classes requires schema modification rights.
 - Answer **A** is incorrect because the authoritative context used for an installation is the user running the installation. Additionally, the Enterprise Administrators group does not have rights to modify the schema. Answer **B** is incorrect because Domain Admins cannot modify the schema. Answer **D** is incorrect because the problem does not lie with GPO creation rights.
7. What FSMO roles should exist in a child domain in a Windows Server 2003 forest? (Choose all that apply.)
- A. Schema Master
 - B. Domain Naming Master
 - C. PDC Emulator
 - D. RID Master
 - E. GC
 - F. Infrastructure Master
 - C, D, and F.** The RID Master, PDC Emulator, and Infrastructure Master all operate at the domain level.
 - Answers **A** and **B** are incorrect for the same reason. The Schema Master and the Domain Naming Master both operate at the forest level only, and only exist in the forest-root domain. Answer **E** is incorrect because the GC is not an FSMO role.

Creating the Forest and Domain Structure

8. Your company is expanding its single location with five new offices, all in different states. Each location will have 10 marketing employees or less to supplement the 25 already employed at the main office. Your security engineer says that all employees will abide by the same company security policy. They will hire another systems administrator at your office to handle the increased workload. When asked about how the company's single-domain Windows Server 2003 Active Directory will be affected by the expansion, you reply that new servers will have to be installed at the remote locations. Your manager wants to know what server hardware and software to budget for. What do you tell him? (Choose one.)

- A. Five servers and five copies of Windows Server 2003 Datacenter Edition
 - B. Five servers and five copies of Windows Server 2003 Standard
 - C. Ten servers and ten copies of Windows Server 2003 Enterprise Edition
 - D. Ten servers and ten copies of Windows Server 2003 Standard
- B.** Each remote location should be a new site in the company's existing domain. The remote sites will all stay within your current administrative, security, and organizational boundaries; therefore, you do not need any new domains. Since they will all be additional DCs in an existing domain, only one DC is needed at each site; the WAN link *and* the site's DC must both malfunction to stop users from logging on, eliminating the single point of failure. Doubling up on DCs at remote sites is not needed in that case. With the new Windows Server 2003 no-GC logon feature, a GC is not required at each site. Windows Server 2003 Standard is all that is needed for a DC.
- Answer **A** is incorrect because the Datacenter Edition is very expensive and overfeatured for the task. Answers **C** and **D** are incorrect because only five servers are needed. Again, the Enterprise Edition is an unnecessary additional expense for the DC role.
9. You have just been hired to install an Active Directory for a new startup company. An e-commerce Web farm has already been set up, waiting to join your root domain. Walleyedhuckle Mullys.com employs seven chemists who appear to split their time evenly between research and development activities, two Web developers, one network engineer, and the owner of the company. The previous MCSE who ordered the computers and software has already left the company, and all you have is a pile of boxes and a crude sketch of a logical Active Directory forest. The owner of the company has already approved the design and wants you to work all weekend to make it work. After you clean up the drawing somewhat, it looks like Figure 4.57. Among the boxes you find six computers, three copies of Windows Server 2003 Standard, and three copies of Windows Server 2003 Web Edition. Under the pressure of time, you create the forest structure with two computers in each of the three domains. On Monday morning, you approach your new manager to say that the Active Directory design is weak. What is wrong with it? (Choose all that apply.)

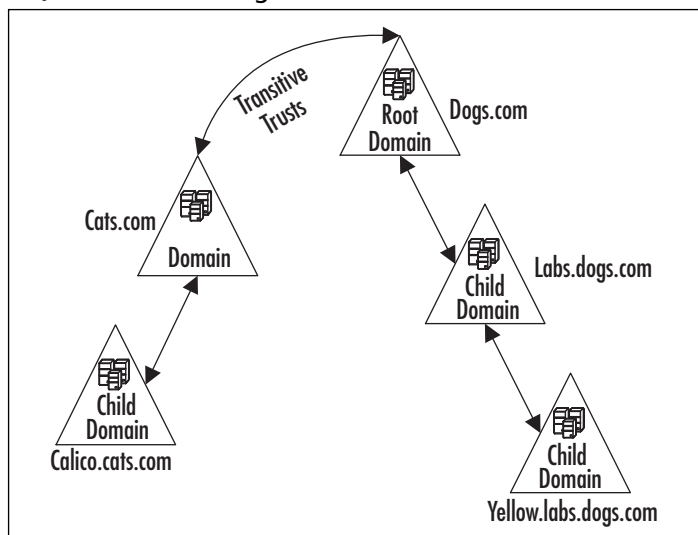
Figure 4.57 Question #9 Diagram



- A. DNS should be split, and the design is full of bad namespace and computer security practices.
 - B. Too many domains.
 - C. You should have installed Windows Server 2003 Enterprise instead of Standard.
 - D. Firewall is in the wrong place, because Web servers should be directly on the Internet.
 - E. Not enough DCs for the number of domains.
- A, B, and E.** Answer **A** is correct, because internal and external DNS zones can have the same name, but they should not share the same logical namespace. It is a bad security practice to place Internet-accessible hosts in your internal network. Answer **B** is correct because there appear to be no appreciable differences between the `Development.WHMs.com` domain and the `Research.development.WHMs.com` domain. The same people appear to use the resources in both domains equally, and there is only one administrator. Answer **E** is correct because Windows Server 2003 Web Edition is incapable of taking the role of DC, leaving only a single DC in each domain. A single DC is susceptible to catastrophic failure, leading to the loss of a domain or the structure of the entire forest.

- Answer **C** is incorrect because no Enterprise Edition services are being used. Windows Server 2003 Standard is more than sufficient to the task. Answer **D** is incorrect because the statement that your Web servers should be directly on the Internet is incorrect; they should be protected. Best practices call for a *second* firewall between the Web servers and your internal network, placing them in a DMZ.
10. The name of your company has changed from Fish to Aquatics. In her zeal to embrace the new company image, your CEO directs you to change the Active Directory domain name to reflect the new company name. Your efforts are unsuccessful. Which of the following would cause the domain rename tool to fail? (Choose all that apply.)
- A. Your DCs are running Windows Server 2000 Enterprise Edition.
 - B. The Active Directory forest functional level is set to Windows Server 2003 interim.
 - C. You are logged in as a Domain Admin.
 - D. Your current mail server is running Exchange 2000.
 - E. Your DFS root server is running Windows 2000 with Service Pack 3.
 - F. You are logged on to a Windows 2003 Server DC as an Enterprise Admin.
- A, B, C, D, and F.** Answer **A** is correct, because the domain rename tool will fail if any DC is running Windows Server 2000. Answer **B** is correct, because the domain rename tool requires the forest functional level to be Windows Server 2003, not interim. Answer **C** is correct, because the domain rename tool must be executed in the Enterprise Admins context. Answer **D** is correct, because the domain rename operation is not supported in an Active Directory forest that includes Exchange 2000. Answer **F** is correct, because the domain rename tool cannot be run from a DC.
- Answer **E** is incorrect, because all DFS root servers with domain-based DFS roots must be running Windows Server 2000 with Service Pack 3 or higher, or Windows Server 2003. Since the DFS root server meets these requirements, this is not a reason for the domain rename tool to fail.
11. Your network operations center has identified excessive bandwidth utilization caused by authentication traffic in the root domain subnet, especially between Calico.cats.com and Labs.dogs.com. Your logical network is set up as shown in Figure 4.58. What type of trust or trusts would you set up to alleviate the situation?

Figure 4.58 Question #11 Diagram



- A. Set up a bidirectional transitive parent and child trust between Calico.cats.com and Labs.dogs.com.
- B. Set up a shortcut trust between Calico.cats.com and the forest root, and set up a second shortcut trust between Labs.dogs.com and the forest root.
- C. Set up a shortcut trust between Calico.cats.com and Labs.dogs.com.
- D. Set up two shortcut trusts between Calico.cats.com and Labs.dogs.com.
- E. Set up a realm trust between Calico.cats.com and Labs.dogs.com.
- D.** The shortcut trust will greatly reduce authentication traffic seen by the root KDC. The trust path of authentication between Calico.cats.com and Labs.dogs.com will be reduced from three hops to one.
- Answer **A** is incorrect because, although parent and child trusts are bidirectional and transitive, they are automatically created as domains are joined to the forest. They cannot be created manually, and Calico.cats.com is not a child of Labs.dogs.com. Answer **B** is incorrect because although you might potentially reduce authentication traffic slightly, it would not be reduced at the root KDC where the problem lies. Answer **C** is incorrect because shortcut trusts are unidirectional. Although authentication would be partially reduced, and the problem would be helped some, a one-way trust in each direction between the two domains would be the appropriate solution, making the trust relationship bidirectional. Answer **E** is incorrect, because while the purpose of a realm trust is to channel Kerberos authentication traffic, it is designed to interoperate between a non-Windows Kerberos realm and a Windows Server 2003 domain. There is no mention of a different Kerberos realm in the question.

12. Which of the following factors should have the most influence on the number of domains that you install? (Choose all that apply.)
- A. The need for different security policies for different groups of users.
 - B. The need for different schemas for each user group.
 - C. The possibility that one company division might eventually be sold.
 - D. The separation of IT management between different major physical sites.
 - E. The structure of your organizational chart.
- A and D.** Answer **A** is correct because security policies tend to apply at the domain level. Although workarounds exist, the need for a security boundary is one of the primary design considerations for multiple domains. Answer **D** is correct because the need for geographically separated administration groups is another primary consideration for the placement of domains.
- Answer **B** is incorrect because all domains within your forest share the same schema. Changing the schema for one domain changes it for all domains. Answer **C** is incorrect because although vague future requirements should be considered, they should not be a primary factor, while current requirements for security and administrative boundaries have a clear and distinct influence on the number of domains that you should implement. Answer **E** is incorrect because organizational charts are fluid and should not be a basis for relatively permanent network infrastructure decisions.
13. Your company has just elected a new chairperson who believes in a sound information infrastructure. Currently, your network is a product of years of least-cost and least-effort development. Having the opportunity to build a new enterprise, you sit down and begin your design. What is the most critical design question to answer before beginning the installation of your new Windows Server 2003 forest?
- A. The name of the root-level domain.
 - B. The number of domains that you need to support the enterprise.
 - C. The locations of the operational master roles and GCs.
 - D. Whether to use application directory partitions for your DNS servers.
- A.** The first question you should answer is the name of your root-level domain.
- Answer **B** is incorrect because domains can be added as needed. Answer **C** is incorrect because the locations of the operational master roles and GCs can be adjusted as needed at any time. Answer **D** is incorrect because DNS can be incrementally upgraded from Standard to Active Directory-integrated, and from Active Directory-integrated to application partition enabled.

Implementing DNS in the Active Directory Network Environment

14. Your DNS expert is not an MCSE, and has mostly UNIX experience. Your enterprise runs on Windows Server 2003 Active Directory application partition enabled DNS servers. You want your DNS expert to have access to the DNS servers, but you do not want to grant the ability to configure other DCs. To which group should you add your expert?
- A. Enterprise Admins
 - B. DNSAdmins
 - C. Administrators
 - D. Domain Admins
- B.** DNSAdmins allows the proper access while adhering to the principle of least privilege.
- Answers **A**, **C**, and **D** are all incorrect for the same reason. While they will work for DNS administration, they grant more rights than are needed for that role.
15. A help desk ticket is routed to you with the diagnosis that your DNS server is failing to renew the lease on a user's IP address. After reviewing the roles that DNS plays in the Active Directory, you think you know where the problem lies. What are the roles of DNS? (Choose all that apply.)
- A. Locating DCs
 - B. Assigning IP addresses
 - C. Defining the Active Directory namespace
 - D. Name resolution
- A**, **C**, and **D**. Answer **A** is correct because DNS is required for the correct functioning of the logon process. Answer **C** is correct because an Active Directory tree name is derived from the DNS name of the tree-root domain. Answer **D** is correct because name resolution is the primary function of DNS.
- Answer **B** is incorrect because DHCP is the service that performs that function, while DNS does not.

Chapter 5: Working with Trusts and Organizational Units

Working with Active Directory Trusts

1. You are administering two domains, mycompany.com and denver.hr.mycompany.com. Users in denver.hr.mycompany.com need to access resources in mycompany.com. You want to optimize the trust relationships. What type of trust should you create to allow this?
 - A. Cross-domain trust
 - B. Shortcut trust
 - C. External trust
 - D. None

D. In this case, you do not need to create any additional trust to accomplish your purpose. The default parent-child trust will automatically connect the users in denver.hr.mycompany.com to mycompany.com.

Answer **A** is incorrect because there is no cross-domain trust. Answer **B** could work, but would duplicate the default parent-child trust. Answer **C** is incorrect because the domain exists in the same tree.
2. Your company, mycompany.com, is merging with the yourcompany.com company. The details of the merger are not yet complete. You need to gain access to the resources in the yourcompany.com company before the merger is completed. What type of trust relationship should you create?
 - A. Forest trust
 - B. Shortcut trust
 - C. External trust
 - D. Tree Root trust

C. This case calls for an external trust because you are joining two separate forests.

Eventually, you might want to create a forest trust, but that should wait until the details of the merger are solid or until the merger is complete; therefore, Answer **A** is not appropriate at this time. Answer **B** is incorrect because it would be used where an existing forest trust was in place. Answer **D** is incorrect because this is the default trust created inside a forest when a new domain tree is created.

3. Your boss just informed you that your company will be participating in a joint venture with a partner company. He is very concerned about the fact that a trust relationship needs to be established with the partner company. He fears that an administrator in the other company might be able to masquerade as one of your administrators and grant himself privileges to resources. You assure him that your network and its resources can be protected from an elevated privilege attack. Along with the other security precautions that you will take, what will you tell your boss that will help him rest easy about the upcoming scenario?
- A. The permissions set on the Security Accounts Manager (SAM) database will prevent the other administrators from being able to make changes.
 - B. The SIDHistory attribute tracks all access from other domains. Their activities can be tracked in the System Monitor.
 - C. The SIDHistory attribute from the partner's domain attaches the domain SID for identification. If an account from the other domain tries to elevate its own or another user's privilege, the SID filtering removes the SID in question.
 - D. SID filtering tracks the domain of every user who accesses resources. The SIDHistory records this information and reports the attempts to the Security log in Event Viewer.
- C. SID filtering blocks access by removing a SID. The domain's SID is attached using the SIDHistory attribute. This attribute records the domain's SID so that SID filtering can determine which domain an account is coming from. This prevents security principals from another domain from masquerading as an administrator.
- Answer **A** is incorrect because SID filtering is used to prevent this type of an attack, not permissions. Answer **B** is incorrect because the activity is not tracked in System Monitor. Answer **D** is also incorrect because SID filtering activity is not recorded in the Event Viewer.
4. You recently completed a merger with yourcompany.com. Corporate decisions have been made to keep the integrity of both of the original companies; however, management has decided to centralize the IT departments. You are now responsible for ensuring that users in both companies have access to the resources in the other company. What type of trust should you create to solve the requirements?
- A. Forest trust
 - B. Shortcut trust
 - C. External trust
 - D. Tree Root trust
- A. Because of the decision to keep both namespaces, you will need a forest trust.

- Answers **B** and **C** are incorrect because shortcut or external trust will not give you full access to the entire tree structure of the other company's namespace. Answer **D** is incorrect because a tree root trust is created when a new domain tree is created within a forest, to provide trust between the different domain trees of the forest.
5. You recently created a trust relationship with a partner company for collaboration on a joint project. This partner company has many such joint projects and has many trust relationships with other companies. You created a share containing all the files needed for the joint project. You worked with the partner company's administrator and added your project members to one of his existing universal groups that contains all of the members in his domain who need access to the project files. You added them to the permissions on the folder and the permissions on the share. You granted the universal group Read access to the share permission and Read & Execute access to the folder via NTFS permissions. The users in the universal group are now complaining that they cannot gain access to the project's files. What do you need to do to fix the problem?
- A. You need to upgrade the level of permissions on the folder to **Modify** so that the universal group can have access.
 - B. You need to upgrade the level of permissions to **Change** on the share so that the universal group can have access.
 - C. You need to break the trust relationship and recreate it; it has a corrupted file.
 - D. You need to have the domain administrator from the partner domain verify that only members from his domain are in the universal group.
- D.** The problem in this case is that the universal group contains members from another domain that is not trusted by your domain. This will prevent the members of the universal group from having access to the resource.
- The issue at hand is not the level of permission, so Answers **A** and **B** are incorrect. Answer **C** is incorrect because the trust relationship is not corrupted.

Working with Organizational Units

6. The development team of your company has started a new research project. They want to ensure that only the members of their project team are allowed to see the new directories that they create. You created a new OU that contains the user accounts of the development team, the computers they will be using, a shared folder where they are going to place their research documents, and several printers that are to be isolated from the rest of the company. They are concerned about who will have access to the new directories. How will you protect the directories from unauthorized access?

- A. Create a GPO that will limit access to the directories. Apply the GPO to the new OU.
 - B. Create a GPO that will limit access to the directories. Apply the GPO to the domain.
 - C. Create a security group that contains the members of the research group. Remove the **Everyone** group from the ACL. Add the new group to the ACL and grant it the appropriate permissions.
 - D. Do nothing. Since the directories and files are part of an OU, no one outside the OU can access them.
- C. Since this is an issue dealing with access to a resource, a security group with the appropriate permissions is the most appropriate course of action.
- Answers **A** and **B** are incorrect because GPOs are not used to control access to resources such as files and directories. Answer **D** is incorrect because the resources need to have their permissions set to restrict access to users outside the group. Doing nothing is not an appropriate course of action.
7. You created three OUs for your domain: one called *Corp*, and two child OUs called *Sales* and *Tech*. You create two GPOs, one called *Desktop* the other called *Network*. The Desktop GPO specifies the desktop settings for all users. The Network GPO specifies the network and Registry policies. The Registry policy prohibits users from being able to edit the Registry. You first apply the Desktop GPO to the Corp OU and then apply the Network GPO to the Corp OU. You want the members of the Tech OU to be able to modify Registry settings. What should you do?
- A. Nothing; because the GPOs were not applied to the Tech OU, they will not affect the users.
 - B. Nothing; because you applied the Desktop GPO first, the Desktop GPO will not take effect.
 - C. You should set **No Override** on the Tech OU so that its settings are not overridden.
 - D. You should set **Block Inheritance** on the Tech OU so that the settings from the parent OU are not applied to the child OU.
- D. Block Inheritance will prevent the settings in the parent OU from being applied to the child OU.
- Answer **A** is incorrect because the GPO was applied to the parent and the settings will be inherited by the child. Answer **B** is incorrect because the order is not relevant. Once a setting is made, it will be applied unless another setting overwrites it. Answer **C** is incorrect because if the No Override setting is used, the child cannot override the setting made in the parent OU. It will also override the Block Inheritance setting.

8. Your Active Directory domain has one site and five OUs. *Marketing* and *Technical* are child OUs to the *Corp* OU. The *Marketing* OU is a parent to the *Sales* and *PR* OUs. You are using GPOs to configure environment and security policies on the network. The following restrictions are in place:

- Corp OU Disable Registry editing tools for all users
- Marketing OU Disable modification of network connections for all users
- Technical OU Corporate logo as desktop wallpaper for all users
- Sales OU 3D Pipes screensaver for all users
- PR OU High Contrast #1 color scheme for all users

Which restriction or restrictions will be in place for users in the *Sales* OU? (Choose all that apply.)

- A. Disable Registry editing tools for all users.
 - B. Disable modification of network connections for all users.
 - C. Corporate logo as desktop wallpaper for all users.
 - D. 3D Pipes screensaver for all users.
 - E. High Contrast #1 color scheme for all users.
- A, B, and D.** Settings applied through GPOs linked to OUs affect the specified users in that OU. In addition, settings are inherited from all parent objects (unless **Block Inheritance** is specified). Because the Sales OU is a child of both the Marketing and Corp OUs, it will inherit properties from both. In this case, users in the Sales OU will not be able to edit their Registries (applied at the Corp OU level), modify network connections (applied at the Marketing OU level), and have the 3D Pipes screensaver (applied at the Sales OU level).
- Answers **C** and **E** are both incorrect because the Technical and PR OUs are not parents of the Sales OU. The settings made there will have no effect on the Sales OU.
9. You have an OU called Support. You have a GPO called RegEdit. The only setting in the RegEdit GPO is that the use of the Registry editing tools has been disabled in the User Configuration node. For performance reasons, the decision has been made to limit the numbers of GPOs that are processed at logon. The decision has been made to remove the requirement to disable the use of the Registry editing tools. What should your course of action be to implement the new decisions?
- A. Remove the RegEdit GPO from the Support OU.
 - B. Create a new GPO that enables the use of the Registry editing tools. Apply the new GPO to the Support OU.

- C. Edit the Registry on the computers used by the Support OU that will allow for use of the Registry editing tools.
 - D. Configure a local GPO to allow the use of the Registry editing tools. Set the No Override option to this policy.
 - A.** Removing the RegEdit GPO from the Support OU fulfills both requirements. It will allow the members of the Support OU to use the Registry editing tools and it will limit the number of GPOs processed at logon.
 - Answer **B** is incorrect because it will increase the number of GPOs processed at logon. Answer **C** is incorrect because it will have no effect. Although you edited the Registry, the GPO will overwrite your new settings the next time a member of the Support OU logs on. Answer **D** is incorrect because it increases the number of GPOs processed at logon.
10. You created three OUs for your domain: one called *Corp*, and two child OUs called *Sales* and *Tech*. You create two GPOs, one called *Desktop* and the other called *Network*. The Desktop GPO specifies the desktop settings for all users. The Network GPO specifies the network and Registry policies. The Desktop policy prohibits users from being able to change their wallpaper. You first apply the Desktop GPO to the Corp OU, and then apply the Network GPO to the Corp OU. You delegated control of the OU to the senior member of the Tech group. Later, the Tech OU manager modifies the Desktop GPO to allow his users to change their wallpaper. What should you do to ensure that their changes will not take effect?
- A. Nothing, since the GPOs were not applied to the Tech OU, they will not affect the users.
 - B. You should set **No Override** on the Tech OU so that its settings are not overridden.
 - C. You should set **No Override** on the Corp OU so that its settings are not overridden.
 - D. You should set **Block Inheritance** on the Tech OU so that the settings from the parent OU are not applied to the child OU.
 - C.** To ensure that a setting made in the parent OU is not changed in a child OU, set the **No Override** setting in the options for the GPO. Once No Override is set, the child cannot change the GPO setting, even if Block Inheritance is set.
 - Answer **A** is incorrect because setting made in parent OUs are applied and have effect unless the child has a setting in its GPO that changes the setting. Answer **B** is incorrect because the No Override setting here would only apply to any child domains of the Tech OU. Since the Tech OU is the child, it will have no effect. Answer **D** is incorrect because the No Override setting in the Corp OU will override the Block Inheritance setting in the Tech OU.

11. Your network consists of a single domain and five OUs. The parent OU is named *Corp*. Corp has two child OUs, *First Floor* and *Second Floor*. The First Floor OU has one child OU, *Sales*. The Second Floor OU has one child OU, *Administration*. All of the company's DCs are members of the Corp OU. The First Floor and Second Floor OUs contain the resources that belong to their respective floors. The Sales OU has nonadministrative computers, users, and groups. The Administration OU has the administration computers, users, and groups. You need to design a domainwide security policy that will accomplish the following goals:

- All users need to have the same password and lockout policy.
- Audit policies are required for only the DCs.
- The nonadministrative computers do not need the same level of security applied to them as is required for the administrative computers.
- The number of group policies to be processed at logon needs to be minimized.

You take the following actions:

- Create a single GPO.
- Import a security template for the DCs.
- Link the GPO to the domain.

Which of the desired results are achieved by your actions?

- A. All users have the same password and lockout policy.
- B. Audit policies are implemented on only the DCs.
- C. The nonadministrative computers have the same level of security applied to them as is required for the administrative computers.
- D. The number of group policies to be processed at logon is minimized.

A and **D**. Answer **A** is correct; since the GPO has been applied to the domain, all users will have the password and lockout policy. Answer **D** is also correct since there will only be one GPO processed.

Answer **B** is incorrect because one GPO applied to all computers will not allow you to create an audit policy that will only be applied to the DCs. The audit policy will apply to all computers. Answer **C** is incorrect because a single policy will not allow you to create different levels of security for the nonadministrative computers.

Planning an OU Structure and Strategy for Your Organization

12. Your Active Directory domain consists of one site. You have three OUs. The Corp OU is a parent OU to the Sales OU and Training OU. You have specified restrictions in various group policies and included them in GPOs. On the Corp OU, there is a linked GPO, which prevents users from using Registry editing tools. The Sales OU has a linked GPO that specifies a company logo as the desktop for all users. The Training OU has a linked GPO that disables users from modifying network connections. All other group policy settings are set to defaults. What restrictions (if any) will users in the Sales OU be under when they log on to the network? (Choose all that apply.)
- A. They cannot edit the Registry.
 - B. They have the company logo as their desktops.
 - C. They cannot modify network connections.
 - D. They will have no restrictions.
- A** and **B**. Settings applied through GPOs linked to OUs affect the specified users in that OU. In addition, settings are inherited from all parent objects. Therefore, users in the Sales OU cannot edit the Registry (applied at the Corp OU level), and will have the company logo as their desktops (applied at the Sales OU level).
- Answer **C** is incorrect because the GPO that sets that users cannot modify network connections has been applied to the Training OU. Since the Training OU is also a child of Corp, its settings are not applied to the Sales OU. Answer **D** is incorrect; the users in the Sales OU will have the settings from both the Corp OU and the Sales OU.
13. You have been tasked to ensure that network security policies are in place, and standards are implemented for users' configurations. The network is a single Active Directory domain network. There are five OUs: Corp, Sales, Marketing, Development, and Technical. The Corp OU is a parent OU to all other OUs. You are given the following list of objectives to meet:
- All users must be prohibited from editing their Registries.
 - All users must have a password of at least eight characters.
 - Users in the Sales and Marketing OUs must not be able to store more than 50MB of data on any server.
 - Users in the Development OU must change their passwords every 30 days.
 - All policy settings should only affect their intended targets.

Which of the following solutions will accomplish all of your objectives?

- A. Create a GPO called Policy, with settings prohibiting users from using Regedit, and requiring passwords of at least eight characters. Link Policy to the Corp OU. Create a GPO called Data, with disk quotas set at 50MB. Link Data to the Sales OU and to the Marketing OU. Create a GPO called Password, making users change their passwords every 30 days. Link Password to the Development OU.
 - B. Create a GPO called Policy, with settings prohibiting users from using Regedit, and requiring passwords of at least eight characters. Link Policy to the domain. Create a GPO called Data, with disk quotas set at 50MB. Link Data to the Corp OU. Create a GPO called Password, making users change their passwords every 30 days. Link Password to the Development OU.
 - C. Create a GPO called Policy, with settings prohibiting users from using Regedit, and requiring passwords of at least eight characters. Link Policy to the Corp OU. Create a GPO called Data, with disk quotas set at 50MB. Link Data to the Corp OU. Create a GPO called Password, making users change their passwords every 30 days. Link Password to the Corp OU.
 - D. Create a GPO called Policy. In Policy, define settings prohibiting users from using Regedit, requiring passwords of at least eight characters, setting disk quotas at 50MB, and a maximum password age of 30 days. Link Policy to the Corp OU.
- A.** The only answer that meets all requirements is Answer **A**. While this solution is a long one to implement, it is the only one that applies all desired policies to their intended targets without affecting other computers or OUs.
- Answer **B** is incorrect because the disk quota setting was applied to the Corp OU. This setting will then be applied to all users, not just the users in Sales and Marketing. Answer **C** is incorrect because it applies the disk quotas to all users, not just those in Sales and Marketing, and it applies the change password to all users as well. Answer **D** is incorrect because it makes all policies apply to all users. The need to apply policy settings to only affect their intended targets is not met.
14. Your Active Directory domain has two OUs. The Corp OU is a parent OU to the Technical OU. You have implemented a GPO linked to the Corp OU. You do not want those settings affecting the users in the Technical OU. How can you accomplish this with minimal effort?
- A. On the GPO linked to the Technical OU, select **Block Policy** inheritance.
 - B. On the GPO linked to the Corp OU, select **Block Policy** inheritance.
 - C. On the GPO linked to the Technical OU, negate any options set in the Corp OU by choosing **Disabled** for those options.
 - D. On the GPO linked to the Technical OU, select **No Override**.

- A.** By blocking policy inheritance on the Technical OU, you effectively state that all objects within that OU should start with a “clean slate,” and not be affected by any policy from a higher level. You could negate all settings in the Corp OUs GPO by selecting **Disabled** for all options, but that would be tedious at best. The **No Override** option is for administrators to prevent other admins or users from effectively using the **Block Policy** inheritance option at a lower level.
- Answer **B** is incorrect because Block Inheritance will not have the desired effect if it applied to the parent OU. It needs to be set on the child OU to block the policy settings made in the parent. Answer **C** is incorrect because this will require constant attention and does not meet the minimum effort requirement. Answer **D** is incorrect because setting the No Override on the Technical OU will ensure that its setting will not be overwritten by any of its child OU. It will not affect the settings that come from its parent OU.
15. John Smith is a junior network administrator for your company. His user account is JSmith. You want him to take charge of linking all network group policies to the appropriate OUs. Because of his experience level, you do not want him to have additional controls over the OUs. What is the easiest way to accomplish this?
- A. Use the Delegation of Control Wizard. Select JSmith, and check **Create, delete, and manage groups**.
- B. Use the Delegation of Control Wizard. Select JSmith, and check **Manage Group Policy links**.
- C. Use the Delegation of Control Wizard. Select JSmith, and check **Create and Modify Group Policy**.
- D. Use the Delegation of Control Wizard. Select JSmith, and check **Apply Group Policy**.
- B.** Using the Delegation of Control Wizard, you can allow users to manage group policy links. Note that by delegating this task, the administrator in question can manage links, but does not necessarily have permission to modify the GPO itself.
- Answers **A**, **C**, and **D** are incorrect because you only want him to have the ability to manage the Group Policy Links. The other options will give him more power than what is desired in this situation.

Chapter 6: Working with Active Directory Sites

Understanding the Role of Sites

1. An Active Directory environment has been configured with multiple sites and has appropriate resources in each site. The administrator of the Active Directory environment tries to choose a protocol for the transfer of replication information between two sites. The connection between the two sites has the following characteristics:
 - The link is unavailable during certain times of the day due to an unreliable network provider.
 - The replication transmission must be carried out whether the link is available or not.
 - Replication traffic must be able to travel over a standard Internet connection.

Which of the following protocols meets these requirements?

- A. Internet Protocol (IP)
 - B. Simple Mail Transfer Protocol (SMTP)
 - C. Remote Procedure Calls (RPCs)
 - D. Dynamic Host Configuration Protocol (DHCP)
- B.** SMTP is suitable for environments that do not have persistent connections. It uses the store-and-forward method to ensure that data is not lost if a connection cannot be established.
- Answer A** is incorrect because IP requires a persistent connection to transfer the data. **Answer C** is incorrect because RPCs are used to transfer information between DCs across remote sites that require persistent connections. **Answer D** is incorrect because DHCP is used to allocate IP addresses and distribute TCP/IP configuration information; it is not used for replication.
2. Julie installs a Windows Server 2003 server that will be used during the installation of the Active Directory structure for her organization. She installs the DNS server, creates the domain, and configures it for dynamic updates. When she attempts to install the first DC, she gets a message that the DC for the domain is not available. She decides to continue the installation and fix the problem later. What problem will she need to fix later?
 - A. The DNS server needs to be restarted.
 - B. The server she is installing needs to point to the DNS server.
 - C. The Active Directory-integrated DNS must be used while installing Active Directory.
 - D. The DNS server needs to be configured for dynamic updates and not to the zones.

- B.** In this case, the most likely cause is that the new DC is not pointing at the right DNS server.
- Answer **A** is incorrect because the switch between modes does not require restarting. Answer **C** is incorrect because the Active Directory-integrated DNS is not mandatory when installing Active Directory. Answer **D** is incorrect because the DNS service can host both dynamic and nondynamic zones. In this question, it is set on the zone level.
3. Robin is managing an Active Directory environment of a medium-sized company. He is troubleshooting a problem with the Active Directory. One of the administrators made an update to a user object and another reported that he had not seen the changes appear on another DC. It was more than a week since the change was made. Robin checks the problem by making a change to another Active Directory object. Within a few hours, the change appears on a few DCs, but not on all of them. Which of the following are possible causes for this problem?
- A. Connection objects are not properly configured.
- B. Robin has configured one of the DCs for manual updates.
- C. There might be different DCs for different domains.
- D. Creation of multiple site links between the sites.
- A.** Misconfiguring the connection objects of the Active Directory might cause a failure in updates.
- Answer **B** is incorrect because configuration of the DCs for manual updates does not cause failure in updates. Answer **C** is incorrect because the presence of different DCs for different domains does not cause failure in updates. Answer **D** is incorrect because creation of multiple site links between the sites does not cause failure in updates.

Relationship of Sites to Other Active Directory Components

4. James is a systems administrator for an Active Directory environment that consists of three sites. He wants to set up site links to be transitive. Which of the following Active Directory objects is responsible for representing a transitive relationship between sites?
- A. Additional sites
- B. Additional site links
- C. Bridgehead servers
- D. Site link bridges
- D.** Site link bridges are designed to allow site links to be transitive. They enable site links to use other site links for transporting replication information between sites.

- Answer **A** is incorrect because additional sites do not ensure that all DCs are kept up to date at a given point in time. Answer **B** is incorrect because additional site links do not allow site links to be transitive. Answer **C** is incorrect because this does not allow site links to use other site links to transfer replication information between sites.
5. Michael, a systems administrator of a medium-sized company, suspects that Active Directory replication traffic is consuming a high amount of network bandwidth. He wants to determine the amount of network traffic that is generated through replication. He plans to carry out the following procedures:
- Find out replication data transfer statistics.
 - Find out details on multiple Active Directory DCs at the same time.
 - Find out other performance statistics, such as server CPU utilization.

Which of the following administrative tools is most useful for meeting these requirements?

- A. Active Directory Users and Computers
 - B. Active Directory Domains and Trusts
 - C. Event Viewer
 - D. Performance
- D.** The Performance administrative tool enables Michael to measure and record performance values related to Active Directory replication.
- Answer **A** is incorrect because Active Directory Users and Computers cannot be used to track the replication traffic of a network. Answer **B** is incorrect because Active Directory Domains and Trusts cannot be used to monitor multiple servers at the same time and to view other performance-related statistics. Answer **C** is incorrect because Event Viewer cannot be used to track the amount of network bandwidth the replication traffic is consuming.
6. Steffi is an administrator of a medium-sized organization responsible for managing Active Directory replication traffic. She finds an error in the replication configuration. How can she look for specific error messages related to replication?
- A. Use the Active Directory Sites and Services administrative tool
 - B. Use the Computer Management tool
 - C. View the System log option in Event Viewer
 - D. View the Directory Service log option in Event Viewer
- D.** The Directory Service event log contains error messages and information related to replication.

- Answer **A** is incorrect because this tool doesn't maintain the error messages. Answer **B** is incorrect because the information related to replication is not tracked by the Computer Management tool. Answer **C** is incorrect because the System log does not contain the error messages and information related to replication.

Creating Sites and Site Links

7. George is in charge of managing Active Directory replication traffic for a medium-sized organization that has installed a single Active Directory domain. The current setup is configured with two sites and consists of default settings that are ideal for replication. Each site consists of 20 DCs. Recently, the administrators have found that the Active Directory traffic is using a large amount of available network bandwidth between the two sites. George now has the task of meeting the following requirements:
- Decrease the network traffic between DCs in the two sites.
 - Decrease the amount of change to the current site topology.
 - Make no changes to the current physical network infrastructure.

George decides that it would be highly efficient to set up specific DCs in each site that will receive the majority of replication traffic from the other site. Which of the following solutions will meet the requirements?

- A. Form additional sites that are intended only for replication traffic, and move the current DCs to these sites.
 - B. Establish multiple site links between the two sites.
 - C. Establish a site link bridge between the two sites.
 - D. Configure one server at each site to act as an ideal bridgehead server.
- D.** Bridgehead servers gather the replication information for a site and transfer this information to other DCs within the site. This plan enables George to ensure that the replication traffic between the two sites is passed through the bridgehead servers, and replication traffic will flow properly between the DCs.
- Answer **A** is incorrect because the replication traffic between the additional sites is passed through the current DCs, and replication traffic will not flow properly between the DCs due to the formation of additional sites. Answer **B** is incorrect because the establishment of multiple site links between the two sites increases the amount of change to the current site topology. Answer **C** is incorrect because it requires changes to the current physical network infrastructure.

8. James is in charge of managing the Active Directory environment for a medium-sized organization. He has to write down the procedures for creating a site for a new administrator who is starting up a new office for his organization. Which of the following is the best method for creating a site?
- A. Create the site, select the site link, add the subnets, and then move in the DCs.
 - B. Move the DCs, create the site, add the subnets, and then select the site links.
 - C. Create a temporary site link bridge, add the DCs, rename the site that's created, and then add subnets.
 - D. Create the subnets and then create a site by grouping them. Next, create the links and then move in the DCs.
- A.** You have to create the site first.
- Answers **B**, **C**, and **D** are incorrect because you are asked for the site link that the site will be part of during the creation of that site. This means that you select the site link as you create it. You can then add subnets and DCs in any order.
9. Sofia, an administrator of a medium-sized organization, has created the site links and site link bridges for the Active Directory network. The replication between the sites is working fine, and all the sites are receiving the updates to the Active Directory. She describes the network she is working on to a colleague, and he tells her that she didn't have to configure site link bridges. Why didn't Sofia have to create site link bridges?
- A. The KCC will create the site link bridges for you.
 - B. The sites will be automatically bridged.
 - C. The Domain Naming Master will handle this for you.
 - D. The GC will handle this for you.
- B.** You do not have to configure site link bridges manually, since they will be automatically bridged while creating them.
- Answer **A** is incorrect because the KCC won't actually create site link bridges. Answer **C** is incorrect because the Domain Naming Master deals with domains. Answer **D** is incorrect because the GC has nothing to do with this.

Understanding Site Replication

10. Peter, an administrator of an organization, has formed a Windows 2003 Active Directory structure. He has installed a single domains containing 700 users and computers. The organization is split into two offices with a 56 Kbps link between them. Peter creates two sites, one for each office, and a site link between them using SMTP. The replication between the sites doesn't seem to be working. What should Peter do?

- A. He has to configure an enterprise CA.
 - B. He has to configure Microsoft Exchange.
 - C. He has to configure an SMTP-based mail system.
 - D. He must have a connection faster than 56 Kbps.
 - A.** If you are using SMTP for your site links, you need to have an enterprise CA. The authority will be used to sign the SMTP packets being sent.
 - Answers **B** and **C** are incorrect because the SMTP packets are sent between servers in the sites involved in the site link and do not actually use mail servers. Answer **D** is incorrect because SMTP (e-mail) can run over a modem that is capable of 56 Kbps.
11. A company uses a single-master domain model, with resource domains for each of its divisions. It has registered two domains under the names `www.dotnetforce.com` and `www.w3force.com`. In this situation, which Active Directory information will be replicated between DCs in the `dotnetforce.com` and the `w3force.com` domains?
- A. Domain-naming context
 - B. Schema-naming context
 - C. Configuration-naming context
 - D. GC
 - E. SYSVOL
 - B, C, and D.** The schema- and configuration-naming contexts are replicated to all DCs in a forest. The GC is replicated to all GC servers in a forest.
 - Answers **A** and **E** are incorrect because both the domain-naming context and SYSVOL replication occur only between DCs in the same domain.
12. Steffie, an system administrator, has implemented two sites that are connected by a site link. The Cost property is set to 100, and the Replicate Every property is set to 50 minutes. How often will the replication occur?
- A. Every 5 minutes
 - B. Every 50 minutes
 - C. Every 180 minutes
 - D. The replication frequency cannot be determined.
 - B.** The Replicate Every property for the site link is set to 50 minutes, which determines how often replication will occur.
 - Answer **A** is incorrect because the Replicate Every property is not set to 5 minutes. Answer **C** is incorrect because the Replicate Every property is not set to 180 minutes. Answer **D** is incorrect because the Replicate Every property is used to determine the frequency.

13. A financial company with branches throughout the United States has hired a consultant to set up the Active Directory sites for their organization. Which of the following structures will he recommend?
- A. Domain structure
 - B. Political concerns
 - C. Geographic distribution
 - D. Physical network infrastructure
- D.** The primary consideration for site structure is always based on physical network infrastructure.
- Answer **A** is incorrect because the domain structure is considered secondary criteria. Answer **B** is incorrect because it is considered a main factor for the replication topologies. Answer **C** is incorrect wrong because geographical distribution is not a barrier for the latest technologies.
14. James, a network administrator, has configured Active Directory sites. He wants to implement intersite and intrasite replication. Which of the following replication protocols uses RPCs for replication?
- A. DHCP
 - B. RPCs
 - C. IP
 - D. SMTP
- C.** The IP replication protocol is used to replicate Active Directory information within sites.
- Answer **A** is incorrect because the DHCP is suitable for environments that do not have persistent connections between sites. Answer **B** is incorrect because RPCs are used to transfer information across remote sites that require persistent connections. Answer **D** is incorrect because SMTP uses the store-and-forward method to ensure replication over site links.
15. Your Active Directory structure consists of five domains running in a single forest with 40,000 users. One domain is the Sales domain. Your organization has opened a branch office with 100 employees who are members of the Sales domain. The branch office is connected to the corporate office by a high-speed WAN link. The link is reliable, and you expect the utilization rate of the link to be low. What should you do to minimize Active Directory-related authentication traffic on the WAN link? (Choose all that apply.)
- A. Add the subnet of the branch office to the corporate site.
 - B. Add a DC from the Sales domain to the branch office and configure it as a GC server.

- C. Add a DC from all five domains to the branch office and configure one DC as a GC server.
- D. Add a DC for the Sales domain at the branch office.
- E. Define the branch office as a site.
- B and E.** By defining the branch office as a site, you can control authentication traffic, because Windows 2003 will search for a DC in the site where the client is logging on. By adding a DC from the Sales domain to the branch office and configuring it as a GC, you can minimize authentication traffic in two ways. By having a GC at the branch site, no traffic will cross the WAN link to query a GC at the other end. The Sales domain's DC will authenticate the client, preventing the authentication traffic from crossing the WAN link to the corporate site because subnets are defined at the site with which they are associated.
- It is incorrect to put a DC from each domain at the branch site. Since all the 100 users are part of the Sales domain, it is only required to put a Sales DC at the branch location. While it is correct to add a Sales domain's DC to the site, this answer alone does not combine with any other answer to give a complete solution to the problem. Therefore, the rest of the choices are wrong.

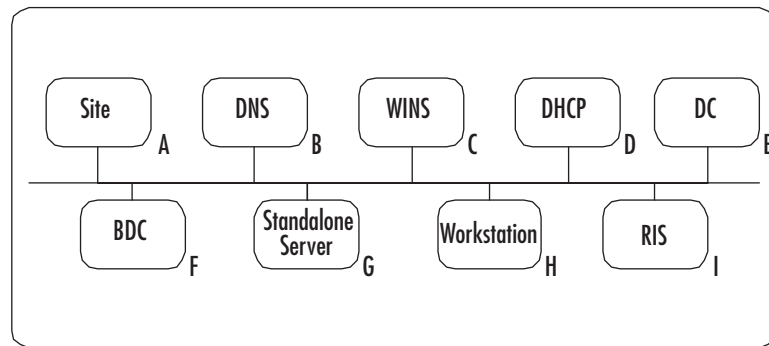
Chapter 7: Working with Domain Controllers

Planning and Deploying Domain Controllers

1. As a domain administrator you have seen the success of other departments using a RAS server to allow remote access to their domains. The other administrators use Windows NT 4 RAS and it has worked well for them. You want the same, so you install a Windows NT 4 RAS server in to your Windows Server 2003 domain. As you test this configuration, you continually get "Access Denied," no matter which user you use to dial in with. What is a likely explanation for the continual failure to allow access?
 - A. Your domain was created using **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems**.
 - B. Your domain mode is set to Windows Server 2003 domain functional level.
 - C. Your domain mode is set to Windows 2000 native domain functional level.
 - D. A Windows NT 4 RAS server cannot authenticate to a Windows Server 2003 domain.
 - A.** Some pre-Windows 2000 services required the use of the anonymous user logon to even begin a session of inquiry with a DC. RAS on Windows NT is one of them. Setting your domain to **Permissions compatible only with Windows 2000 or Windows Server 2003 operating systems** removes the anonymous user capability in lieu of higher security.

- The functional level of your domain only refers to the operating systems of your DCs and not to any of your member servers or workstations; therefore, Answers **B** and **C** are both incorrect. Windows NT 4 RAS server *can* authenticate users on a Windows Server 2003 DC *if* the anonymous account within the domain is activated, so Answer **D** is also incorrect.
2. Using the services depicted in Figure 7.43, select the components that are required to create a domain and place them in order of implementation.
- Site
 - DNS
 - WINS
 - DHCP
 - DC
 - BDC
 - Standalone server
 - Workstation
 - RIS

Figure 7.43 Network Services




- B**, **G**, and **E**. All DCs start out as a standalone Windows Server 2003 (Answer **G**), and then various services are installed to promote the server to a role in the domain. A DNS (Answer **B**) server is required for a domain. It can be created *during* the dcpromo, but even then, the DNS server is created and the appropriate records are added *before* the DC (Answer **E**) is created. It is also preferable to create the DNS first; therefore, the steps of DNS, Standalone Server, DC are the correct components and in that order.
- Although most of the components listed are often an important part of the domain services, none of them are *required* to create a domain. The DHCP (Answer **D**) server

doles out IP addresses to the computers on the network. WINS (Answer **C**) is only required if you have servers and/or workstations running Windows operating systems older than Windows 2000. A default site (Answer **A**) is created with the domain, but it is not required before installing a domain; and sites are used for WAN connections to control the bandwidth of Active Directory replication. A BDC (Answer **F**) is only used on Windows NT 4 networks and is not required for a new domain. The RIS Server (Answer **I**) is used to automate the imaging of new workstations over the network and is not required to create a domain. Finally, a workstation (Answer **H**) is used within a domain, but again, is not required to create the domain.

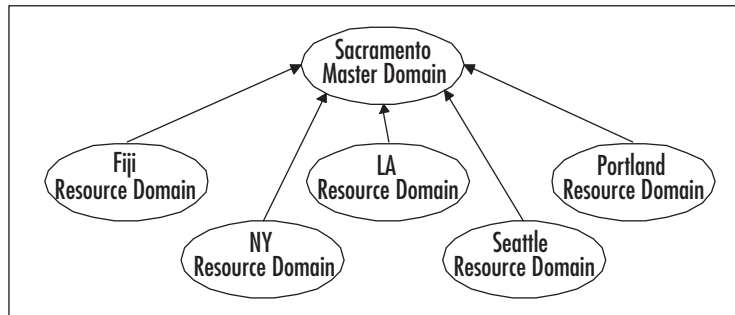
3. DasSchmeckt, the leading food services company outside the United States, has just merged with Yummy, Inc. in the United States. DasSchmeckt's headquarters are in Berlin, and Yummy, Inc. is in Atlanta. Most of the clients they serve are remote and have no need to connect directly to the company's LAN; they just use Internet mail and VPNs to access the intranet. With the merger, it has been decided that you will expand the forest by creating two domains: one in Europe and the other in the United States. To improve performance and accessibility, you will create sites at each major management location and link them all for Active Directory replication. Each management location only has 10 to 30 people, and most are connected with T1 Internet access. Use the information provided in the following table to determine the minimum number of sites and DCs you need.

Europe	WAN Speed	United States	WAN Speed
Berlin	E3	Atlanta	T3
Amsterdam	E1	New York	T3
Paris	E1	Chicago	T1
Zurich	E3	Dallas	T1
Rome	E1	Portland	T1

- A. Ten sites, with two DCs in each site
 - B. Ten sites, with one DC in each site
 - C. Eight sites, with one DC in each site
 - D. Eight sites with two DCs in each site
 - E. Eleven sites with two DCs in each site
- B.** A site should follow the WAN topology. Figure 7.44 shows this configuration. Although two DCs per site is recommended, only one per site is required. With such a small number of local users within each site, it might be difficult to justify more than one DC.

Figure 7.44 Multisite Domain

- ☒ Answer **A** gives you the correct number of sites: one at each end of a WAN link, and two DCs per site. This is not *required*, and with the low number of local users per site, more than one DC is hard to justify. (Note: Push for two anyway. In the real world, losing a DC at one of these sites could prove to be a single point of failure in your domain topology, especially since we are not told in this scenario how the replication is configured.) Answers **C** and **D** are incorrect because there are not enough sites to match the WAN topology. You might have been tempted to make the management sites with E3 or T3 connections into one site. Again, sites should follow the WAN topology. You take advantage of the higher-speed WANs by increasing the frequency of Active Directory replication, not by combining into one site. Answer **E** is incorrect because there are only 10 locations. An extra site is not required. Any replication of forestwide services such as GCs, Schema FSMO, and Domain Naming FSMO can occur by using a site link between Atlanta and Berlin, not by creating a whole site.
4. Currently, the POTC Company uses Windows NT 4. They have a single-master domain structure with five resource domains (see Figure 7.45). The IS in Oakland manages all except the offshore connection in Fiji, where most everything is in French. The POTC Company has a chance to improve on their multidomain network as they migrate to Windows Server 2003. Using the information provided, determine which domains can become OUs and which must remain a domain.

Figure 7.45 Single-Master NT Domain

- A. Create one domain incorporating all five resource domains into OUs.
- B. Create two domains: one for the root, Sacramento, and one subdomain for the five resource domains, creating OUs for each location.
- C. Create three domains: Sacramento with LA, Portland, and Seattle as OUs; Fiji as a subdomain, and New York as another subdomain.
- D. Create two domains: Sacramento with LA, Portland, Seattle, and New York as OUs; Fiji as a subdomain.

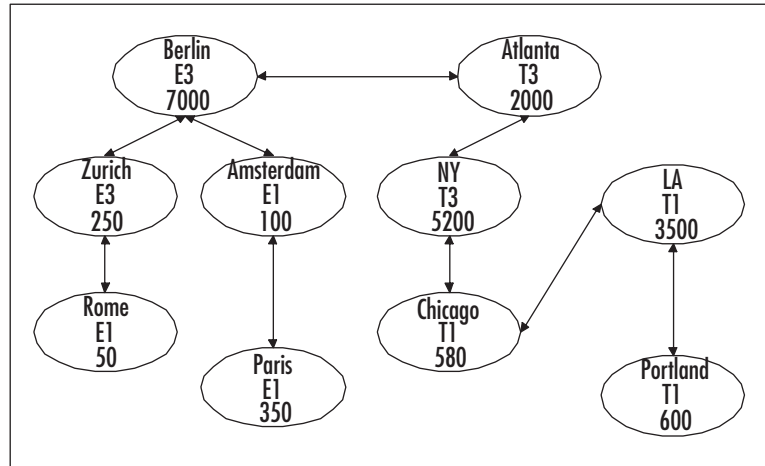
D. There are very few technical reasons to have more than one domain.

Administratively, the Fiji location manages its own users and resources and probably has a lot of French words and foreign spellings that just make it too difficult for the IS in Sacramento to manage, so it should have its own domain. All of the other locations can be incorporated into one domain with five OUs. The WAN connection between New York and Sacramento is not very critical because Windows Server 2003 domains manage Active Directory replication across WAN connections very efficiently. In addition, the IS is centralized in Sacramento, so it is better to keep New York within the one domain.

Making this single-master domain into just one domain is possible, but the deciding factors are twofold: First, the WAN connection between the United States and Fiji is likely inconsistent and slower than a T1. Second, the resources in Fiji are managed by the Fiji administrators, so they should have their own domain. Answer **A** is incorrect for this reason. Answer **B** is incorrect because there is an administrative reason for keeping Fiji separate from the rest of the domains. Fiji is administered by its local staff, and the language used is foreign to the main IS in Sacramento, leading to frustrating naming errors and other miscommunication with Fiji's resources if it were to be managed by the IS in Sacramento. Answer **C** is also incorrect because a third domain is not needed for New York. Windows Server 2003 domains manage Active Directory replication very efficiently across WANs, and the IS manages the resources from Sacramento, so there are no political or administrative reasons to keep New York as a separate domain.

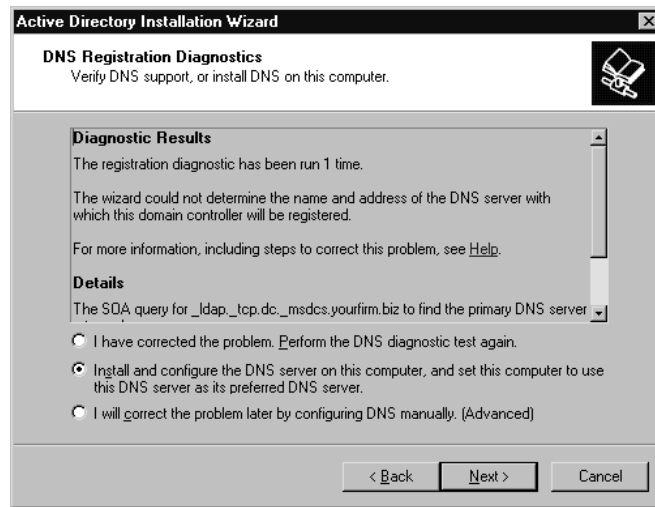
5. Referring to Figure 7.46, determine the minimum number of DCs required. Each oval represents a physical location and lists the WAN connection speed available at that site. The arrows indicate the proposed replication strategy.

Figure 7.46 Site Topology



- A. Ten, one at each site
- B. Twelve, one at each site plus one more in Berlin and Atlanta
- C. Twelve, two at each site plus one more in Berlin and New York
- D. Twenty, two at each site
- D.** Following Microsoft's guidelines and planning for redundancy and fault tolerance, place two DCs at each site and one DC for every 5000 users. Since S5 has 7500 users, you might consider a third DC there, but the minimum requirements are still met with only two at that site as well. Two DCs per site times six sites equals twelve, so **D** is correct.
- Anything less than 12 disregards redundancy and fault tolerance. You should place at least two DCs per site; therefore, Answers **A**, **B**, and **C** are incorrect because there are not enough DCs to place two at each site.
6. You are installing Windows Server 2003 and promoting it to the first DC of your new domain, BusyBees.biz. During the Active Directory Installation Wizard process you get the dialog box in Figure 7.47. What is the solution to the problem? (Choose all that apply.)

Figure 7.47 DNS Diagnostics



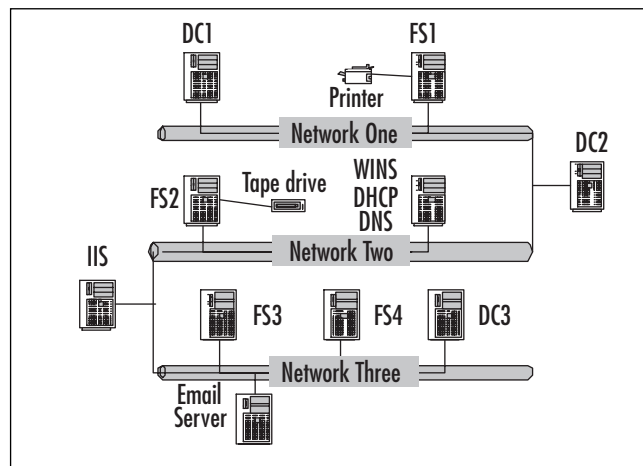
- A. Do nothing. The ADIW will create a DNS server for you.
 - B. Cancel ADIW. Install a DNS server that supports RFC 2136 (dynamic updates).
 - C. Cancel ADIW. Install a Windows 2000 DNS server using the defaults.
 - D. Cancel ADIW. Install a Windows 2000 DNS server. Create a primary zone called BusyBee.biz and enable dynamic updates.
 - E. Cancel the ADIW. Install a Windows 2000 DNS server. Create a primary zone called BusyBee.biz and don't enable dynamic updates.
- A, B, and D.** A DC requires a DNS server that supports both SRV records and dynamic updates. Answer **A** is correct because the ADIW will indeed create a DNS server automatically with all the appropriate records and dynamic updates required. If you know you have a DNS server that meets the RFC requirements, then you are looking at a connectivity issue. In this case, you should cancel ADIW, fix the connectivity issue, and then restart the installation process. Answer **B** is correct because it is possible that the required DNS server is just not there yet and you want to create it yourself, first. Answer **D** is correct because a Windows 2000 DNS server meets the requirements necessary to host the Windows domain services and records, and dynamic updates is enabled, which is also required.
- Answer **C** is incorrect because just installing Windows 2000 DNS server using the defaults leaves the dynamic updates disabled and you need them enabled. Answer **E** is incorrect for the same reason: dynamic updates *must* be enabled or the DC cannot be created.

Backing Up Domain Controllers

7. Mark is the local administrator for the site in Portland. His duties include the backups for the servers in his site, using Windows Backup. His site includes a DC that he does not back up because the DC in San Francisco is backed up and all Active Directory replications come to Portland once a night. When Mark loses his DC in Portland to a lightning strike, he replaces the server and now wants to restore the computer to a DC. What is the simplest and fastest way to do this?
- A. Restore the system state from a DC in San Francisco.
 - B. Promote the server to a DC using ADI/W.
 - C. Ship the server to San Francisco and have the dcpromo run there and then ship it back.
 - D. Set up a VPN and then run dcpromo from San Francisco.
- A.** A backup of a Windows Server 2003 DC can be restored onto any other server, making it a DC. In this case, restoring a backup from a healthy DC in San Francisco is the quickest and easiest method.
- Although promoting the restored server to a DC using ADI/W works, it requires a WAN to connect to a DC in San Francisco and will be quite slow; therefore, Answer **A** is a *better* choice and Answer **B** is incorrect. Shipping the server to San Francisco and promoting it to a DC there eliminates the WAN, but requires time and money to ship the server back and forth, which takes longer than choice **A**; therefore, Answer **C** is also incorrect. Setting up a VPN and running dcpromo from San Francisco is just like option **B**. The WAN is involved and will be slow; therefore, **D** is incorrect.
8. Stephanie is the administrator for the scrapbook company, Book On Over Co. (BOOC), which was recently bought by their competitor, Buecher Sind Toll GMBH (BST, a German company). Consequently, the two Windows Server 2003 domains were brought into one tree with two domains. Manfred, the systems engineer for BST, recently performed an authoritative restore of the Active Directory in his domain successfully and informed Stephanie of it. Now the Managers group in the BOOC domain can no longer access data on the Forms server in the BST domain. Based on the information given, the authoritative restore seems to have caused the problem. What is the likely cause of this problem?
- A. The Managers group was deleted by the authoritative restore.
 - B. The authoritative restore removed the Trust between the domains.
 - C. The authoritative restore replaced the Security ID (SID) of the Managers group to an old SID that makes it no longer valid in either domain.
 - D. The password used by the Trust between the two domains was changed to an old password by the authoritative restore.

- D.** Trusts periodically and automatically change their password. An authoritative restore of the Directory, and not just certain pieces, can restore an old password used by the Trust. This is then out of sync and causes the Trust to fail. Without a valid Trust to allow cross-domain authentication, the users are confined to resources in their local domain only.
- Answer **A** is incorrect because the Managers group exists in the BOOC domain, and the restore was done in the BST domain, so the group still exists. If a password for a Trust is removed or replaced, the Trust itself still exists; therefore, Answer **B** is incorrect. The Managers group is not in the BST domain where the restore took place. In addition, once an object is assigned its SID, there is no such thing as an old SID. The object can be renamed, but the SID remains the same; therefore, Answer **C** is incorrect.
9. Using the diagram in Figure 7.48, determine which data can be included in the daily backup routine to the tape device connected to FS2.

Figure 7.48 Backing Up a Domain



- A. Net 1: AD on DC1; FS1
 6. Net 2: FS2; IIS
 7. Net 3: Email; FS3; FS4
- B. Net 1: FS1
 8. Net 2: AD on DC2; FS2; IIS
 9. Net 3: Email; FS3; FS4

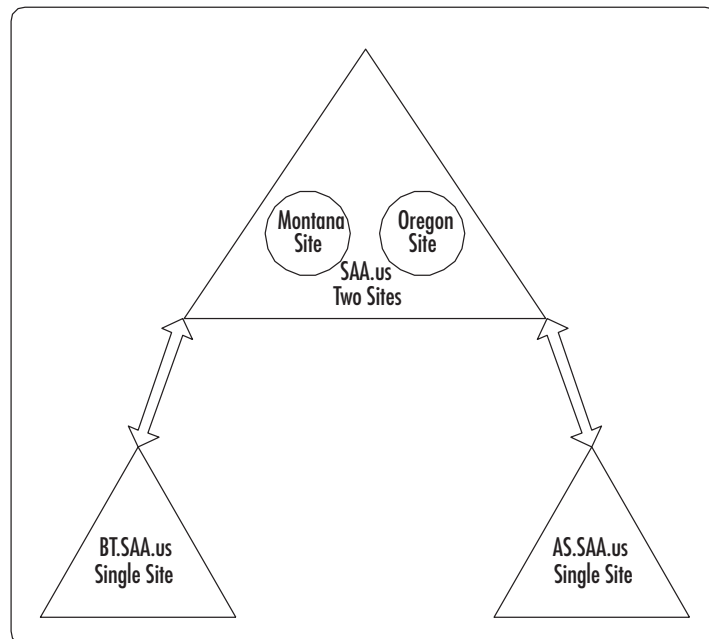
- C. Net 1: FS1
- 10. Net 2: FS2; IIS
 - 11. Net 3: AD on DC3; Email; FS3; FS4
- D. Net 1: FS1
- 12. Net 2: FS2; IIS
 - 13. Net 3: Email; FS3; FS4
- D.** The only answer possible is **D** because the other three are not possible. Why? They each list the Active Directory as part of the backup to tape. The system state, of which Active Directory is a part, requires a direct local connection to the device you are backing up. The tape device in Figure 7.48 is locally connected to FS2, which is *not* a DC; therefore, Active Directory cannot be backed up.
- Answers **A**, **B**, and **C** are not possible for the same reason stated in the correct answer. The Active Directory cannot be backed up in Figure 7.48, because the tape device is not locally connected to any DC.
10. Brayden is the domain administrator for a multisite Windows Server 2003 domain. The headquarters is located in South Bend, Indiana. A new branch is being opened remotely in San Jose, California. Brayden needs two DCs to place at the new San Jose site. The WAN link won't be up for two more weeks, but he wants to get the DCs online and in place this week so his San Jose technicians can begin setting up the workstations in San Jose right away. What can Brayden do to create those DCs before the WAN is installed?
- A. Create the two DCs in South Bend, and then ship the servers to San Jose.
 - B. Create a backup of a DC in South Bend to a CD or DVD and ship it to San Jose.
 - C. Create the DCs in San Jose, and then when the WAN link is installed, synchronize them with the DCs in South Bend.
 - D. Nothing. Brayden must wait for the WAN link before creating the remote site's DCs.
- A.** To create a DC within the same domain, you must have an available existing DC to authenticate to and then synchronize with.
- Answer **B** is a nice option except that once the DC is restored, it requires a connection to an existing DC, which is unavailable without a WAN link. Answer **C** is incorrect because you cannot create an additional DC within the same domain without an existing DC available. Answer **D** is incorrect because there *is* something Brayden can do: install the two DCs in South Bend and express ship them to San Jose.

Managing Operations Masters

11. James comes to work on Monday and opens the Active Directory for Users and Computers. His task today is to create three new users and create a new group. James attempts this and it fails repeatedly. He knows that one DC went down over the weekend, but he is not connected to that DC and can see all the objects in Active Directory. Users are logging on just fine as well. What is a possible explanation for not being able to create new objects in Active Directory ?
- A. James is not logged on as a Schema Admin.
 - B. The DC that went down had the Domain Naming FSMO on it.
 - C. The DC that went down had the RID FSMO on it.
 - D. The DC that went down had the Schema FSMO on it.
 - E. The DC that went down had the PDC Emulator on it.
- C.** All new objects in a domain require a unique SID. The RID keeps a pool of unique SIDs to give out when an object is created.
- Answer A** is incorrect because a member of the Schema Admins grants rights to modify the schema and has nothing to do with creating new objects in Active Directory. James only needs to be a Domain Admin to create new objects. The Domain Naming FSMO manages the names of domains within a forest, not the names of objects in a domain, so **Answer B** is incorrect. **Answer D** is incorrect because the Schema FSMO controls the extension of the schema only and has nothing to do with creating objects in Active Directory. **Answer E** is incorrect because the PDC Emulator controls Active Directory replication to Windows NT 4 BDCs, is the master copy of all passwords for each domain, and has nothing to do with creating objects in Active Directory.
12. Ryan is a domain administrator for Astronauts Ltd. It is a multidomain tree with five sites. Today, he must add some users to the Marketing group. He uses ADUC to open the group and adds the users, Brayden and Hannah, from the SD.CA.COM domain. The users Rebecca and McKay are already members of this group from the LA.CS.COM domain. In testing the access of these users to the Contact database used by the Marketing department, Ryan finds that the users Brayden and Hannah are still unable to access the database, while Rebecca and McKay can. Which of the following is an applicable troubleshooting step in diagnosing this problem?
- A. Verify that the group is a distribution list.
 - B. Verify that the group is a local group.
 - C. Verify that the RID FSMO is online and available.
 - D. Verify that the Infrastructure FSMO is online and available.

- D.** A symptom of the Infrastructure FSMO missing is group creation failure or the inability to add users to a group. The Infrastructure FSMO is responsible for Active Directory objects being updated between domains and can be a cause of inconsistent group membership problems. All FSMOs are important to the management of your domain, and making sure they are online and available is a good troubleshooting step. (Note: The Infrastructure FSMO is a helper in this role and, given enough time, the group would probably catch up and work fine.)
- Answer **A** is incorrect because a distribution list cannot have permissions assigned to it, and this group needs permissions. Answer **B** is incorrect because the group needs to be a global group and local group in a domain are used when multiple domains exist and you wish to assign groups from the other domains permissions in your domain. Answer **C** is incorrect because the RID FSMO manages the unique SID required for new object creation, and in this scenario, all the objects already exist; therefore, the RID FSMO has nothing to do with this problem.
13. As an enterprise administrator for the Sports Agents of America (SAA), you must migrate the newly acquired agency's domain into your existing forest as a child domain to SAA.us. The new agency is called Alternative Sports, Inc. The new Windows Server 2003 domain is called AS. Figure 7.49 shows the current domain and site topology of SAA.us. To set up the migration, your first step is to create the child domain, AS.SAA.us. This fails repeatedly. What is a possible reason for this?

Figure 7.49 Sports Agency of America Domain Tree



- A. The Domain Naming FSMO located in the Montana site is offline.
 - B. The Schema FSMO in the Montana site is offline.
 - C. The FSMOs for AS.SAA.us need to be created before you can create a child domain.
 - D. The Infrastructure FSMO is unavailable.
- A.** A possible reason for the child domain to fail creation is if the Domain Naming FSMO is unavailable. This FSMO's role is to ensure unique domain names within a forest and must be available when creating a new domain. This FSMO should be on a DC that is highly accessible and well connected.
- The Schema FSMO must be available when the schema needs to be extended. Creating a new domain does not modify the schema, so Answer **B** is incorrect. Answer **C** is incorrect because all FSMOs are created on the first DC and then exist automatically in your domains. The FSMOs are already there, and since the only one that must exist *before* you create a child domain is the Domain Naming FSMO, you do not need to *create* any FSMOs, but you must ensure that the Domain Naming FSMO is available. Answer **D** is incorrect because the Infrastructure FSMO deals with updating changes made to the Active Directory user and group objects, and not naming a domain.
14. Michael is an enterprise administrator for NuttyNuts, Inc. He is installing Microsoft Exchange 2000 into his domain. His domain, nuttynuts.biz, has two sites and one child domain: CA.nuttynuts.biz, a subsidiary in Sacramento, California. Michael logs on to the domain with his focus on a local DC and as a member of the Enterprise Admins group. During the Exchange installation, he runs across errors that restrict him from completing the installation. Which is a possible reason for this problem?
- A. Exchange 2000 cannot run on Windows Server 2003 domains because the schemas are incompatible.
 - B. The RID FSMO is unavailable.
 - C. The Domain Naming FSMO is unavailable.
 - D. Michael must log on as a member of the Schema Admins group.
- D.** A user in the Schema Admins group can only modify the schema, and MS Exchange 2000 requires access to the schema in order to modify it. Therefore, even an Enterprise Admin cannot install Exchange or any other Directory-enabled application.
- Answer **A** is incorrect because Exchange does not have its own schema; rather, it is Directory enabled. This means that it modifies the schema of Active Directory, so the schemas cannot be incompatible. Answers **B** and **C** are incorrect because neither the RID nor the Domain Naming FSMO affect the installation of an Active Directory-enabled application installation. The RID FSMO keeps a pool of SIDs to issue to DCs as needed. If the installation of an application tried to create more ADS objects than its current pool, the DC would attempt to get more SIDs from the RID FSMO. If the RID FSMO were unavailable, the object creations would fail. The Domain Naming FSMO comes into play only when creating a new domain, which in this case we are not.

15. Heather has been hired to come into your company and install a customized Directory-enabled application. Only the users in your branch office located in Fresno, California use this application. Your headquarters is in Santa Rosa, California, and you created a site for each location and set up directory replication over the slow WAN link to occur only at night. Access between the sites occurs at that time, but occasionally you allow the sites to connect during the day when a certain threshold of requests is reached. You create a temporary account for Heather and place the new account in the Schema Admins group. Heather begins to install the application but soon realizes that the schema will not let her extend it, as the application requires? Which is a possible reason for this?
- A. She must install the application in Santa Rosa and then set up Terminal Services for the users in Fresno to access the application remotely.
 - B. She needs to wait for the schema extension requests to be processed between the two sites.
 - C. The Schema FSMO is unavailable.
 - D. The schema can only be extended on the DC that holds the Schema FSMO.
- C.** A Directory-enabled application implies that it will interact with your Directory. This means it will create a new object in your Active Directory that has properties unique to its function. To create a new object, the schema is extended. To extend the schema requires the knowledge and approval of the Schema FSMO. If the Schema FSMO is unavailable during this installation, it cannot get approval to extend the schema.
- Remote access to the application does not change the fact that the installation fails because the schema cannot be extended. Setting up terminal services, Answer **A**, then is incorrect. Answer **B** is incorrect because there is no such thing as a schema extension requests. To extend the schema, the Schema FSMO *must* be present and accounted for at the time of extension. The Schema FSMO can reside anywhere in a forest. It only approves and manages the schema extension, but the schema is part of every piece of the Directory and can be extended from any DC, so Answer **D** is also incorrect. It is a good idea to be connected via a fast link to the Schema FSMO to improve the efficiency of the application installation, but only its immediate presence is required regardless of how long it takes to access it.

Chapter 8: Working with Global Catalog Servers and Schema

Working with the Global Catalog and GC Servers

1. You are working on your DC and want to be able to run the Schema snap-in. You click on Start and select Run. You type MMC and press Enter. When you go to add the snap-in, you don't see it listed as one you can add. Why?
 - A. The DC you are on is not the GC server, so the Schema Admin snap-in would not be available on that DC.
 - B. You are not a member of the Schema Admins group, so you cannot install the snap-in.
 - C. The DC you are logged on to doesn't serve the role of Schema Master, so the snap-in will not run.
 - D. The schmmgmt.dll file has not been registered.

D. The schmmgmt.dll has to be registered with the following command: **regsvr32 schmmgmt.dll**, before you can add the Schema management snap-in to the MMC.

Answer **A** is incorrect because the DC does not have to be the GC server to run the snap-in. Answer **B** is incorrect because being a member of the Schema Admins group determines whether you have write ability to modify or extend the schema, but has nothing to do with running the Schema Admin tool or installing the snap-in. Answer **C** is incorrect because the DC does not have to serve the Schema Master Operations role to run the snap-in.
2. You just finished setting up a forest containing three DCs. Server DC1 is the forest root DC. Servers DC2 and DC3 will serve as DCs also. You want to assign the GC responsibility to DC2. How do you determine which DC is serving as the GC server now? (Choose all that apply.)
 - A. You can look in the Properties of each *Server* object within the Active Directory Sites and Services administrative tool to determine if the server is the GC server.
 - B. You know that DC1 is the GC because the first DC set up in the forest automatically takes the role of GC.
 - C. You can look at the Properties of NTDS Settings under each *Server* object within Active Directory Sites and Services.
 - D. You know that DC3 is the GC server because the third DC takes role of GC away from the forest root server upon being added to the domain.

- B** and **C**. The forest root DC is assigned all the roles within Active Directory because it is the only DC that exists during initial forest setup. You can also look at the Properties of NTDS Settings in the *Computer* object if you run Active Directory Sites and Services. The GC setting will be checked if the server is assuming that role.
- Answer **A** is incorrect because the Properties of the server will not show any GC settings. Answer **D** is incorrect because the GC role is manually enabled or disabled, and no automatic reassignments happen during addition of DCs.
3. You have a new attribute that needs to be added to the GC. You have the Schema Admin snap-in open. How do you make sure an attribute is included in the GC?
- A. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Replicate this attribute to the GC** is selected.
- B. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Allow this attribute to be shown in advanced view** is selected.
- C. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Index this attribute for containerized searches in the Active Directory** is selected.
- D. Expand the **Attributes** section, right-click the attribute you want to include, and select **Properties**. On the **General** tab, make sure that **Index this attribute in the Active Directory** is selected.
- A**. Answer A is correct and will include the attribute in the GC as long as the check is in the box that says **Replicate this attribute to the GC**.
- Answer **B** is incorrect because allowing the attribute to be shown in advanced view doesn't add the attribute to GC. Answers **C** and **D** are incorrect because those options deal with indexing the object for faster queries.
4. You recently made your new staff member a member of the Universal Group named Enterprise Admins. The new staff member is located at a branch office. When the user logs off and then back on, he notices that he cannot get to some of the Administrative tools. You recently added the user to the Universal Group and you have a 56K link between your branch office and your main office. Your GC server is at the main office. What could be the problem? (Choose all that apply.)
- A. You cannot add users to Universal Groups, only to Global and Domain Local groups.
- B. You have Universal Group caching turned on and the cache information hasn't refreshed since this morning.
- C. Transmission of GC data is failing across the WAN link.
- D. GC replication doesn't support 56K links.

- B** and **C**. You could have Universal Group caching enabled, which could cause old information to be retrieved from the cache on the DC at the branch office. The GC has to be able to transmit data on port 3268, and this might be failing if the workstations or DCs have no way to transmit information between the corporate office and branch office.
- Answer **A** is incorrect because users can be members of Universal Groups. It is not a recommended method to keep users organized, however. Generally, you would make Domain Local or Global Groups members of the Universal Group. Answer **D** is incorrect because the speed of the link doesn't prevent replication or sending of data. However, the slower the link, the slower the query process will be.
5. You have a network with a main office and a satellite office. The functional level of your network is Windows 2000 Native. The satellite office has a DC. The main office has a DC and a GC server. You encounter a problem with the link between the main office and the satellite office. You are concerned that users will not be able to log on at the satellite office because they cannot access the GC. To your surprise, they are still able to log on to the domain. How is this possible?
- A. The DC at the branch office could be set to cache Universal Group information, allowing clients to still log on.
- B. The GC isn't required for logon, simply for searching the directory after you are logged on.
- C. The DC at the satellite office is operating in the role of Schema Master and can authenticate without a GC server.
- D. The users are logging on locally and not authenticating to the domain.
- A**. The caching of Universal Group information can be cached on the DC at the satellite office, allowing authentication to still function.
- Answer **B** is incorrect because the GC has a lot to do with logging on. Answer **C** is incorrect because the role of Schema Master has nothing to do with accessing the GC. Answer **D** is incorrect because as the questions states, the users are logging onto the domain.
6. You have multiple locations that are part of the Default-First-Site-Name site. These locations are in Florida, Oregon, and Iowa. You have instituted GC servers at each location. While monitoring your network, you are noticing a lot of replication traffic between the locations. How can you remedy the amount of replication traffic and how that traffic is handled?
- A. Implement the use of *Subnet* objects
- B. Implement the use of *Object* classes
- C. Implement the use of sites

- D. Implement the use of site connectors
- C. The use of sites will help optimize the replication of traffic by using compression for intersite replication.
- Answers **A**, **B**, and **D** are incorrect. These are other objects in the Active Directory, but not the type of objects we were looking for. The *Subnet* object defines a subnet on your network. The *Object* class is a component of the schema that defines the object type. Site connectors help in connecting two different site objects.

Working with the Active Directory Schema

- 7. You are working with the Schema Admin snap-in and cannot make any changes. You created a network administrator equivalent account in the forest root domain but cannot modify the schema. Why?
 - A. You must be a member of the Enterprise Admin group to modify the schema.
 - B. You must be a member of the Schema Admin group to modify the schema.
 - C. You must be a Domain Admins member in each domain in the forest to modify the schema.
 - D. Only the initial Administrator account during forest creation can modify the schema.

B. You must be a member of the Schema Admins groups to make changes to the schema.

Answer **A** is wrong because the Enterprise Admins group cannot modify the schema. Answer **C** is incorrect because even if you are a Domain Admin in every domain, you still cannot modify the schema without being a member of Schema Admins. Answer **D** is incorrect because, although the initial Administrator account is part of the Schema Admins by default, it isn't the only account allowed to make schema changes; other accounts can as long as they are members of the Schema Admins group.
- 8. You are a network administrator and you want to modify an attribute that is associated with one of your user accounts. How do you do this?
 - A. Open **Active Directory Users and Computers** and change to **advanced** view. This will allow you to modify the properties of the attributes in the user account for which you need to make the change.
 - B. Open **Active Directory Sites and Services**. Open the **Properties** for the site containing the attribute and make the modifications.
 - C. Open the **Schema** Snap-in, expand **Objects**, and select the *User* object to modify the associated attributes.
 - D. Open the **Schema** Snap-in, expand **Attributes**, and find the attribute you want to modify.

- D.** The **Schema** snap-in is used to make changes to attributes, by selecting the attribute you want to modify and selecting **Properties**.
- Answers **A** and **B** are incorrect because ADUC and ADSS do not allow for schema modifications. Answer **C** is incorrect because it references the Objects section instead of the Attributes section.
9. You are explaining the various attributes to a fellow network administrator. You are showing her the properties of a User account, and your new network administrator asks what the Other button means with regard to various attributes. What do you tell her?
- A. Those attributes are multivalued attributes.
- B. Those attributes are single-value attributes.
- C. Those attributes are actually *Object* classes.
- D. Those attributes are Index attributes.
- A.** Attributes of an object with the Other button allow you to input more than one value, making the attribute a multivalue attribute.
- Answer **B** is incorrect because this type of attribute does not have the Other button next to it. Answer **C** is incorrect because attributes and *Object* classes are two different components of the schema. Answer **D** is incorrect because attributes that are indexed do not necessarily have an Other button by them unless they are multivalued attributes.
10. As a network administrator, you are responsible for making sure that various attributes are indexed for optimal performances for queries. What steps do you take to make an attribute indexed?
- A. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Index this attribute in the Active Directory**.
- B. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Replicate this attribute to the GC**.
- C. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Allow this attribute to be shown in advanced view**.
- D. Using the **Schema** snap-in, right-click the attribute you want to index and select **Properties**. Select **Attribute is Active**.
- A.** These steps will add the attribute to the index.
- Answer **B** is incorrect because that setting deals with an attribute being added to the GC and not the index. Answer **C** is incorrect because allowing an attribute to be shown in advanced view doesn't add the attribute to the index. Answer **D** is incorrect because this setting either activates or deactivates the attribute in the schema.

11. You are working with *Schema* objects and you need one component that has to be supplied by a third-party. Which component is supplied by a third party so standards can be followed?
- A. LDAP name
 - B. Common name
 - C. OID
 - D. Object GUID
- C.** The OID is supplied by either the ISO or ANSI based on their standards.
- Answer **A** is incorrect because the LDAP name, although based on the X.500 standards, can be customized to your environment. Answer **B** is incorrect because, as with Answer A, the Common Name is customized to your organization; it is just a simplified way to identify objects. Answer **D** is incorrect because the Object GUID isn't something you would work with in schema management.
12. You make a mistake while setting up new classes in your schema. You want to correct the mistake so you can have the appropriate name and configuration for the class. How do you do this?
- A. You must deactivate the class that was added with the mistake and then rename it. You then can create a new class with the appropriate name and configuration.
 - B. You must delete the class that has the mistake and simply create the appropriate *Class* object.
 - C. You must wait 24 hours before you can delete any new classes in the schema. You can then delete the class and create the corrected *Class* object.
 - D. You can go in and fix the existing *Class* object without having to recreate the object.
- A.** You can deactivate the object and then rename it, allowing you to create a new class with the appropriate name and configuration.
- Answers **B** and **C** are incorrect because classes and attributes in schema cannot be deleted. Answer **D** is incorrect because you cannot modify the class after it is created; you must deactivate it and recreate a new one.
13. You have an office with three locations separated by 56K WAN links. You are experiencing slow queries when looking for objects in the Active Directory. You have one GC server at your main office. What can you do to improve the query performance?
- A. Add GC servers to your other two locations.
 - B. Add DCs that are not GC servers to your other two locations.
 - C. Add a DNS server for faster resolution at your other two locations.
 - D. Add another OU to the directory to separate the locations by OU.

- A.** Having a GC at each location will help with query response time. Sites might be another consideration to ensure that you don't max out your WAN line with GC replication traffic.
- Answer **B** is incorrect because a DC that is not a GC server will not help with query response. Answer **C** is incorrect because a DNS server will have nothing to do with GC response. Answer **D** is incorrect because another OU will not help query response.
14. You have been experiencing a large amount of processor utilization on your GC server. Your network consists of one location with 2500 users. You currently have three DCs for fault tolerance and load balancing. What can you do to help with your GC server processor utilization?
- A. Add a fourth DC to the network.
- B. Add another GC server to the network to offload some of the traffic.
- C. Remove one DC from the network.
- D. Split your network into three OUs with less than 1000 users each.
- B.** The GC traffic will be more balanced having two GC servers and should cut down on the processor utilization because the original server is overloaded with GC traffic.
- Answer **A** is incorrect because another DC would not help offload some of the GC traffic on the overloaded GC server. Answer **C** is incorrect because removing a DC would cause more traffic to the other two DCs. Answer **D** is incorrect because splitting the networks into OUs will not help offload the traffic from the overloaded server.
15. You are working on updating the schema and cannot associate an attribute with a class. What can you do to resolve this?
- A. Add yourself to the schema Admins group.
- B. Makes sure the Schema Operations Master is online and reachable.
- C. Reload the schema in the Schema admin tool.
- D. Move the role of Schema Operations Master.
- C.** You need to reload the schema so the schema cache is updated.
- Answer **A** is incorrect because if you are not part of the schema Admins group, you will not be able to do anything with the schema other than view it. Answer **B** is incorrect because the Schema Operations Master role would prevent any changes to the schema. Answer **D** is incorrect because moving the role will not help with outdated schema cache.

Chapter 9: Working with Group Policy in an Active Directory Environment

Understanding Group Policy

1. You have just set up a Windows Server 2003 Active Directory network, and you want to use group policies to control user configuration. You have configured local policies on some of the machines in your domain, and you also want to configure some site and OU policies for more granular control, but you are concerned about policies at different levels “canceling each other out.” Which of the following types of GPOs will override settings applied at the domain level? (Choose all that apply.)
 - A. Local
 - B. OU
 - C. Site
 - D. Domain

B. Only GPOs assigned at the OU level can override settings applied at the domain level.

Answers **A** and **C** are incorrect because local and site policies are processed before domain policies. Answer **D** is incorrect because multiple GPOs processed at the domain level are processed concurrently.
2. You have been asked to set up a group policy environment in a new Windows Server 2003 Active Directory network. Your supervisor has asked if local computer settings will override settings applied in a domain GPO. You explain to him that policies applied later in the processing order generally take precedence over policies set earlier. In what order are group policies applied?
 - A. OU policies, domain policies, site policies, local policies
 - B. Site policies, domain policies, OU policies, local policies
 - C. Local policies, site policies, domain policies, OU policies
 - D. Local policies, OU policies, domain policies, site policies

C. Policies are always processed starting with the local computer policy, then following all directory policies from the farthest GPO from the object to the closest GPO to the object.

Answer **A** is incorrect because local policies are always processed first, and site policies are processed before OU policies. Answer **B** is incorrect because local policies are always processed first. Answer **D** is incorrect because site policies are processed before OU policies.

3. Your department has just hired a new junior system administrator and has asked you to train him. The trainee has worked some with Active Directory, but has never used Group Policy before. He has been running RSoP in planning mode to get an understanding of where different group policy settings are stored, but he keeps getting confused because he is not seeing the same groupings between the computer settings and user settings in the report. What are the main types of policies for user and computer configurations he should see in the report, as represented by nodes in the console tree?
- A. Assign scripts, Manage applications, Redirect folders, and Change Registry settings
 - B. Software settings, Windows settings, and Administrative templates
 - C. Security settings, Account settings, and Software settings
 - D. Local settings, Site settings, Domain settings, and OU settings
- B.** User configuration and computer configuration settings are collected into these three nodes.
- Answer **A** is incorrect because these items describe some of the specific tasks you can achieve with group policy, but none are specifically listed in the three groups under user and computer configuration. Answer **C** is incorrect because there are security settings under Windows settings, but there are no groups labeled Account settings and Software settings. Answer **D** is incorrect because these terms describe the locations where you can apply GPOs.
4. You work for a large company that has just acquired another company in a merger. The acquired company has merged its Active Directory structure into your forest. The new group wants to maintain control over their portion of the directory, but you want to make sure certain that domain policy settings are not changed by GPOs applied at the OU level. How will you achieve this?
- A. Set the No Override option on the domain GPO.
 - B. Set the Block Policy Inheritance option on the domain GPO.
 - C. Set the Disable Domain Inheritance option on the domain GPO.
 - D. Unlink the domain GPO from the domain container.
- A.** Setting the No Override option on the domain GPO will prevent any lower-level GPO settings from being applied.
- Answer **B** is incorrect because the Block Policy Inheritance option prevents a container from inheriting GPO settings from a higher level. It does not have any effect on settings at a lower level. Answer **C** is incorrect because there is no Disable Domain Inheritance option. Answer **D** is incorrect because unlinking the GPO from the domain container will prevent the domain GPO settings from being applied at all.

Planning a Group Policy Strategy

5. You have been asked to implement group policy for a large, geographically diverse company. The users in the company are used to being able to log on very quickly, and you do not want to slow the logon process significantly when adding group policy settings. Which of the following are ways to reduce the processing time for group policy when a user logs on? (Choose all that apply.)
- A. Apply the Block Policy Inheritance setting on the OU closest to the logon object to keep all other policies from processing.
 - B. Set the Disable Computer Configuration Settings or Disable User Configuration Settings options in the GPO options.
 - C. Filter access to the GPO with WMI settings.
 - D. Filter access to the GPO with security group permissions.
- B** and **D**. Disabling the processing of computer or user configuration settings will reduce the time needed to process the GPO. Users and/or computers that do not have security permissions to see or process a GPO will not be able to process the settings.
- Answer **A** is incorrect because blocking policy inheritance will not prevent GPOs higher in the directory from processing; it will just override the settings once the GPO for that container is processed. Answer **C** is incorrect because using WMI filters in a GPO will actually increase the time it takes to process the policy settings, not decrease it.
6. You have been asked by your supervisor to duplicate the group policy settings of the Sales department for the Marketing department. A coworker suggests that instead of creating a new GPO for the Marketing OU, you can just link the existing Sales GPO to the Marketing OU. What are the guidelines for linking GPOs to a container?
- A. Each GPO can be linked to only one container.
 - B. Each GPO must be linked to a container within the same domain.
 - C. Only one GPO can be linked to the root domain container.
 - D. Each GPO should be linked to a single container only one time.
- D**. You can link a GPO to a container more than one time, but doing so can cause significant policy problems.
- Answer **A** is incorrect because GPOs can be linked to more than one container. Answer **B** is incorrect because you can link a GPO in one domain to a container in another domain, but you really should not do this. Answer **C** is incorrect because you can have more than one GPO tied to the root domain container.

7. You are the administrator for the corporate Active Directory network. There are four business units that are separated into individual domains that are rather large. How should you approach managing group policy for the corporation?
- A. Limit each business unit to one Default Domain Policy object in the root of each domain, and apply all policy settings for the domain in that object.
 - B. Identify one or more users in each domain and delegate control to them to create and manage group policy for the domain while retaining the ability to manage policy for each domain.
 - C. Give all users rights to manage group policy for themselves.
 - D. Only allow the administrator to manage group policy for the company.
- B.** Delegating control for group policy to trusted users in each domain can aid in the management of the needs for each domain. You should retain the ability to edit and manage policy in each domain in case of problems.
- Answer A** is incorrect because limiting each domain to a single GPO will likely not meet the needs of the users in the environment. **Answer C** is incorrect because granting all users rights to manage group policy would end up giving users too much power in the domain and represents a security risk. **Answer D** is incorrect because a sufficiently large directory will be difficult for a single administrator to manage effectively.

Implementing Group Policy

8. You just took over as network administrator for a company. Your network consists of a single domain. The previous administrator had set up a group policy for the domain that allowed six unsuccessful logon attempts before an account would be locked out. A series of new computers has been purchased and deployed in the environment, and the local policy on these systems is set to allow three unsuccessful logon attempts before locking an account. You decide that you want to enforce account lockout to occur after three unsuccessful logon attempts across the company. How would you achieve this?
- A. Set the local policy on each PC to lock out accounts after three attempts, and set No Override on the local policy.
 - B. Set group policy in a domain GPO to lock out accounts after three unsuccessful logon attempts.
 - C. Set the Block Policy Inheritance on the group policy.
 - D. Remove the local policies from each PC.
- B.** Since non-local group policy always overrides local policy, setting the account lockout threshold in group policy for a domain GPO will force all systems to have this new setting.

- Answer **A** is incorrect because you cannot set the No Override option on local policy. Local policy will always be overridden by non-local group policy. Answer **C** is incorrect because setting Block Policy Inheritance only affects non-local group policy, not local policy, and nothing has been done to change the number of unsuccessful logon attempts. Answer **D** is incorrect because you cannot remove local policy from computers, and even if you could, this would be a massive undertaking for nothing, as the remaining domain policy would still allow six unsuccessful logon attempts.
9. You need to create a new GPO to enable settings for a particular OU. You open Active Directory Users and Computers and select the OU in the tree. What is the next step in the process of creating a GPO for this OU?
- A. From the **Actions** menu, select **Create New GPO**.
 - B. Right-click on the OU and select **Create New GPO**.
 - C. Right-click on the OU and select **Properties**.
 - D. From the **Actions** menu, select **Group Policy Object Editor**.
- C**. Within the properties of the OU is the Group Policy tab, where you can create and edit GPOs for the OU.
- Answer **A** is incorrect because there is no Create New GPO option in the Actions menu. Answer **B** is incorrect because there is no Create New GPO option in the pop-up menu for the OU. Answer **D** is incorrect because there is no Group Policy Object Editor option in the Actions menu.

Performing Group Policy Administrative Tasks

10. You want to enforce minimum password lengths for all users in a particular domain. What is the best approach to doing this?
- A. Set the minimum password length policy in Computer Configuration | Windows Settings | Security Settings | Account Policies in the Default Domain Policies GPO.
 - B. Set the minimum password length policy in User Configuration | Windows Settings | Security Settings | Account Policies in the Default Domain Policies GPO.
 - C. Set the minimum password length policy in User Configuration | Windows Settings | Security Settings | Account Policies in the local policy for each computer on the network.
 - D. Set the minimum password length policy in User Configuration | Windows Settings | Security Settings | Account Policies for each OU in the network.

- A.** Password length settings are stored in the Computer Configuration settings. Applying these settings in the Default Domain Policy object will apply the settings to all computers in the environment.
- Answer **B** is incorrect because the password length settings are not located in the User Configuration settings. Answer **C** is incorrect because local policy settings for password length can be overwritten by group policy settings. Answer **D** is incorrect because password policies cannot be set at the OU level; they can only be set at the local or domain levels.
11. You have been asked to set up folder redirection for a particular set of users. Upper management wants these particular users to have a consistent interface on their computers, specifically the appearance of the Desktop and Start menu. These users will not be contained in a separate OU, and management does not want a separate policy created for this function. How will you accomplish this task?
- A. Set up Basic folder redirection settings in an existing GPO for the Desktop and Start Menu folders, and filter access to the redirection settings based on security group.
- B. Set up Basic folder redirection settings for the Start Menu, and Advanced folder redirection settings for the Desktop folder.
- C. Set up Advanced folder redirection settings for the Start Menu, and Basic folder redirection settings for the Desktop folder.
- D. Set up Advanced folder redirection settings for both the Desktop and Start Menu folders, specifying the specific security groups that should have the folder redirections.
- D.** When setting up Advanced folder redirection settings, you can select the specific security groups to which the folder redirection settings will apply, not impacting any of the other settings of the GPO.
- Answer **A** is incorrect because setting security filters on the entire GPO will restrict access to all GPO settings, not just the folder redirection settings. Answers **B** and **C** are incorrect because Basic folder redirection settings will apply to all users who access the policy, not just specific groups.

Applying Group Policy Best Practices

12. You have been asked by your project team to draft a policy document for managing group policy within your Active Directory environment. This policy document needs to include a summary of the best practices for implementing group policy. Which of the following statements would you include in your policy document? (Choose all that apply.)
- A. Keep the number of GPOs being processed to a minimum.
- B. Change Registry settings through Group Policy wherever possible.

- C. Assign security permissions on GPOs to individual users.
 - D. Maintain standard processing order whenever possible.
 - A** and **D**. Keeping the number of GPOs to a minimum helps reduce the amount of time needed to process policies. Maintaining the standard processing order makes it easier to troubleshoot policy problems.
 - Answer **B** is incorrect because using group policy as the primary method to change Registry settings can cause problems when those policies are removed. Answer **C** is incorrect because filtering group policy for individual users is more difficult to manage and troubleshoot.
13. One of the best practices for redirecting the My Documents folder is to let group policy create a folder for each user in a common path. Why should you avoid redirecting the My Documents folder to the user's home folder on the network? (Choose all that apply.)
- A. You cannot set exclusive rights on the user's home folder through group policy.
 - B. After you redirect the My Documents folder to the user's home folder, you will not be able to change the folder redirection settings.
 - C. You cannot redirect the user's My Pictures folder to the home folder.
 - D. Users must belong to the Redirected Folder Users security group, a setting that is often overlooked by system administrators.
 - A**. When redirecting the My Documents folder to the user's home folder, the existing permissions on the folder remain intact. The system cannot grant exclusive rights for the user to the folder.
 - Answer **B** is incorrect because you can always change the folder redirection settings, no matter which folder is redirected or where the folder has been redirected. Answer **C** is incorrect because the My Pictures folder can be redirected with the My Documents folder, no matter what location is chosen for the folder redirection. Answer **D** is incorrect because there are no special security groups that a user must belong to in order to participate in folder redirection.

Troubleshooting Group Policy

14. You have been asked to create a special policy environment for testing. You have been given the following requirements: Create a GPO called Test Settings in the root domain container. The settings of the Test Settings GPO should not apply to any users in Active Directory. You should be able to apply and remove the settings to/from an OU with minimal effort. Which of the following options meets these requirements? (Choose all that apply.)

- A. Set No Override at the domain level.
 - B. Rename the Test Settings GPO to break the link to other containers.
 - C. Set Block Policy Inheritance at the domain level.
 - D. Remove the link to the Test Settings GPO from the domain container.
- D.** Removing the link from the domain container will prevent the GPO from being processed by any users in the domain. You will still be able to link the GPO to other OUs as needed later.
- Answer **A** is incorrect because setting the No Override at the domain level will force the settings in Test Settings to apply to every user in the domain, regardless of any other GPO settings below. Answer **B** is incorrect because renaming the GPO will not break any links to any containers. Answer **C** is incorrect because setting Block Policy Inheritance at the domain will only impact GPO settings processed before the domain. It will not remove the Test Settings GPO setting from application.
15. A user complains that when he tries to save files to his My Documents folder, he keeps getting an error that he does not have permissions to write to the folder. He also tells you that when he looks at the files in his My Documents folder, he doesn't see any files that he recognizes. The domain policy you created redirects the My Documents folder to a secured share on the network. You suspect that someone has made a change to group policy elsewhere in the domain. How can you find the policy that is impacting folder redirection? (Choose all that apply.)
- A. Run an RSoP logging query for the user with his computer and look in the results for the policy objects applied to the computer.
 - B. Run an RSoP logging query for the user's OU and look in the results for the policy objects applied to the user.
 - C. Run an RSoP logging query for the user and his computer and look in the results for the policies applied to the user.
 - D. Run an RSoP planning query for the computer, ignoring the user settings, and look in the results for the policy objects applied.
- C.** Folder redirection policies are set at the user level. Looking through the results in the user configuration settings will tell you which folder redirection policy has been applied and which GPO applied the policy.
- Answer **A** is incorrect because folder redirection policies are in the user configuration, not in the computer configuration. Answer **B** is incorrect because you cannot run a logging query on an OU, only on a specific user. Answer **D** is incorrect because the folder redirection settings are in the user configuration settings, not the computer configuration settings.

Chapter 10: Deploying Software via Group Policy

Understanding Group Policy Software Installation Terminology and Concepts

1. As a network administrator, you need to deploy software to your user's workstations. What must you create on the network so the workstations will be able to locate, read, and execute the .msi file associated with an application?
 - A. A distribution point
 - B. An Admin C\$ Share
 - C. A share named "SFDeployment"
 - D. Nothing has to be created; everything necessary for software deployment is created by default when you select to deploy a package.

A. A distribution point is a share on a server or within the Dfs structure that holds the .msi files for users to install. You must create this network share and ensure that appropriate permissions are assigned; otherwise, software deployment will fail.

Answer B is incorrect because the C\$ admin share is created by default for administrators to use to access the root of the C: drive remotely. Each drive has a similar admin share that is created by the system. These shares are hidden and can only be used by administrators. You do not create the admin shares; in fact, if you remove them, they are recreated when the system reboots. In addition, you do not use them for software distribution because doing so would require giving administrative privileges to all users to whom the software will be deployed. **Answer C** is incorrect because there is no requirement that the shared folder have a specific name. The distribution point can have whatever folder name you want to give it. **Answer D** is incorrect because the Software Installation process does not create this share for you. It is necessary for you to set up the shared distribution point manually. This gives you the ability to locate the share in the best location to prevent deployment from taking place across a slow link.
2. When you assign an application to a user, which of the following actions on the part of the user will cause the software to be installed? (Choose all that apply.)
 - A. Clicking on the shortcut that represents the application in the Start menu or Desktop.
 - B. Opening **Add/Remove Programs** in Control Panel and selecting the application from the list.
 - C. Double-clicking a file that has an extension that is associated with the application.
 - D. Contacting an administrator to request pre-staging of the user's workstation.

- A** and **C**. When an application is assigned to a user, Software Installation places a shortcut to the application on the user's Desktop or in the Start menu. When the user clicks the icon for this shortcut to try to use the application for the first time, the software will automatically install. Installation of an assigned application can also be started via document invocation. When a user double-clicks a file that has an extension associated with the assigned application to open it, the application will be installed and then will open the file.
- Answer **B** is incorrect because an assigned application is not added to the list in Add/Remove Programs. This option is only available when the application is published. Answer **D** is incorrect because pre-staging a workstation is not a function of Group Policy software deployment.
3. What term describes what happens when a user double-clicks on a file with an associated extension that launches the installation of a package configured in Group Policy?
- A. Folder redirection
 - B. Document invocation
 - C. Blocking inheritance
 - D. No override
- B**. *Document invocation* is the process of invoking installation based on the appropriate file extension. This is also sometimes called *file extension activation*. Document invocation works for both published and assigned applications. When a user attempts to open a file, the system searches the Registry for file association information. If no information is found there, the system queries Active Directory for the information. If there is information there that points to a software distribution point, the Windows Installer package will be invoked, the application will be installed, and then it will open and open the file.
- Answer **A** is incorrect because *folder redirection* is a different component of Group Policy that doesn't deal with Software Installation. Answer **C** is incorrect because *blocking inheritance* deals with permissions, not software installation. Answer **D** is incorrect because *no override* is an option you set with the GPO itself and doesn't deal with Software Installation specifically.
4. You have configured Group Policy Software Installation to deploy several assigned and published applications. Which of the following is created automatically for each deployed application and stored in the GPO in Active Directory to contain advertisement information about the application configuration?
- A. Microsoft Installer Package
 - B. Logon Script
 - C. Application Assignment Script
 - D. Microsoft Software Transform

- ☑ **C.** The Application Assignment Script, with the file extension .aas, is created automatically for each application that is published or assigned within a Group Policy. These .aas files are scripts that contain advertisement information about the application configuration.
- ☒ Answer **A** is incorrect because it is the actual installer package file and doesn't contain the advertisement information. Answer **B** is incorrect because the logon script doesn't have any part of the software installation, but is used for drive mappings, printer mapping, and so forth. Answer **D** is incorrect because it is the file used to make modifications to an existing package and doesn't contain advertisement information.

Using Group Policy Software Installation to Deploy Applications

5. You want to add a new package to deploy to users in the Marketing OU. What are the steps required to complete this task?
 - A. Open **Active Directory Users and Computers**. Navigate to the Marketing OU, right-click, and select **Properties**. Click the **Group Policy** tab and select **Edit**. In the GPO Editor window, expand the **User Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.
 - B. Open **Active Directory Users and Computers**. Navigate to the Marketing OU, right-click, and select **Properties**. Select **Edit**. In the GPO Editor window, expand the **User Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.
 - C. Open **Active Directory Users and Computers**. Navigate to the domain node, right-click, and select **Properties**. Click the **Group Policy** tab and select **Edit**. In the GPO Editor window, expand the **User Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.
 - D. Open **Active Directory Users and Computers**. Navigate to the Marketing OU, right-click, and select **Properties**. Click the **Group Policy** tab and select **Edit**. In the GPO Editor window, expand the **Computer Configuration** node, select **Software Installation**, right-click, choose **New**, and then choose **Package**.
- ☑ **A.** To open the GPO and add the application to the Marketing OU so the software will be deployed at the OU level, you must use **Active Directory Users and Computers** to access the appropriate GPO. You do this by accessing the **Properties** of the OU, then selecting the appropriate GPO on the **Group Policy** tab and editing it. This opens the GPO Editor. From here, you select **User Configuration** because you want to deploy the software to users rather than computers. To add a package, you select **New** from the context menu, and then select **Package**.

- Answer **B** is incorrect because you have to select the **Group Policy** tab before you can select **Edit**, and this step is left out. Answer **C** is incorrect because you don't want to deploy the software to the entire domain, which is what you will accomplish if you follow these steps. Answer **D** is incorrect because following these steps would deploy the software to computers instead of to users, as specified in the scenario.
6. What steps should you take to set up a new category for software distribution?
- Right-click **Software Installation** in the GPO Editor and select **Properties**. Go to the **Categories** tab. Click **Add**.
 - Right-click **Software Installation** in the GPO Editor and select **Properties**. Go to the **General** tab. Click **Add**.
 - Right-click the package in the right pane of the GPO Editor, select **Properties**, and then click the **Categories** tab. Click **New**.
 - Right-click the package in the right pane of the GPO Editor, select **Properties**, and then click the **General** tab. Click **New**.
- A**. The Categories tab in the Properties of Software Installation is the proper interface for creating and managing categories into which you can place applications to make them easier for users to locate and install.
- Answer **B** is incorrect because the General tab does not allow you to administer category information. Answers **C** and **D** are incorrect because right clicking the package in the right pane of the console will allow you to configure its properties, but does not allow you to manage categories.
7. Which of the following tabs and options are used to force applications to update when deploying upgrades, regardless of whether the user wants to upgrade the application?
- “Required upgrade for existing package” on the Upgrades tab
 - “Mandatory upgrade” on the Upgrades tab
 - “Required upgrade” on the Modifications tab
 - “Mandatory upgrade” on the Modifications tab
- A**. The “Required upgrade for existing package” option exists on the Upgrades tab. Checking this check box will force an upgrade of the application on each computer on which it is installed, regardless of whether the user wants to perform an upgrade.
- Answers **C**, **B**, and **D** are incorrect because these options reference the wrong tab or the wrong options on the correct tab.
8. You are a network administrator and you have a number of legacy applications that need to be repackaged. You will be using WinINSTALL LE 2003 to create .msi packages for these applications. You have decided to set up a workstation that will be dedicated to cre-

ating these application packages. Which of the following is the best type of machine to use for this purpose?

- A. An existing computer configured with an optimized operating system that has been in service for at least two years.
 - B. A newly installed computer with a clean Registry and default operating system configuration.
 - C. An existing computer that has all of your organization's applications installed.
 - D. A new computer running only the critical applications but the original Registry settings.
- B.** A newly installed computer with a clean Registry and default operating system configuration is the best choice for this purpose. To avoid potential software conflicts and problems that can interfere with the creation of application packages, it is best to have a clean machine with a newly installed operating system, a clean Registry, and no extra applications installed. This makes the behavior of the system more predictable and makes it less likely that you'll have problems.
- Answer A** is incorrect because an existing machine that has been in service for two years could present problems, because it is likely that there have been programs installed and uninstalled, Registry settings modified, and other unknown configuration changes made during the time it has been in use. **Answer C** is incorrect because you do not want all of the organization's applications installed. Installing applications makes changes to the Registry and can install files that might conflict with or interfere with your ability to create clean package files. **Answer D** is incorrect because the original Registry settings would cause potential issues with application running properly.
9. You need to deploy an older application to users on your network, and you want to use Group Policy Software Installation to make the deployment easier. However, you are undecided as to whether you should use a program such as WinINSTALL to repackaging the application or use a .zap file to install it. In making this decision, you consider the limitations of using .zap files and whether these limitations will affect your deployment. Which of the following are limitations of using .zap files that should be factored into your decision? (Choose all that apply.)
- A. Applications deployed with .zap files cannot be installed by double-clicking a shortcut.
 - B. Automatic repair and removal doesn't work for applications deployed with .zap files.
 - C. Creating .zap files requires more programming expertise than repackaging applications does.
 - D. .Zap files install applications with elevated user privileges.

- A** and **B**. One limitation of deploying applications with .zap files is that document invocation does not work; thus the installation cannot be started by double-clicking a file associated with the application, as can be done with applications deployed with .msi packages. Another, perhaps more significant limitation, is that the automatic repair and removal features don't work with applications deployed via .zap files. If files for such an application become corrupt or are missing, the system will not automatically replace or repair them as it can with applications deployed via .msi packages. Moreover, you cannot select to have the application automatically removed if the GPO ceases to apply to the users to which the application is deployed.
- Answer **C** is incorrect because the creation of a .zap file requires less programming expertise than repackaging applications does. A .zap file is a simple text file created in a text editor such as Notepad that consists of one required section called [Application] and one optional section called [Ext]. Only two entries are required in the [Application] section: a FriendlyName entry and a SetupCommand entry. Answer **D** is incorrect because, unlike applications deployed via .msi packages, those deployed with .zap files cannot take advantage of the elevated user privileges that some applications use. This means that if an application requires administrative privileges to install, users to whom the software is deployed will have to be given administrative privileges, or the installation will fail. This is one of the most important reasons for repackaging legacy applications as .msi files instead of using .zap files.

10. You are the network administrator for a medium-sized financial services company. The users in the Accounts Receivable department all use a popular spreadsheet application that was deployed via Group Policy Software Installation. A new version of the spreadsheet program has been released and includes features that will be useful for some of your users; however, these new features are not required by all users. You want users who are comfortable with the old version and don't need the new features to be able to continue using their current version, but there are several new employees coming on board and you want them to start out with the latest version of the application. How can you disable new installations of an application but not remove the old application from users' workstations?
- A. Right-click **Software Installation** in the left pane of the GPO Editor and click **Removal**.
 - B. Right-click the application in the right pane of the GPO Editor and select **Remove**.
 - C. Right-click the application in the right pane of the GPO Editor and select **Uninstall**.
 - D. Right-click **Software Installation** in the left pane of the GPO Editor and select **Delete**.
- B**. The Remove option is available if you right-click on the application or package you want to remove. Selecting the optional removal option by clicking **Allow users to continue to use the software, but prevent new installations** will keep users from

installing the package, but will not prevent users who already have the application installed from continuing to use it.

- Answer **A** is incorrect because there is no Remove option in the context menu when you right-click **Software Installation**. Applications are removed individually by right-clicking the specific application you want to remove. Answer **C** is incorrect because Uninstall is not an option when you right-click the application. Answer **D** is incorrect because there is no Delete option in context menu when you right-click **Software Installation**.
11. You are the network administrator for a company that manufactures housewares. You need to deploy a particular software application to all members of the Sales department. All members of this department are already members of the Sales group in Active Directory. Now that you are setting up distribution of software to this same group of people, you would like to use their membership in this security group to define to whom the software will be deployed, if possible. How can you ensure that only the members of the Sales group will receive the software?
- A. You can move all their accounts to a newly created OU for deployment purposes.
 - B. You can use the Security tab on the GPO to configure the appropriate permissions for the Sales group only.
 - C. You can associate the GPO used to deploy the software with the domain so that new users will also receive the application.
 - D. You can remove the existing policy and create a new one that is applied to the group in question.
- B**. You can set permissions on the GPO that is used to deploy the software, so that only the appropriate individuals receive the software. Permissions can be assigned to individual user accounts or to group accounts. In this instance, you can give the necessary permissions only to members of the Sales group.
- Answer **A** is incorrect. Moving users to a new OU for deployment purposes will work, but it involves unnecessary work in this case. Answer **C** is incorrect. Associating the GPO with the domain might cause extra network traffic and could result in users receiving the software updates over a slow link, depending on your network topology. Answer **D** is incorrect. The original policy can be used, and creating a new one just results in unnecessary work.
12. You are the network administrator for a medical billing company. You want to deploy a new billing program to all the users in your organization. It is mandatory that every user have this new software installed. You decided to deploy the application via Group Policy Software Installation, but when you initially added the new package, you mistakenly configured the application to be published. You realize this will not accomplish your purpose because users can choose not to install the application. You want to change the application to assigned status. What is the best way to accomplish this?

- A. Right-click the application in the right pane of the GPO Editor, select **Properties**, and edit the **General** tab.
 - B. Right-click the application in the right pane of the GPO Editor, select **Properties**, and edit the **Deployment** tab.
 - C. Right-click **Software Installation** in the left pane of the GPO Editor, and click **Assigned**.
 - D. You will have to remove the package and add it again to change the status from published to assigned.
- B.** You can change an application's deployment type from published to assigned by editing the **Deployment** tab on the application's Properties sheet. To access this tab, right-click the application in the right pane of the GPO Editor, select **Properties**, click the **Deployment** tab, and select the **Assigned** option button.
- Answer **A** is incorrect because the only information that you can edit on the General tab of an application's Properties sheet are the package name and a support URL. There is no option on this tab for changing (or even viewing) the deployment type. Answer **C** is incorrect because there is no option to assign or publish when you right-click **Software Installation** in the left pane. To change an application's status, you must edit the properties of that specific application. Answer **D** is incorrect because, although it will work, it involves extra work and it is not necessary to go through the process of removing the package and adding it again, because you can change the deployment type after the package is added by editing the **Deployment** tab on the application's Properties sheet.

Troubleshooting Software Deployment

13. You are a network administrator who handles software installation for your organization. You are having problems with software installation for all of your applications. How would you turn on verbose logging for Software Installation when troubleshooting problems?
- A. Make a change to the GPO by selecting **Verbose Logging** on the **Security** tab.
 - B. A Registry change is required to HKLM\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics.
 - C. Make a change to the package by selecting **Verbose Logging** on the **General** tab of the package.
 - D. A Registry change is required to HKLM\Software\Policies\Windows\Installer.
- B.** The Registry needs to be modified with a `Appmgmtdebuglevel = 0000009b` DWORD to HKLM\Software\Microsoft\Windows NT\CurrentVersion\Diagnostics.

- Answer **A** is incorrect because there is no Verbose Logging check box on the Security tab. Answer **C** is incorrect because there is no Verbose Logging check box on the Security tab. Answer **D** is incorrect because it references the wrong Registry key for this type of verbose logging.
14. You need to deploy an application to all the users in an OU via Group Policy Software Installation. As part of the deployment process, you have set up a distribution point in which to place the .msi package that will be used to install the software. Users are having trouble installing the program, however, and you suspect that it is because of a permissions issue on the share. You need to ensure that you have set the appropriate permissions on the share point for the users. What are the minimal permissions that should be configured for users to be able to connect to the share and install the software?
- A. Read
 - B. Read and Execute
 - C. Modify
 - D. Full Control
- B**. Users must be able to access and execute software from this location in order to install it on their computers. This requires Read and Execute permissions. If users do not have these minimal permissions, the installation will fail.
- Answer **A** is incorrect because the Execute permission is needed to successfully deploy software. If users have only Read permission, they will be able to access the share, but they won't be able to run the executable program that installs the software; therefore, the installation will fail. Answers **C** and **D** are incorrect because both options give users higher permissions than those needed to install the software from the distribution point. This could result in users deleting folders from your distribution point or performing other changes that you don't want made. The question specified that you give users the minimal permissions needed to install the software, and these answers do not meet that requirement.
15. You are one of several network administrators who have deployed various software packages and updates within your domain. You have a user who is not receiving software installation as other users are in his department. He can log on and perform normal network functions. You are looking through your configuration to determine the possible cause. Which of the following are potential causes for this user not receiving the software that is associated with the group policy linked to his OU? (Choose all that apply.)
- A. Look for group policy conflicts that would cause this user to not receive the software.
 - B. Check permissions on the GPO to ensure that the user has Read and Apply Group Policy permissions.

- C. Check that the user account isn't disabled in Active Directory.
- D. Check permissions on the location of the software itself. The user needs at least Read permission to the distribution point.
- A, B, and D.** Group Policy has to have proper permissions assigned to the GPO for the user to be affected. In addition, having multiple group policies apply to a user can sometimes cause conflicts. Checking higher up in the hierarchy for conflicts to see if the user is affected. Proper NTFS permissions assigned to the distribution point are essential to users being able to access the needed files.
- Answer **C** is incorrect because if his user account were disabled, he wouldn't be able to log on and perform any functions at all. The question states that he can perform other network functions, so it has to be a problem with Group Policy itself.

Chapter 11: Ensuring Active Directory Availability

Understanding Active Directory Availability Issues

1. Your supervisor asks you how Active Directory knows which transactions have been committed to the database. You explain that this is tracked in a file known as:
 - A. Edb.log
 - B. Ntds.dit
 - C. Edb.chk
 - D. Edb00001.log
 - C.** The Edb.chk file, also known as the checkpoint file, keeps track of which transactions have been fully committed to the database. Transactions are “checked off” in this file only after they have been fully written and verified, not when the write process begins. If a failure occurs, this ensures that partially written and unwritten transactions are known and can be committed to the database fully upon recovery.
 - Answer **A** refers to the log file that is actively receiving new transactions, so this answer is incorrect. Answer **B** is incorrect because it refers to the Active Directory database file. Answer **D** is incorrect because it refers to a log file that is full and still has transactions that have not been committed to the database. Although all of these components have a hand in the process, none are primarily responsible for tracking what has and has not been committed to the Active Directory database.
2. You are studying the Windows Server 2003 Resource Kit in preparation for deployment of your company's new infrastructure. You learn that updates to the Active Directory database are transaction based. Which of the following are components of a transaction? (Choose all that apply.)

- A. Existing Active Directory data that has been altered by an administrator or by the replication process
 - B. New Active Directory data that has been altered by an administrator or by the replication process
 - C. All unchanged information that describes the object
 - D. Metadata
- A, B, D.** A transaction, as defined by Microsoft, consists of two components. The first is the changed data. This can consist of additions, deletions, or modifications to data stored in the Active Directory database. These changes can be made by an administrator or by the replication process, so Answers **A** and **B** are correct. Answer **D** is also correct. Each change is also accompanied by metadata. This can include the GUID for the object, property, or attribute that is being modified or other descriptive information, such as its version number.
- Answer **C** is incorrect. Only changed information and metadata comprise an Active Directory transaction. Existing but unchanged data is not replicated. This ensures maximum replication efficiency by minimizing the amount of data that is sent across the network. It also ensures maximum DC performance through the use of smaller write operations to commit the transaction data to the Active Directory database.
3. You are explaining the importance of system state data to your supervisor. Among other things, you tell her that the information it contains depends on the components installed in the underlying operating system. Which of the following do you tell her are always part of system state data on a Windows Server 2003 DC? (Choose all that apply.)
- A. The certificate database
 - B. The Registry
 - C. The Active Directory database
 - D. The metabase
 - E. SYSVOL
- B, C, E.** The Registry, COM+ Class Registration database, and boot/system files are always part of the system state data, so Answer **B** is correct. Answers **C** and **D** are also correct. The Active Directory database and related components such as SYSVOL are backed up as part of the system state data on a Windows Server 2003 DC.
- Answer **A** is incorrect. The Certificate Services database is only backed up on Windows Server 2003 computers that have the Certificate Services component installed. Although a DC can also be a certificate authority (CA), this is not true of all DCs; thus the certificate database is not always part of system state data on a DC. Answer **D** is incorrect. The metabase is an IIS component and is only backed up as part of system state data if IIS is installed. IIS is not installed on Windows Server 2003 computers by default, except on Web Server Edition, and computers running Web Server Edition cannot be DCs.

4. You are trying to explain the Active Directory deletion process to one of your coworkers. Despite your best efforts, he can't seem to grasp the concept of the Garbage Collection process. Which of the following key points do you reiterate? (Choose all that apply.)
- A. The Garbage Collection process runs every 12 hours by default.
 - B. An offline defragmentation of the database runs automatically at the end of the Garbage Collection process.
 - C. The Garbage Collection process deletes all expired objects from the Active Directory database.
 - D. The Garbage Collection process deletes all tombstoned objects from the Active Directory database.
- A, C.** Answer **A** is correct; the Garbage Collection process runs every 12 hours by default, although you can change this interval. Answer **C** is also correct, because the Garbage Collection process deletes expired objects from the Active Directory database. Objects are marked as expired at the end of the tombstone interval.
- Answer **B** is incorrect. Offline defragmentation must be performed manually by an administrator. Because the Garbage Collection process deletes objects from the Active Directory database, it produces fragmentation. The Active Directory database is defragmented at the end of the Garbage Collection process by an automatically run *online* defragmentation. Answer **D** is also incorrect. When an object is deleted from the Active Directory database, it is first marked as tombstoned and moved to the deleted objects container. This provides time for any other DCs to learn of the deletion. After the default tombstone interval (60 days) expires, the object is marked as expired. Although the objects removed by Garbage Collection have been tombstoned, they must have been tombstoned for 60 days and marked as expired to be deleted by Garbage Collection.

Performing Active Directory Maintenance Tasks

5. Your Windows Server 2003 DC is running out of disk space on the partition containing the database and log files. Which of the following steps will best rectify the situation before critical failure and additional performance degradation occurs? (Choose all that apply.)
- A. You can compress the database file using, NTFS compression.
 - B. You can install an additional hard drive.
 - C. You can move the Active Directory database and log files, using Windows Explorer, and reboot the server.
 - D. You can move the Active Directory database and log files, using the Ntdstutil utility

- B, D.** The best option is to either delete data from the hard drive to free up space, or move the database to another disk drive. Answer **B** is correct because it specifies adding a hard drive, which can be used in relocating the Active Directory database and log files. Answer **D** is also correct. The proper way to move the database is to boot into Directory Services Restore Mode and use the Ntdsutil command-line utility. This utility moves the database and log files and updates all Registry references to them.
- Answer **A** is incorrect. The Active Directory database is already heavily optimized for size. Using NTFS compression will have little effect on it and is not the best solution. Answer **C** is also incorrect. It is not possible to move the Active Directory database and active log files while Active Directory is initialized on the server. Booting to the Directory Services Restore Mode will allow it, but simply copying the files to a new location does not update their location pointers in the Registry. Upon reboot, the Windows Server 2003 will not be able to find the files and initialize Active Directory.
6. You recently started a new job as a network administrator. Upon reviewing the servers for which you are responsible, you notice that Active Directory is failing to start on one of your DCs. Further investigation reveals that Active Directory cannot find the location of the log files. You verify that the log files are located in the D:\ADLogs directory. Which of the following actions is the best to use in resolving this problem?
- A. You can delete the log files so that they are automatically recreated.
- B. You can boot into Directory Services Restore Mode and use the *set path logs* command within the Ntdsutil utility.
- C. You can boot into Directory Services Restore Mode and use the *move logs to* command within the Ntdsutil utility.
- D. You can directly modify the Registry so that it points to the new location.
- B.** The *set path logs* command within Ntdsutil allows you to tell Active Directory where the logs are currently located. It is designed to be used in situations exactly like the one described in the question. The command must be run from within Directory Services Restore Mode so that the Active Directory database is not initialized during the operation.
- Answer **A** is incorrect. In most situations like this one, Active Directory does not recreate the log files when they are deleted. Answer **C** is also incorrect. The *move logs to* command within Ntdsutil cannot be used when Active Directory is not aware of the current location of the log files. Answer **D** is incorrect. Although it might work, it is important to remember that Microsoft always recommends using a provided utility in lieu of directly modifying the Registry. Thus, this is not the best solution.

7. You have been promoted to network engineer and are training someone from the help desk to move up to your previous network admin role. You stress the importance of monitoring the log files each day to spot Active Directory problems early. Which of the following logs do you specify as important for the new admin to monitor? (Choose all that apply.)
- A. The System event log
 - B. The Application event log
 - C. The Directory Service event log
 - D. The File Replication Service event log
 - E. The DNS Server event log
- A, B, C, D, E.** All of these event logs are important to monitor. Answer **A** is correct because the directory service depends on the core functionality of the operating system, which writes messages to this event log. Answer **B** is also correct. The Application event log contains Group Policy and other related AD events. Answer **C** is correct because it is the primary log file to which AD writes messages. Answer **D** is also correct. The File Replication Service log deals with directory connection objects and replication of Group Policy. Finally, Windows 2000 and later clients depend on the DNS server service to locate Active Directory DCs, so Answer **E** is correct.
- All answers are correct.
8. The network engineer for whom you work is concerned about the impact of monitoring too many performance counters. She asks that you limit your monitoring to the core directory service performance object and its counters only. Which of the following will this allow you to monitor?
- A. DS Search sub-operations/sec
 - B. DRA inbound and outbound counters
 - C. Context Switches/sec
 - D. % DPC Time
- A, B.** The primary directory service performance object is the NTDS object. Answers **A** and **B** are counters within this object. DRA inbound and outbound counters are used to track the amount of replication information that flows into and out of a site. Significant changes can indicate a major increase in the amount of replication traffic or a shift in the site replication topology. Significant changes in the DS Search sub-operations/sec counter can indicate an application that is incorrectly targeting a DC, or performance problems involving the DC.
- Answer **C** is incorrect because this counter is tracked using the System object, not the NTDS object. Context Switches/sec tracks when the processor switches between waiting processes. This counter can indicate too many applications (including operating

system applications) for the processor to service, or applications that are too active for the processor to keep up with. Answer **D** is incorrect because this counter is tracked using the Processor object. % DPC Time is valuable in alerting you to delayed execution of processes resulting from the DC being too busy to execute them.

Backing Up and Restoring Active Directory

9. You've just taken over a new job and have been reviewing the backup logs. On one of your new DCs, the log shows that the Active Directory database is failing to back up. It also shows that several key components of the operating system, such as the Registry, are not even selected for backup. Which of the following is the best option for ensuring that these critical system components get backed up?
- A. Verify that Volume Shadow Copy is enabled
 - B. Select to back up system state data
 - C. Enable full system backup
 - D. Verify that there is sufficient free space on the drive to hold a shadow copy
- B.** Based on the question, it appears that the Active Directory database has been individually selected for backup. This is further proven by the fact that the Registry is not being backed up. The recommended method of backing up key system configuration components, such as Active Directory and the Registry, is to back up system state data, rather than the individual components.
- Answers **A** and **D** are incorrect because they are not the best options. It is possible that, because the Active Directory database appears to be selected for individual backup, Volume Shadow Copy is not enabled or that it is failing because there is not sufficient free space on the volume. However, this is not the best method to use in ensuring that these critical components are backed up. Answer **C** is also incorrect. Although this would solve the problem, nothing in the question suggests that a full system backup is required.
10. Your coworker has been working to script an automated restore of system state data but has not been successful. He asks for your help. What do you tell him?
- A. A macro is required to script a restore using Ntbackup.
 - B. A CMD file must be used for scripting a restore.
 - C. A BAT file must be used when scripting a restore.
 - D. It is not possible to script a restore using the Ntbackup utility.
- D.** Although Ntbackup can be used to back up system state data using the *systemstate* option, it cannot be used to restore data from the command line or a script.

Answers **B** and **C** are incorrect. Ntbackup cannot be used to script a restore, so neither of these standard files that are often used for scripting will work. You might be able to configure a macro to work with the graphical Ntbackup utility for restoring system state data. However, this is not scripting, so Answer **A** is incorrect.

11. Over the weekend, a tornado ripped through your company's Kansas office. You arrived to find no functioning DCs. The office has its own domain within your company's forest. Duplicate hardware is due to arrive in two days, and offsite backup tapes should arrive tomorrow. Which of the following restore methods will you use for the three DCs at the facility? (Choose all that apply.)

- A. Primary
- B. Normal
- C. Non-authoritative
- D. Authoritative

A, B. Answer **A** is correct. You should use the primary restore method to restore the first server that is brought back online for this domain. Answer **B** is also correct. Once the initial DC is reinstalled, the second and third DCs should be installed, using the normal restore method. This is because they will have a DC with a valid copy of the Active Directory database with which they can replicate after their restores are complete.

Answer **C** is incorrect. *Non-authoritative* was a term used in the Windows 2000 product line. This is now called the normal restore method in Windows Server 2003. Answer **D** is also incorrect. Authoritative restores are often used to reintroduce objects back into Active Directory that were accidentally deleted. This is not needed in the scenario provided.

12. One of the help desk employees has been assigned to assist with backups. She just sent you an e-mail stating that although she can back up her own files and some of the files in her department, she is unable to back up any other files on the network. Which of the following can enable her to back up any file? (Choose all that apply.)

- A. Her user account can be added to the local administrators group.
- B. Her user account can be granted the *Back up files and directories* right.
- C. Her user account can be added to the local backup operators group.
- D. Her user account can be added to a group that has been granted the *Back up files and directories* right.

A, B, C, D. Answer **A** is correct. Although not the ideal solution because of security concerns, adding her account to the local administrators group would allow her to back up any file on the system. Answers **B** and **D** are also correct. The *Back up files and direc-*

ories user right can be used to allow users and groups to back up files. Answer **C** is also correct, because membership in the local backup operators group allows a user to back up files on the local system.

All answers are correct.

Troubleshooting Active Directory Availability

13. You've been asked to train a new junior administrator on the basics of Active Directory troubleshooting. Which of the following do you tell him you will be covering? (Choose all that apply.)
- A. The Performance console
 - B. Event Viewer
 - C. DNS
 - D. DCPROMO
- A, B, C.** Answer **A** is correct, because the Performance console contains a number of counters that relate to Active Directory. Answer **B** is also correct. Most of the Active Directory related critical errors, warnings, and general messages are written to various Event Viewer logs. DNS is used by Windows 2000 and later clients to locate Active Directory, so Answer **C** is also correct.
- Answer **D** is incorrect. The DCPROMO command is used to create DCs and is not used for troubleshooting.
14. Your well-honed administrative skills are telling you that something is not quite right with your Active Directory database. You've reviewed the log files and performance monitor data. No users are complaining about connectivity problems. What can you do next to attempt to verify your hunch?
- A. You can run *Ntdsutil* with the *integrity* option.
 - B. You can run *esentutil* with the */g* switch.
 - C. You can run *esentutil* with the */p* switch.
 - D. You can run *Ntdsutil* with the *Semantic database analysis* option.
- A, B, D.** Answers **A** and **B** are correct. When the *Ntdsutil* command is run with the *integrity* option, it executes the *esentutil* command with the */g* switch. This command performs an integrity check on the **esentutil** portions of the AD database. Answer **D** is also correct. *Ntdsutil*'s *semantic database analysis* option is used to perform a full Active Directory database integrity check.

Answer **C** is incorrect. *Esentutl /p* performs a full binary repair of the database. It should be used only as a last resort repair effort, not when you simply suspect that there might be a problem.

15. The administrator with whom you work believes that the Active Directory problem she's troubleshooting could be solved if she could manually re-run the Active Directory log files. Which of the following commands does exactly that? (Choose all that apply.)

A. The *Ntdsutil integrity* command

B. The *Ntdsutil recover* command

C. The *esentutl /r* command

D. The *esentutl /g* command

B, C. Answer **B** is correct because the *Ntdsutil recover* command calls the *esentutl /r* command. Answer **C** is correct because the *esentutl /r* command performs a “soft” recovery of the database by re-running the log files.

Answer **A** is incorrect. The *Ntdstuil integrity* command calls the *esentutl /g* command. Answer **D** is also incorrect. The *esentutl /g* command performs an integrity check of the database but does not re-run the log files.

A

- .aas files, 652
- access
 - shared resources over networks, 12
 - to user account information, 3
- access control
 - Active Directory authentication, 61
 - group policy settings, 635
 - implementing in Active Directory, 55–60
 - role-based, 60
 - security groups and, 132
 - setting for OUs, 407
 - SIDS and, 86, 88
 - strong passwords, 184
- access control entries (ACEs), DACL listing, 46
- access control lists (ACLs)
 - adding users to, 395
 - and authentication process, 427
 - modifying, 220
 - security and discretionary, 56
 - use described, 86
- access levels, MMC User mode, 38
- access rights, 254
- access tokens, 480
- account lockout policies, 190–192, 217
- Account Operators group, 136
- Account tab, User's Properties (fig.), 118–120
- accounts
 - built-in, 101
 - built-in group, 135–140
 - computer. *See* computer accounts
 - group, 131–150
 - local and domain, 99
 - user. *See* user accounts
 - Windows Server 2003 assignment of, 84
- ACEs (access control entries), 46
- ACLs. *See* access control lists (ACLs)
- Active Directory
 - administrative tool types, 34–35
 - authentication. *See* authentication
 - availability. *See* Active Directory availability
 - backing up, 720–745
 - command-line tools, 35
 - components, 26–34
 - directory structure overview, 19–34
 - distribution of services information, 429–430
 - functionality levels, 65–66
 - Global Catalog. *See* GC (Global Catalog)
 - Group Policy integration in, 583
 - hierarchical structure of (fig.), 7
 - implementing security and access control, 55–64
 - Installation Wizard (DCPROMO), 15
 - installing to create domain controller, 15–19
 - introduction to infrastructure, 2
 - namespace hierarchical structure (fig.), 93
 - naming schemes, 11
 - performing maintenance tasks, 701–720
 - permissions, vs. share permissions, 56
 - relationship with DCs, members servers and clients (fig.), 6
 - replication service described, 32
 - restoring, 733–745
 - schema. *See* schema
 - sites. *See* sites
 - terminology, concepts, 5
 - Windows Server 2003 new features, 65–71
 - Windows Server 2003 version support, 9

- working with computer accounts, 150–162
- working with group accounts, 131–150
- working with user accounts, 99–130
- Active Directory availability
 - data commitment process (fig.), 693
 - Exam Objective Summary, 759
 - Exam Objectives Fast Track, 760
 - introduction, 690
 - Self Test, 764–769
 - troubleshooting, 745–758
- Active Directory Domains and Trusts
 - console described, , 40–43
 - creating, verifying, removing trusts with, 398–401
 - raising domain and forest functionality, 65–66
 - working with trusts, organizational units, 390–398
- Active Directory in Application Mode (ADAM), 249
- Active Directory Installation Wizard, 484–492
- Active Directory Object Manager tool, installing, 167
- Active Directory sites, Self Test, 468–474
- Active Directory Sites and Services
 - console described, 44
 - creating sites with, 436
 - creating subnets with, 442–444
- Active Directory Users and Computers
 - changing permissions, enabling, 57
 - console described, , 38–40
 - creating and managing user accounts with, 99–100
 - creating computer accounts with, 152–154
 - creating groups using, 141–143
 - creating user accounts with, 105–109
 - creating user objects with, 109–110
 - managing group accounts with, 143–149
 - moving account objects with, 164–165
 - opening, 308
 - Security tab, 76
- ADAM (Active Directory in Application Mode), 249
- adding
 - application package modifications, 673
 - computer to domain, 151
 - groups by nesting, 258
 - objects with dsadd, 49
 - snap-ins to MMCs, 36, 45
 - UPN suffixes, 162
 - users to ACLs, 395
- Address tab, User's Properties (fig.), 116
- ADIW. *See* Active Directory Installation Wizard
- administering application directory partitions, 314, 372
- administration
 - Group Policy, 612–617
 - policy-based administration, 9–10
 - schema, 556
- administrative logons, 206
- administrative rights, domain function and, 254
- administrative templates, Group Policy, 579
- administrative tools
 - GPO Editor, 659
 - Windows Server 2003 vs. Windows NT (table), 366
- Administrator account, 102
- Administrators group, 137
- ADSI Edit, 333, 694
- ADUC. *See* Active Directory Users and Computers
- Advanced tab, software installation properties, 665

- advertising software, 642
 - Alerts option, 714
 - application partitions
 - creating default DNS, 371
 - described, using, 8, 249
 - and DNS zones, 367–368
 - using, 313
 - applications
 - assigning with Group Policy, 644
 - assignment scripts, 652
 - categories of, 647
 - categorizing, 673
 - Group Policy deployment, 642
 - managing properties, 670
 - publishing via Group Policy, 646
 - removing managed, 669
 - upgrading, 667
 - ASR (Automated System Recovery), 361
 - assigning
 - applications, 644
 - applications via Group Policy, 662
 - master roles, 300–310
 - smart cards, 214
 - software to groups, 663
 - associating sites with subnets, 444–447
 - Asynchronous Transfer Mode (ATM), 446
 - attributes
 - and classes, 555, 567
 - default in Active Directory (table), 556
 - described, 541
 - in directories, 3
 - in GC, 547–548
 - indexing, 558
 - and object creation, 555
 - reused and defunct, 247
 - AttributeSchema objects, 555
 - authentication
 - Active Directory, described, 61–64
 - ADAM and, 249
 - creating user strategies, 192–195
 - cross-forest, forest-wide, 254
 - described, 427
 - digest, 199
 - Digest, 360
 - interactive logon, 428
 - Kerberos. *See* Kerberos
 - LDAP/SSL, 63
 - mutual, 196
 - need for, 193
 - network, 195, 428
 - Passport, 200–202
 - password, 193
 - PKI, 64
 - public and private keys, 64
 - server, 197
 - smart card, 203–217
 - strategies. *See* user and group strategies
 - trusts and, 317
 - types of, 195–203
 - UPN, 541–542
 - authorization
 - required levels for domain rename procedure (table), 321
 - roles, 60
 - Authorization Manager snap-in, 60
 - autoenrollment, 209, 612
 - Automated System Recovery (ASR), 361
 - automatic bridging, 459
 - Auxiliary objects, 554
- ## B
- backing up
 - Active Directory, 720–745
 - DCs (domain controllers), 361, 527
 - root domains, 247
 - system state, 503
 - Windows Backup, 298
 - Backup Operators group, 137

Backup or Restore Wizard, 721–732
 Backup utility, using, 721
 backups
 creating DCs from, 500
 scheduling, 727–732
 system state data. *See* system state data
 backward compatibility
 of naming schemes, 11
 to previous versions of Windows
 operating systems, 94
 bandwidth
 and GC servers, 548–549
 site requirements, 20
 Banyan, 4
 BDCs (Windows NT), 478
 best practices
 applying Group Policy, 622–623
 passwords, 491
 security groups, 224–229
 software restriction policy
 implementation, 622
 .bkf files, 298
 bridgehead servers, configuring, 459
 Builtin container, 39, 100, 136

C

caching, Universal Group, 248, 550–551
 Cacs tool, 46–47
 canonical names, 14, 147
 CAs (Certificate Authorities)
 and autoenrollment, 612
 and PKI, 206
 preparing, 331
 Categories tab, software installation
 properties, 666
 Cert Publishers group, 139
 Certificate Authorities. *See* CAs
 certificate revocation lists (CRLs), 207
 certificate rule, Group Policy, 621

certificates
 autoenrollment of, 612
 and smart cards, 206–208
 X.509, 63
 certtmpl.msc, 208
 certutil.exe, 359
 changing
 See also modifying
 group scope, 222
 permissions, 57
 characters
 not allowed in group names, 96
 in strong passwords, 102–103
 child and parent
 domains, 23–24
 objects and policy inheritance, 582
 circular transaction logging, 692
 classes
 and attributes, 555, 567
 disabling ability in Windows Server
 2003, 70
 dynamically linked auxiliary, 70
 object. *See* object classes
 object, described, 554
 reused and defunct, 247
 ClassSchema objects, 554
 clients
 logging onto domains, 11–12
 virtual private network (VPN), 182
 cmdkey tool, 47
 command-line tools
 Active Directory, 34–55
 command reference, 45–54
 DSADD, creating computer accounts,
 155–156
 getting help for, 50
 gpresult.exe, 627
 commands
 See also specific command

- Active Directory command reference, 45–54
- common names, 559
- compressed data, 550
- computer accounts
 - creating, 150–156
 - deleting and recreating, 173
 - introduction, 84
 - managing multiple, 156–163
 - naming, 12
 - resetting, 169
 - troubleshooting, 168
 - working with Active Directory, 150–162
- computer policies, 577–578
- computers
 - adding to domain, 151
 - deploying software via Group Policy to, 652–653
 - restricting machines users can log on from, 120–121
- Computers container, 39
- configuration data, and partitions, 7–8
- configuration partition, site information store, 435
- Configure Your Server Wizard, 478, 484
- configuring
 - bridgehead servers, 459
 - computer environment for Group Policy, 590
 - DNS servers for use with Active Directory, 370–373
 - GPOs, 596–604
 - OUs, 401–406
 - replication between sites, 457
 - site link availability, 458
 - site link costs, 449–452
 - software installation packages, 662
 - Universal Group caching (fig.), 551
 - user and computer security with Group Policy, 616
 - user environment for Group Policy, 588
- connecting domains in forests, 41
- consistency checker, 250
- consoles
 - Active Directory Domains and Trusts, 40–43, 390
 - Active Directory Sites and Services, 44
 - Active Directory Users and Computers, 38
 - MMCs, 35–38
 - Performance, 713
 - Replication Monitor tool, 461
 - WinINSTALL LE 2003, 656
- containers
 - automatically created with Active Directory installation, 100
 - created by default with MMC, 38–40 and OUs, 401
- contiguous namespaces, 246, 259
- copying volume shadow, 698
- costs, site link, configuring, 449–452
- counters, performance, 714–717
- creating
 - child domain in existing domain, 288–289
 - computer accounts, 150–154
 - DC from backup, 500
 - DCs, additional, 494
 - default DNS application directory partitions, 371
 - domain controllers, 15–19
 - domain password policy, 188–190
 - forest and domain structure, 275–364
 - forest root domain, 278
 - GC servers, 545
 - GPOs, 596–604
 - group accounts, 140–143
 - new groups, dialog box (fig.), 141

- new object classes, 561
- objects in Active Directory, 554
- objects with ldifde tool, 51–52
- OUs, 401–406
- replication topology, 455–456
- separate domains, 23
- site links, 446
- sites and site links, 434–438
- smart card enrollment station, 210–211
- subnets, 441–444
- system key, 185
- trusts, 395, 398
- user accounts, 105–109
- user authentication strategies, 192–195
- Windows Installer packages, 654
- .zap setup files, 658
- cross-forest authentication, 254
- cross-forest trusts, 377
- csvde tool, 47–48
- customizing GC with Schema MMC
 - snap-in, 543–544

D

- DACLs, use described, 86
- data, compressed, 550
- data access, 702
- data commitment process (fig.), 693
- data seeking, 702
- data store, directory, 5
- database, Active Directory
 - defragmenting, 702
 - event logging, 711–713
 - monitoring, 711
 - moving, 11–17–710
- database, transactional, 690
- DC locators, 327
- DC renaming tool, 70
- dcgpofix tool, 49
- dclist.xml, 342

- DCPROMO (Active Directory Installation Wizard)
 - described, 15–18
 - installing Windows Server 2003 as DC, 75
 - preparing to create forest and domain structure, 275–278
 - troubleshooting failure of, 293
- DCs (domain controllers)
 - backing up after domain rename, 361
 - creating additional, 293–300, 494
 - deciding when to create, 275–277
 - demoting to standalone server, 477
 - determining number of, 481
 - enabling GC functionality, 566
 - Exam Objective Summary, 523
 - Exam Objectives Fast Track, 524
 - function described, 480–481
 - installation of Active Directory and, 4
 - installing, 275–278
 - intra-site replication, 436
 - introduction, 476
 - master roles and, 27–28
 - outside of sites, 526
 - password control, 185
 - placing within sites, 502
 - promoting server to, 494
 - promoting Windows Server 2003 to, 29
 - relationship with Active Directory,
 - member servers and clients (fig.), 6
 - renaming, 256
 - in replication topology (fig.), 33
 - Self Test, 528–537
 - setting up DNS first, 489
 - upgrading, 500–502
 - upgrading before raising functional level, 272
 - Windows Server 2003 editions that can be used as, 75
- deactivating

- Schema classes and attributes, 562
 - schema or class, 247
- default trusts, 395
- defense in depth security, 183
- defragmentation, offline, 703, 762
- defragmenting Active Directory database, 702
- defunct classes and attributes, 562–563
- delegating
 - Group Policy administrative control, 608
 - OU control, 407
 - permissions and new system state backup method, 299
- Delegation of Control Wizard, 106, 410
- Delegation tab, computer account properties, 159
- deleting
 - objects with ldfide tool, 51–52
 - partitions, 315
- demoting DC to standalone server, 477
- Designated file types policy, 620
- designing
 - DNS infrastructure, 378
 - GC topology, 547
 - group strategy for domain forests, 225, 226
 - new forest, 338
- developers application, ADAM and, 249
- DFS (distributed file system)
 - renamed domains, effects of, 328
 - repairing topology after domain rename, 350
- dfsutil.exe, 350
- diagnostics tool, software installation, 678
- diagraming policies, 624
- Dial-in tab
 - computer account properties, 161
 - user account properties, 127–128
- digest authentication, 199, 360
- digital signing, 206
- direction of trust relationships, 315–316, 391
- directories
 - described, 2
 - information trees, 13
- directory data store, 5
- directory partitions, 313–315
- directory searches, 476
- directory services
 - Active Directory logical and physical components, 27–28
 - entries described, 13
 - history of, 4–5
 - introduction to, 2–3
- Directory Services Restore Mode, 703, 734, 758
- directory structure overview, 19–34
- disabling
 - classes, 70
 - Windows Server 2003 default security, 493
- discontiguous namespaces, 246
- discretionary access control list (DACL), 86
- disk mirroring, 700
- disk partitions, 8
- disk striping with parity, 701
- displaying information about logged-on users, 54
- distinguished name (DN), 13, 255
- distributed file system. *See* DFS
- distribution groups, 131–133, 218, 258–259
- distribution of services information in Active Directory, function, 429
- distribution points, 659, 681
- DNs (distinguished names), 13, 97–98
- DNS (Domain Name System)
 - and Active Directory namespaces, 366
 - checking primary suffix, 329

- configuring for redirecting alias entry, 354
 - described, 11
 - implementing in Active Directory
 - network environment, 365–373
 - infrastructure security, 367–368
 - interoperability with older servers, 276
 - namespaces, 245
 - naming process, 12
 - preparing for forest restructuring, 327
 - required SRV resource records (table), 345–346
 - securing, 373
- DNS registration, troubleshooting, 493
- DNS servers, configuring for use with Active Directory, 370–373
- DNS zones and Active Directory integration, 367
- DnsAdmins group, 139
- dnscmd, administering application partitions, 372
- dNSHostName attribute, 355
- DnsUpdateProxy group, 139
- document invocation, 646
- domain accounts, interactive logons and, 194–195
- Domain Admins group, 139
- domain and forest functional levels, 220
- domain and local accounts, 99
- Domain Computers group, 139
- domain controller rename, 256, 299
- domain controllers. *See* DCs
- Domain Controllers container, 39
- Domain Controllers group, 139
- Domain Guests global group, 103
- domain local groups, 135, 173, 220, 223
- Domain Name System. *See* DNS
- domain names and UPN suffixes, 95
- Domain Naming Master, 29, 301
 - FSMO, 506
 - locating DC, 305
 - seizing, 515
- domain password policies, 188–190
- domain rename utility, 70
- domain security IDs (SIDs), 30
- domain trees
 - creating in existing forest (fig.), 285
 - described, 23, 259
 - in forest (fig.), 25
 - reasons for using, 22
- domainlist.xml, 340–341
- domains
 - adding computer to, 151
 - authentication process described, 61
 - creating child in existing, 288–292
 - creating separate, 23
 - DC renaming tool, 70
 - described, 21
 - difference from OUs, sites, 415
 - DNS names, 11
 - and DNS zones, 367–368
 - establishing time standard for, 508
 - Exam Objective Summary, 374
 - function described, 254
 - functionality, 244–246, 260–265, 526
 - Group Policy settings, 580–581
 - managing with Active Directory
 - Domains and Trusts, 40–43
 - Microsoft use of term, 476
 - moving objects within, 50–51
 - moving within namespace, 248
 - and multiple sites (fig.), 21
 - name change guidelines, 339–340
 - OUs in (fig.), 26
 - post rename steps, 354
 - raising functional level, 71, 269–274
 - reasons for installing additional, 255
 - relationship to sites, 431–434
 - rename conditions and effects, 322–329

- renamed, distributed file system effects, 328
 - replication, 7, 331
 - restructuring renamed, 318–320
 - Self Test, 379–387
 - trust relationships between parent and child, 24
 - Universal Group capabilities (table), 258
 - vs. OUs, 409
 - Domains and Trusts Properties dialog box (fig.), 163
 - dragging and dropping objects into containers, 69
 - DSADD command
 - creating computer accounts with, 155–156
 - creating group accounts, 140, 142–143
 - creating user accounts with, 110–113
 - dsadd tool, 49
 - dsget tool, 49
 - dsmod tool, 50
 - dsmove tool, 50–51, 165–166
 - dsquery utility, 344
 - dumpfsmos.cmd, 511
 - duplexing, disk, 700
 - dynamic DNS, 245
 - dynamically linked auxiliary classes, 70
- E**
- ebay, 200–202
 - edb.chk file, 691
 - Edb*.log, 691
 - educating users about security, 202
 - elevation of privilege attack, 400
 - enabling
 - certificate enrollment in renamed domains, 355
 - certificate templates, 209
 - loopback mode, 590
 - permission-changing, 57
 - encryption
 - and PKI, 64
 - and restoring, 741
 - and smart cards, 203, 206
 - SMB, 493
 - SSL/TLS and, 197
 - syskey capability, 185
 - Enforcement policy, 620
 - enrollment stations, 209
 - Enterprise Admins group, 140
 - Environment tab, user account properties (fig.), 123
 - error events, 713
 - ESE (Extensible Storage Engine), 692, 693, 762
 - esentutl.exe, 747–750, 756, 763
 - establishing trust relationships, 315–317
 - event logging
 - Active Directory database, 711–713
 - with EventViewer, 462–463
 - EventViewer
 - monitoring Active Directory, 711
 - tool described, 462
 - Exam Objective Summaries
 - Active Directory availability, 759
 - Active Directory infrastructure, 72
 - Active Directory sites, 465
 - domain controllers, 523
 - forests and domains, 374
 - GC servers and schema, 564
 - Group Policy, 632
 - Group Policy software deployment, 679
 - trusts and organizational units, 412
 - user and group strategies, 230–231
 - user, group, and computer accounts, 170
 - Exam Objectives Fast Track
 - Active Directory availability, 760
 - Active Directory infrastructure, 73–74

- Active Directory sites, 465
- domain controllers, 524
- forests and domains, 376–377
- GC servers and schema, 565
- Group Policy, 633
- Group Policy software deployment, 679
- trusts and organizational units, 413
- user and group strategies, 232–234
- user, group, and computer accounts, 171
- exercises
 - Active Directory Schema snap-in setup, 544–545
 - adding and using alternative UPN suffixes, 163–164
 - checking primary DNS suffix configuration, 329
 - configuring DNS for redirecting alias entry, 354
 - creating account lockout policy, 191–192
 - creating Active Directory site, 436
 - creating an OU, 403–406
 - creating and modifying group accounts, 149–150
 - creating computer accounts using Active Directory Users and Computers, 154
 - creating DC in existing domain, 294–300
 - creating DC using new system state backup method, 297–300
 - creating domain password policy, 188–190
 - creating new domain in new forest, 278–284
 - creating new domain tree in existing forest, 285–288
 - creating smart card enrollment station, 210–211
 - creating subnets, 442–444
 - creating system key, 185–187
 - creating transitive, one-way incoming realm trust, 398–401
 - creating user objects in Active Directory Users and Computers, 109–110
 - deactivating classes and attributes, 562–563
 - delegating control for Group Policy to non-administrator, 605–608
 - folder redirection, 618–620
 - forcing synchronization of changes made to Domain Master DC, 344
 - generating current forest description file, 337
 - installing MOVETREE with Active Directory support tools, 168
 - locating domain operations Master, 306
 - locating schema operations master, 303
 - modifying garbage collection interval, 695–697
 - monitoring Active Directory with System Monitor, 717
 - performing offline defragmentation of database, 704
 - performing primary restore, 744
 - performing semantic database analysis, 755
 - promoting server to DC, 494–500
 - raising domain functional level, 271
 - raising forest functional level, 273
 - repairing DFS topology after domain rename, 351
 - repairing GPOs and links after domain rename, 352
 - RSoP querying in logging mode, 631
 - running RSoP planning query, 591
 - setting up control station for domain rename, 336
 - setting up smart card for user logon, 213
 - software modifications, working with, 674

- transferring a FSMO, 517–522
 - transferring Domain Naming Master role, 306
 - transferring Schema Master role, 304
 - unfreezing forest configuration, 350
 - verifying DCs' readiness, 346
 - verifying DNS service locator records, 345
 - verifying domain functional level, 270
 - viewing and modifying SIDs, 91–92
 - explicit trusts, 393–394, 414
 - extending schema, 560–561
 - Extensible Storage Engine (ESE), 692, 693
 - extensive defense security model, 183
 - extensions snap-ins, 36
 - external trusts, 43, 315, 396–397, 414
 - Extranet Access Management (EAM) solutions, 249
- F**
- FAT32, 69
 - fault tolerance and performance, 699
 - federated forest, 251, 254
 - file extension activation, 646
 - File Extensions tab, software installation properties, 666
 - filtering scope by security group membership, 582
 - filters, Windows Management Instrumentation (WMI), 586
 - finding. *See* locating
 - First Name attribute, 556
 - Flexible Single Master Operation (FSMO), 29, 257
 - folder redirection, 613, 618–620, 636
 - ForeignSecurityPrincipals container, 100
 - forest and domain functionality, 244–246
 - forest restructuring, 248
 - forest root domain
 - creating, 278
 - described, 246
 - renaming, 247–248
 - forest trusts, 41–42, 251–252, 393–394, 397
 - forests
 - described, 8, 24–25, 435
 - documenting, 336
 - federated, 251, 254
 - freezing configurations, 335
 - functionality levels, 65–66
 - functionality levels, 265–274
 - master roles and, 29–30
 - new domain in new (fig.), 278
 - raising functional level, 71, 269–274
 - restructuring, 248, 318–320
 - role of, 246–254
 - structure (fig.), 246
 - trust relationships and, 40–43
 - unfreezing configuration, 349–350
 - working with, Exam Objective Summary, 374
 - working with, Self Test, 379–387
 - forests trusts, 70
 - forward slashes (\\) and UNC paths, 13
 - freezing forest configurations, 335
 - frequency replication, configuring, 457
 - FSMOs (Flexible Single Master Operation)
 - forestwide, 505
 - reason for, 29, 300
 - roles, spreading across machines, 378
 - seizing master role, 310
 - transferring, 302
 - transferring roles, 509–522
 - functional levels
 - forest and domain, 259–274
 - switching domain, 526
 - functionality levels of Active Directory, 65–66

G

- garbage collection interval, 694
- GC (Global Catalog)
 - in Exchange server environment, 548
 - and forest trust trusted domain object, 252–253
 - function described, 541
 - and group membership changes, 32
 - introduction, 540
 - locating, 312
 - location in configuration partition, 8
 - and naming contexts, 256
 - number needed for sites, 483
 - replication, 547
 - and Schema Master, 256
 - search query (fig.), 542
 - Self Test, 567–573
 - troubleshooting, 551–552
 - UPN storage, 95
- GC servers
 - Exam Objective Summary, 564
 - Exam Objectives Fast Track, 565
 - placement of, 566
- General tab
 - group account properties, 144
 - NTDS settings, GC, 545–546
 - software installation package, 665
 - User's Properties (fig.), 115
- Global Catalog. *See* GC (Global Catalog)
- global groups, 173, 221–224, 223
- global security groups, 134
- Globally Unique Identifier (GUID), 247
- Gpfixup.log, 353
- GPO Editor, 645, 659
- GPOs (Group Policy Objects)
 - applied when, 415
 - creating, configuring, managing, 596–604
 - described, 10, 254, 401
 - evaluating, 411
 - local and non-local, 580
 - loopback mode, 590
 - number needed for sites, 635
- gpresult.exe, 627
- grfixup.exe, 336
- group accounts
 - built-in, 135–140
 - creating, 140–143
 - introduction, 84
 - managing multiple, 162–163
 - troubleshooting, 168
 - types, 131
- group policies, uses for, 9–10
- Group Policy
 - administrative templates, 579
 - applying best practices, 622–623
 - applying to OUs, 406
 - configuring application of, 604
 - controlling primary DNS suffix for computers, 330
 - delegating administrative control, 608, 636
 - Exam Objective Summary, 632
 - Exam Objectives Fast Track, 633
 - implementing, 595–612
 - integration in Active Directory, 583
 - introduction, 576–584
 - objects, 580–583
 - order of application, OUs, 411
 - performing administrative tasks, 612–617
 - placement of GPOs, 607
 - planning strategy, 584–590
 - propagation and replication, 583–584
 - repairing objects, links after domain rename, 352
 - replication management, 332
 - script capability, 121
 - Self Test, 637–639

- software deployment. *See* Group Policy software deployment
 - software restriction policies, 620
 - software settings, 578–579
 - troubleshooting, 624
 - using to predefine primary DNS suffix prior to domain rename, 334
 - verifying, 608
 - Windows settings (table), 579
 - .zap setup files, 656–659
 - Group Policy Creator Owners group, 140
 - Group Policy Object Editor, 595
 - Group Policy Objects. *See* GPOs
 - Group Policy software deployment
 - assigning and publishing software matrix (fig.), 647
 - configuring to users or computers, 652–653
 - Exam Objective Summary, 679
 - introduction, 642
 - opening or creating a GPO for, 660
 - patches, applying, 651
 - planning, 643–654
 - Self Test, 682–687
 - troubleshooting, 675–678, 681–682
 - vs. SMS software deployment, 648
 - Group Policy Wizard, 596
 - group scope, 133–135, 218, 219, 222
 - groups
 - See also specific group*
 - adding users to, 173
 - assigning software to, 663
 - membership limit, 259
 - naming, requirements for, 96
 - nesting, 223, 224
 - new user accounts, 130
 - scope conversions, 257
 - security, and OUs, 401
 - strategies. *See* user and group strategies
 - Universal Groups and group conversions, 257
 - Guest account, 103
 - Guests group, 137
- ## H
- hard disks, partitioning, 8
 - hardware redundancy, 699
 - hash rules, 198, 621
 - help for command-line tools, 50
 - Help and Support Service, SUPPORT_388945a0 account, 104
 - HelpAssistant account, 104
 - hierarchies
 - Active Directory namespace (fig.), 93
 - security group, 410–411
 - hosts
 - Active Directory vs. DNS knowledge of, 245
 - configuring member computers for name changes, 328
 - HTTP (HyperText Transfer Protocol) and URL names, 12
- ## I
- IAS (Internet Authentication Service)
 - group account, 140
 - identifying which computers are holding master roles, 169
 - IETF (International Organization for Standardization), 11
 - IIS_WPG group, 140
 - implementing
 - Active Directory replication, 467
 - DNS in network environment, 365–373
 - Group Policy, 595–612
 - OUs, 402–403
 - smart cards, 206–217
 - implicit trusts, 393–394, 414

implied trusts, 414
 Incoming Forest Trust Builders group, 137
 indexing attributes, 558
 InetOrgPerson accounts, 104–105
 InetOrgPerson objects, 69
 information events, 713
 Infrastructure Master, 30, 300, 301, 308–312, 509
 inheritance

- object, 58
- permissions, 129
- policy, 582

 Install from backups feature, 250
 installing

- Active Directory, 15–19
- Active Directory as DC, 4
- Active Directory Object Manager tool, 167
- MOVETREE.EXE, 166
- new DC, 275–278
- Schema snap-in, 544–545
- smart card reader, 212
- software via Group Policy. *See* Group Policy software deployment
- Support Tools, 461

 instmsi.exe, 649
 Integrity command, 747–750
 IntelliMirror technology, 642
 Inter-Site Transports container, 44
 interactive logon authentication, 205, 428
 International Organization for Standardization. *See* ISO
 International Telecommunications Union, The, 4
 Internet as DNS namespace, 93
 Internet Authentication Service. *See* IAS
 Internet Engineering Task Force (IETF), 11, 105, 197
 Internet Protocol. *See* IP

Internet zone rule, Group Policy, 621
 intersite replication, 453–455, 549
 intra-site replication, 32–34, 436, 453, 549
 IP subnets, 20, 426
 ISO (International Organization for Standardization), 4–5, 63

K

Kerberos

- authentication, 254
- authentication, and DCs, 480
- delegating across forests, 254

 Kerberos version 5

- authentication type, 195–197
- and forest trusts, 251
- protocol described, 62–63
- trusts and, 42

 Key Distribution Center (KDC), 62
 keys, public and private, 64
 Knowledge Consistency Checker (KCC)

- changing run time, 467
- creates replication topology, 453
- described, 32, 250
- and replication, 427

L

LDAP

- Active Directory naming scheme support, 11
- directory access protocol described, 10–11
- protocol described, 559
- standard described, 13

 LDAP/SSL authentication described, 63
 ldifde tool described, 51–52
 leaf objects (Novell), 26
 legacy servers, interoperability, 276

Lightweight Directory Access Protocol.
See LDAP

linked value replication, 250

links
 GPO, viewing, 604
 site, creating, 434–436

local accounts and domain accounts, 99

local area networks (LANs), sites and, 433

Local Security Authority (LSA), 61

locating
 Domain Naming Master DC, 305
 FSMOs, 511
 GCs, 312
 Infrastructure, RID, PDC Operations
 Master roles, 308–312
 objects within directories, 13
 user information, 541

lockout policies, 190–192, 217

log files
 and recover command, 750
 transactional, 690–692

logging
 setting Active Directory levels, 745–747
 transactional, 692
 verbose, 677

logical components
 of Active Directory, 245
 vs. physical components in Active
 Directory, 27

Logon Hours dialog box, user account
 properties (fig.), 120

Logon script field, user account
 properties, 121

logon scripts vs. group policy, 635

logons
 directory services and, 3
 faster with Universal Group caching,
 248
 interactive and network, 194–195
 KCC process and replication, 427

process over WAN links, 550

remote access and smart cards, 205
 and resource access, 61
 setting up smart card for users, 213
 smart card, 215
 user policies applied upon, 581
 viewing information about users, 54

logs, event, 462–463

loopback mode, 590

LostAndFound container, 39, 101

LSA (Local Security Authority), 61

M

Macintosh computers, Set Primary Group
 tab, user account properties, 130

Managed By tab, group account
 properties, 146–147

managing
 Active Directory with Ntdutil tool, 53
 application properties, 670
 certificate templates, 208
 computer accounts with Active
 Directory Users and Computers,
 156–162
 domain replication, 526
 domains and trust relationships, 40–43
 GC servers, 545
 GPOs, 596–604
 group accounts with Active Directory
 Users and Computers, 143–149
 multiple accounts, 162–163
 Operations Masters, 505–522
 OUs, 401–406
 replication topology, 456–457
 SIDS, 90–92
 Universal Groups, 257
 user accounts through policy
 administration, 9–10
 user accounts with Active Directory
 Users and Computers, 113–118

- mapping
 - See also* topology
 - policy problems, 624
 - master roles
 - assigning and transferring, 300–310
 - identifying computer holding, 169
 - used in Active Directory forests, 29
 - Media Player, 580
 - Member Of tab
 - computer account properties, 157–158
 - group account properties, 146
 - user account properties, 130
 - members servers, 477
 - Members tab, group account properties, 145
 - metropolitan area network (MAN), 446
 - Microsoft directory services, history of, 4
 - Microsoft Management Consoles. *See* MMCs
 - migration
 - ADAM and, 250
 - DC upgrade, 501
 - mirrored volumes, 700
 - MMCs
 - Authoring and User mode, 37
 - described, 35–38
 - snap-ins. *See* snap-ins to MMCs
 - modifying
 - access control lists (ACLs), 220
 - GC (Global Catalog), 543–544
 - objects with *ldifde* tool, 51–52
 - password policy, 190
 - schema, 560–561, 566, 567
 - monitoring
 - Active Directory, 763
 - Active Directory database, 711
 - performance counters, 714–717
 - MOVETREE command, moving objects with, 166–168
 - moving
 - account objects in Active Directory, 164–167
 - Active Directory database or log files, 711–710
 - objects within domains, 50–51
 - servers to different sites, 467
 - .msc files, 37
 - msDS-AllowedDNSSuffixes attribute, 332
 - msDS-UpdateScript attribute, 342, 347
 - .msi files, 649, 651
 - MsiInstaller messages, 675
 - .msp files, 651
 - .mst files, 650, 651
 - multidomain DNS Conditional Forwarding, 367–368
 - multimaster replication, 32
 - multiple-forest trusts, 251
 - multivalue attributes, 557
 - mutual authentication, 196
- ## N
- names
 - common, 559
 - relative distinguished and canonical, 14
 - URL, 12
 - namespaces
 - Active Directory vs. DNS, 245
 - contiguous vs. discontinuous, 246
 - DNS and Active Directory, 366
 - domain tree in contiguous (fig.), 23
 - hierarchies, 244
 - moving domains within, 248
 - need for understanding, 379
 - naming
 - See also* renaming
 - GPOs, 596–604
 - Schema objects, 559
 - naming contexts, 249, 255, 313
 - naming schemes

- DNS (Domain Name System), 11
 - schema, 553
 - NBTStat command, 363
 - nesting
 - to add groups, 258
 - groups, 223, 224
 - Net Share command, 300
 - net time command, 254
 - NetBIOS
 - name for new user account, 106–107
 - name of computers, 94
 - name resolution, 254
 - names and Active Directory, 367
 - naming scheme, 11
 - suffixes, Knowledge Base article about, 363
 - netdom query fsmo command, 168
 - Netdom.exe, 317, 363
 - Netlogons.dns, 276
 - NetMeeting, 580
 - Netscape, revocation check mechanism, 358
 - NetWare Directory Services, 4
 - network authentication, 195, 198, 428
 - Network Configuration Operators group, 137
 - Network Identification tab, System Properties dialog, 151
 - network traffic and GC server placement, 549–552
 - networks
 - accessing shared resources over, 12
 - physical vs. logical structure, 433
 - schema, importance of, 552
 - New Object Computer dialog boxes, 152–154
 - New Trust Wizard, 317, 395
 - nodes, Users and Domain Controllers, 100–101
 - non-local policies, 577
 - non-transitive trusts, 42, 391–392
 - Notebook and renaming domains, 339
 - Novell leaf objects, 26
 - NT LAN Manager and network authentication, 198
 - NTDS
 - and Active Directory, 3
 - GC settings, 545–546
 - monitoring performance objects, 717
 - Site Settings Properties box, 550, 564, 566
 - Ntds.dit, 277, 690
 - ntdutil tool, 53, 90–92, 314, 710, 747–750, 763
 - NTFS file system and DACLs, 87
 - NTLM authentication, 198, 254
- ## O
- object classes
 - creating new, 561
 - described, 554
 - object identifier (OID) , 247, 559
 - object inheritance, 58
 - Object tab, group account properties, 148
 - objects
 - Auxiliary, 554
 - ClassSchema, 554
 - in directories, 3
 - directory service, 84
 - distinguished name (DN) of, 13
 - DN used to show location (fig.), 98
 - dragging and dropping into containers, 69
 - Group Policy, 580–583
 - modifying with dsmod, 50
 - quotas, 69, 250
 - schema, defunct, 247
 - setting permissions on, 57
 - site link, 33, 446

- viewing properties with dsget, 49
- offline defragmentation, 703
- OID (object identifier), 559
- Open Systems Interconnect (OSI), 5
- opening Active Directory Users and Computers tool, 308
- Operating System tab, computer account properties, 157–158
- Operations Masters (OMs)
 - assigning and transferring master roles, 300–310
 - described, 32, 505–522
 - valid authorization levels (table), 301
- operations master roles, transferring, 509–522
- Organization tab, User's Properties (fig.), 118
- organizational units. *See* OUs
- OSI (Open Systems Interconnect), 5
- Other Organization SID, 254
- OUs (organizational units)
 - applying Group Policy to, 406
 - containers and, 401
 - delegation, 409, 415
 - described, 25–26
 - difference from sites, domains, security principals, 415
 - Group Policy settings, 580–581
 - RDNs and, 14
 - vs. domains, 409

P

- packages, Windows Installer, 649
- parent and child domains, 23–24
- parent and child trusts, 324–325
- parity, 701
- partitions
 - application, using, 249, 313–315
 - described, 7

- directory and disk, 8
 - naming contexts, 255
 - replication and, 65–66
- Passport authentication, 200–202
- passwords
 - See also* permissions, user rights
 - account lockout policies, 190–192
 - Administrator account, 102
 - applying a policy, 187–190
 - best practices, 491
 - changing Directory Services Restore Mode, 758
 - creating policy for domain users, 182–192
 - and new system state backup method, 298
 - orchestrating reset for digest authentication, 360
 - replication and, 30
 - strong, 102–103, 184
- patches, applying with Group Policy, 649, 651
- path rule, Group Policy, 621
- paths, trust, 316–317
- PDC Emulator
 - described, 30, 300, 301
 - locating, transferring, seizing Operations Master roles, 308–312
 - seizing, 514
 - transferring, 513
 - troubleshooting, 527
 - working with DCs, 507
- performance
 - DC functions affecting (table), 482
 - fault tolerance and, 699
 - issues related to number of users, 479
- Performance Log Users group, 137
- Performance Monitor Users group, 137
- permissions
 - See also* passwords, user rights

- access control, 55–60
- changing, 57
- configuring object's, 128
- normal domain default, changing, 150
- OUs, setting, 407
- setting in Active Directory, 75–76
- setting on Active Directory objects, 58–59
- setting security, 208
- standard and special, 56
- vs. rights, 219
- physical components
 - of Active Directory, 245–246
 - vs. logical components in Active Directory, 27
- physical vs. logical structure of networks, 433
- PIN numbers, 215
- PKI authentication, 64, 206
- placement of GPOs, 607
- planning
 - Active Directory sites, 434–436
 - DCs (domain controllers), 476–484
 - domains or OUs, 409
 - GC server placement, 548
 - Group Policy, 583
 - for Group Policy software deployment, 643–653
 - Group Policy software deployment, 654
 - Group Policy strategy, 584–590
 - OU structure and strategy, 408–411
 - policies with RSoP, 625
 - replication topology, 455–456
 - security group strategy, 217–229
 - smart card support, 216
- policies
 - account lockout, 190–192, 217
 - applying password, 187–190
 - local and non-local, 577
 - precedence of, 621
 - processing setting at user logon (fig.), 581
 - restoring with dcgppfix, 49
 - setting on OUs, 401
 - Software Installation, 589
 - user and computer, 577–578
 - viewing gpresult results, 627–630
 - viewing sample policy map (fig.), 624
- policy-based administration, 9–10
- policy inheritance, 582
- Pre-Windows 2000 Compatible Access group, 138
- primary restore method, 743
- Print Operators group, 138
- privileges
 - See also* passwords, permissions, user rights
 - and access rights, 132
 - using .zap files, 656
- Profile tab, user account properties, 212–222
- profiles, roaming, 328, 588
- propagation, policy, 583–584
- properties
 - application, managing, 670
 - attribute (fig.), 555
 - GPOs, viewing (fig.), 604
 - objects, viewing with dsget, 49
 - OU's, 403
 - software installation, configuring, 664
- protocols
 - See also specific protocol*
 - directory access, 10–11
 - Windows Server 2003 authentication support, 62
- Published Certificates tab, user account properties, 126–127
- publishing
 - applications via Group Policy, 646, 662
 - services in Active Directory, 430
 - software using .zap files, 658

pull model, Group Policy, 648
 push instructions for DC rename, 344
 push model, SMS, 648

Q

queries

GC search (fig.), 542–543
 RSoP planning, 591

quotas

described, 250
 of objects in Active Directory, 69

R

RAID, 277, 699–701

raising

domain and forest functional levels,
 65–66, 71, 269–274

functional level, issues, 377

RAs (registration authorities), 64

RAS and IAS servers group, 140

RAS (Remote Access Service) default
 group, 140

RDNs (relative distinguished names), 14

realm trusts, 41–42, 315, 396, 414

rebooting computers, 219

recover command, 750

redirecting folders, 613–616, 618–620,
 636

registration authorities (RAs), 64

Registry

and application publication, 646
 changing on DCs to reflect domain
 rename, 356
 enabling verbose logging, 677
 setting to apply to local computers
 through Group Policy, 576

relationships

of sites and domains, 431–434

trust. *See* trust relationships

relative distinguished name (RDN), 14,
 95

Relative ID (RID) Master, 85

remote access

smart card authentication, 205
 VPN, 127

Remote Access Service. *See* RAS

Remote Control tab, user account
 properties, 124–125

Remote Desktop Users group, 138

Remote Installation Services. *See* RIS

Remote Procedure Calls (RPCs), 453

remote setup, creating DCs using new
 system state backup method,
 297–300

removing

application package modifications,
 673

managed applications, 669

trusts, 398–401

renaming

DCs, cautions about, 363

DCs (domain controllers), 256

domain controller rename, 299

domains, conditions and effects,
 322–329

domains in forests, 318–320

forest root domain, 247–248

objects with dsmove, 50–51

sites, 438–441, 468

Windows Server 2003 DC, 363–364

rendom.com, 70

Rendom.exe, 320

rendom.exe, 336, 337

replicas, 313, 315

replication

application partitions and, 249

configuring between sites, 457

forcing Active Directory, 344

- GC, 547
 - Group Policy, 583–584
 - intersite and intrasite, 436, 453–455, 549
 - linked value, 250
 - new to Windows Server 2003, 70
 - PDC Emulator and password, 30
 - process described, 7
 - restricting traffic to specific port, 467
 - ring topology (fig.), 454
 - service described, 32
 - site. *See* site replication
 - site default tombstone lifetime, 503
 - storms, 330
 - synchronizing data, 427
 - topology, 32–33
 - troubleshooting failure of, 460
 - units of, 431–432
 - universal security groups and, 134
 - USNs and, 147
 - Replication Monitor tool, 461
 - replmon.exe, 462
 - resource sharing between Active Directories, 252
 - restarting computers after domain rename, 362
 - Restore Wizard, 298
 - restoring
 - Active Directory, 733–745
 - authoritative, 741–743
 - DCs (domain controllers), 504
 - Directory Services Restore Mode, 703
 - new system state backup method, 298
 - normal restore, 734–741
 - system state data, 503, 698
 - restricting
 - forestwide replication, 134
 - machines that a user can log on from, 120–121
 - replication traffic to specific port, 467
 - software through Group Policy, 620
 - restructuring forests, pre-creating parent-child trusts (fig.), 324
 - Resultant Set of Policy (RSOP) tool, 584–588
 - revoking smart cards, 215
 - RID Master
 - assigning and transferring master roles, 300–302
 - described, 30
 - locating, transferring, seizing
 - Operations Master roles, 308–312
 - one per domain, 90
 - transferring, 513
 - troubleshooting, 527
 - ring topology for replication (fig.), 454
 - RIS (Remote Installation Services) and creating computer accounts, 153
 - roaming profiles, 328, 588
 - role-based access control, 60
 - roll back ability, Windows Installer, 649
 - rolling out domain name change, 334
 - root domains, forest, 25
 - routing hints for forest trusts, 253
 - RSA smart cards (figs.), 204
 - RSOP tool
 - described, 584–588
 - querying in logging mode, 631
 - using to plan policy, 625
 - which computers to run on, 635
 - RSOP Wizard, 631
- S**
- SAM (Security Account Manager), 12
 - SASL protocol, 199

- SCEP (Simple Certificate Enrollment Protocol), 358
- scheduling backups, 727–732
- schema
 - administration, 556
 - components described, 553
 - deactivating classes and attributes, 562
 - defunct objects, 247
 - extending, 51
 - InetOrgPerson object class, 105
 - modifying and extending, 560–561
 - object classes, 554
 - partitions, 7–8
 - troubleshooting issues, 563
 - working with Active Directory, 552
- Schema Admins group, 140
- Schema FSMO, transferring, 510
- Schema Master
 - described, 29, 301
 - FSMOs, transferring, 515
 - and GC, 256
 - locating operations, 303
 - use described, 506
- Schema Microsoft Management Console snap-in
 - customizing, 543–544
 - described, 540, 560–563
- Schema objects, naming, 559
- Schmmgmt.dll, 303, 544
- scope described, 133
- scope conversions, allowable, 257
- scripts
 - ADSI, 694
 - application assignment, 652
 - in Group Policy, 576
- searching
 - directory information, 542
 - for objects, 69
 - user information, 541
- securing
 - Administrator account, 102
 - DNS deployment, 373
 - trusts using SID filtering, 400
- security
 - Active Directory authentication, 61
 - boundaries and OUs, 401
 - configuring user and computer settings, 616–617
 - creating extensive defense model, 183
 - delegation issues, 409
 - disabling Windows Server 2003 default, 493
 - educating users, 202
 - GPO settings, 605
 - implementing in Active Directory, 55–60
 - permissions and restoring, 741
 - setting permissions, 57
 - types of, 218
 - user account settings, 126–130
- Security access control list (SACL), 56
- Security Account Manager (SAM), 12, 99
 - granting access to standalone servers, 477
 - interactive logon authentication, 428
- security descriptor reference domain attribute, 313–315
- security groups
 - best practices, 224–229
 - filtering scope by membership, 582–583
 - hierarchy, 410–411
 - nesting, 258
 - planning strategy, 217–229
 - use described, 131–133
- Security ID Histories (SIDHistory), 65
- security permissions, settings, 208
- security principals
 - difference from OUs, 400, 415
 - and SIDs, 85

- Security Sockets Layer (SSL), 63
- Security tab
 - Active Directory Users and Computers snap-in, 76
 - computer account properties, 160–161
 - group account properties, 148
- seeking, data, 702
- segregating subgroups, 288
- seizing
 - FSMO master role, 310
 - FSMO roles, 505
 - Infrastructure, RID, PDC Operations
 - Master roles, 308–312
 - operations master roles, 509–522
 - PDC Emulator, 514
 - RID Master, 515
 - role, issues about, 378
 - schema master role, 303
- selective authentication, 254
- Self Tests
 - Active Directory availability, 764–769
 - Active Directory infrastructure, 76–81
 - Active Directory sites, 468–474
 - GC servers and schema, 567–573
 - Group Policy, 637–639
 - Group Policy software deployment, 682–687
 - trusts and organizational units, 416–423
 - user and group strategies, 234–241
 - user, group, and computer accounts, 174–180
 - working with DCs, 528–537
 - working with forests and domains, 379–387
- semantic database analysis command, 752–755
- server authentication, 197
- Server Message Block (SMB) signing, 493
- Server Operators group, 138
- servers
 - authentication, 476
 - DNS, interoperability, 276
 - GC, working with, 540–552
 - member, 477
 - moving to different sites, 467
 - promoting to DC, 494
 - service location resource record (SRV), 365
 - service tickets, 62
 - servicePrincipalName attribute, 332
 - Sessions tab, user account properties, 123–124
 - setting permissions on Active Directory objects, 58–59
 - settings
 - Active Directory logging levels, 745–747
 - security permissions, 208
 - user account, 118–120
 - share permissions, 56
 - shared resources
 - between Active Directories, 252
 - names, and accessing, 12
 - shortcut trusts, 41, 315, 323, 392, 395–396, 414
 - SID filtering, 400, 414
 - SIDs (domain security IDs)
 - access control use (fig.), 86
 - described, 30
 - naming conventions and limitations, 92–98
 - Other Organization, 254
 - and security principals, 85
 - tracking account information, 480
 - viewing and managing, 90–92, 173
 - well-known (table), 88–90
 - sign-on, single, 427
 - Simple Certificate Enrollment Protocol (SCEP), 358

- Single Instance Store (SIS), feature
 - described, 251
- single sign-on, 194–195, 427
- single-value attributes, 556
- site link bridges, 458
- site link costs, 449
- site link objects, 33
- site links
 - configuring availability, 458
 - creating, 436–438, 446
 - disabling transitivity, 468
- site replication, purpose of, 453
- sites
 - associating with subnets, 444–447
 - configuring replication between, 457
 - DCs outside of, 526
 - described, 20–21
 - difference from OUs, domains, 415
 - Exam Objective Summary, 465
 - Exam Objectives Fast Track, 465
 - Group Policy settings, 580–581
 - inactive, deletion of inactive objects, 503
 - introduction to Active Directory, 426
 - minimum DC, GC, RAM, and CPU
 - (table), 483
 - with multiple domains (fig.), 21
 - placing DCs within, 502
 - placing GCs within, 548–549
 - planning, 434–436
 - relationship to other Active Directory
 - components, 431–434
 - renaming, 438–441, 468
 - replication and, 32–33
 - with several client computers (fig.), 434
- slashes, forward (\\) and UNC paths, 13
- smart cards
 - enrollment stations, 209
 - implementing, 206–217
 - installing readers, 212
 - logon procedure, 215
 - planning authentication strategy, 203
 - planning for support, 216
 - revoking, 215
 - supported under Windows Server 2003
 - (table), 212
- SMS software development vs. Group Policy software deployment, 648
- SMTP replication, 458
- snap-ins to MMCs
 - adding, 15, 35–38
 - Authorization Manager, 60
 - Group Policy Object Editor, 595
 - Schema, working with, 560–563
 - Security Policy, 360
- SNTP, establishing time standard for
 - domain, 508
- software
 - installing via Group Policy. *See* Group Policy software development
 - package, 643
 - removing managed applications, 669
 - SMS software development, 648
 - transforms, installer packages, 650
- Software Installation Diagnostics tool, 678
- Software Installation policy, 589
- software restriction policies, 620–622, 637
- source DC, 456
- Special Permissions dialog box, user
 - account properties, 129–130
- split DNS, 367–368
- SSL (Security Sockets Layer)
 - described, 63
 - Transport Layer Security (TLS), 197
- standalone servers, 476
- standalone snap-ins, 36
- standards
 - Windows Server 2003 authentication
 - support, 62
 - X.500/LDAP, 13
- storms, replication, 330

- strategies
 - designing group, for single domain forest, 225
 - planning security group, 217–229
 - planning smart card authentication, 203–206
 - for raising functional levels, 273–274
 - StreetTalk, 4
 - strong passwords, 102–103, 184
 - structure, network, physical vs. logical, 433
 - subgroups, segregating, 288
 - subnets
 - associating with sites, 444–447
 - creating, 441–444
 - described, 426
 - well connected sites, 20
 - Subnets container, 44
 - Sun Microsystems, 182
 - Support Tools, 461, 464
 - SUPPORT_388945a0 account, 104
 - synchronizing
 - changes made to Domain Naming Master, 344
 - data, and replication, 427
 - syskey, 183, 185
 - system access control lists (SACLs), 86
 - System container, 40, 101
 - system key utility, 183, 185
 - System Monitor, 714
 - system state data, 503, 698, 721, 762
 - System State restoration, Windows Backup, 298
- T**
- TAPI (Telephony API), 314, 340, 359
 - target DC, 456
 - TCP/IP
 - communication between sites, 444
 - and SSL, 197
 - TDO (trusted domain object), 253
 - Telephones tab, User's Properties (fig.), 117
 - Telephony API (TAPI), 314
 - templates
 - certificate, 207
 - Group Policy administrative, 579
 - Terminal Services
 - administrative templates, 580
 - ending sessions, 124
 - tabs, user account properties, 122–126
 - testing
 - patches, 651
 - software restriction policies, 622
 - Ticket Granting Ticket (TGT), 62
 - tickets (Kerberos), 62
 - time service, 508
 - TLS (Transport Layer Security), 63
 - tombstone process, interval, 694
 - topology
 - of Active Directory, 8
 - described, 27
 - establishing site, 435
 - KCC process and replication, 427
 - physical components of networks, 27–28
 - replication, 32–33
 - replication, KCC analysis, 453
 - Universal Groups, effects, 547
 - trace logs, 714
 - traffic, network
 - and GC server placement, 549–552
 - troubleshooting, 566
 - transactional database system, 690
 - transactions, 692
 - transferring
 - domain-naming FSMO, 512
 - Domain Naming Master role, 306
 - FSMOs, 302, 510

- Infrastructure, RID, PDC Operations
 - Master roles, 308–312
 - master roles, 300–310
 - operations masters roles, 509–522
 - RID, PDC, or infrastructure FSMOs, 513
 - Schema Master role, 304
 - transforms, installer packages, 650
 - transitive site links, 458, 468
 - transitive trusts, 390–391
 - transitivity of trust relationships, 315–316
 - Transmission Control Protocol/Internet Protocol. *See* TCP/IP
 - Transport Layer Security (TLS)
 - described, 63
 - SSL (Security Sockets Layer), 197
 - tree-root trust, 317, 325–326
 - trees described, 246
 - troubleshooting
 - account problems, 168
 - Active Directory availability, 745–758
 - dcpromo failure, 293
 - distributed applications, 682
 - DNS registration, 493
 - GC issues, 551–552
 - Group Policy, 624
 - misconfigurations, 276
 - network traffic issues, 567
 - OM failures, 514
 - PDC Emulator, RID Master, 527
 - raising functional levels, 274
 - replication failure, 460
 - schema issues, 563, 566
 - software deployment, 675
 - trust relationships
 - See also* trusts
 - establishing, 315–317
 - managing with Active Directory Domains and Trusts, 40–43
 - nontransitive (fig.), 316
 - types of, 394–398
 - trusted domain object (TDO), 253
 - Trusted Publishers policy, 620
 - trusts
 - See also* trust relationships
 - direction of, 391
 - Exam Objective Summary, 412
 - Exam Objectives Fast Track, 413
 - explicit and implicit, 393–394
 - external, realm, and shortcut, 414
 - reestablishing after domain rename, 350
 - Self Test, 416–423
 - SID filtering, 414
 - transitive and non-transitive, 316, 391–392
 - two-way transitive, 24
 - types of, 317, 394
 - working with, introduction, 390
- ## U
- unfreezing forest configuration (domain rename), 349
 - Uniform Resource Locators. *See* URLs
 - units of replication, 431–432
 - Universal Groups
 - caching, 248, 550–551
 - and group conversions, 257–258
 - membership information, and GC, 543, 547
 - Universal Naming Convention (UNC), 11
 - universal security groups, 134, 221–224
 - updates
 - configuring required, 668
 - replicating across DCs, 7
 - updating replication topology, 454–455
 - upgrading
 - applications via Group Policy, 667
 - before raising functional level, 272

- DCs (domain controllers), 500–502
 - UPN authentication, 541–542
 - UPN names and domain names, 95
 - UPNs, implementing, 162–163
 - URLs
 - accessing shared resources using, 12
 - Active Directory naming scheme support, 11
 - and UNC paths, 13
 - USB token (fig.), 205
 - user accounts
 - access to information, 3
 - callback options, 127–128
 - creating, 105–109, 554
 - creating with DSADD, 110–113
 - groups, default, 130
 - introduction, 84
 - lockout policies, 190–192
 - logon name and UPN suffix, 11
 - managing multiple, 162–163
 - managing with Active Directory Users and Computers, 113–118
 - security-related tabs, 126–130
 - setting logon hours, 120
 - settings, 118–120
 - troubleshooting, 168
 - use in Active Directory, 99
 - user and group strategies
 - authentication types, 195–203
 - creating password policy for domain users, 182–192
 - creating user authentication strategies, 192–195
 - Exam Objective Summary, 230–231
 - Exam Objectives Fast Track, 232–234
 - implementing smart cards, 206–217
 - introduction, 182
 - planning security group strategy, 217–229
 - planning smart card authentication strategy, 203–206
 - Self Test, 234–241
 - user objects, creating, 109–110
 - user policies, 577–578, 581
 - user principle names (UPNs), 11
 - user rights
 - See also* passwords, privileges
 - vs. permissions, 219
 - users
 - adding to groups, 173
 - creating new, 110–113
 - deploying software via Group Policy to, 652–653
 - educating about security, 202
 - Group Policy configuration, 617
 - number of, and performance, 479
 - restricting logon machines of, 120
 - smart card enrollment, 210–211
 - Users container, 100, 138–140
 - Users group, 138
 - USNs and replication, 147
- ## V
- verbose logging, 677
 - verifying
 - Active Directory database integrity, 752–755
 - forest functional level, 272
 - Group Policy, 608
 - trusts, 398–401
 - use of UPNs after domain rename, 355
 - viewing
 - GPO general properties (fig.), 604
 - GPO security settings, 606
 - gpresult results (fig.), 627
 - logged-on user information, 54
 - RSoP planning results (fig.), 588

SIDS (domain security IDs) , 90–92, 173
 virtual private network (VPN) clients, 182
 volume shadow copying, 698, 721
 VPN remote access, 127

W

W32Time service, 254, 508
 WANs. *See* wide area networks (WANs)
 warning events, 713
 Web sites, federated forest information, 251
 Web, surfing via URL names, 13
 well-known SIDs (table), 88–90
 Whoami tool, 54–55, 90, 173
 wide area networks (WANs)
 and domain changes, 480–481
 network traffic considerations, 549–552
 Windows 2000
 Active Directory and, 3
 domain functional levels, 261–265
 domain rename, Registry changes, 357
 forest functional level, default, 267
 KCC, 250
 mixed and native modes, 500–502
 Windows 2000 mixed, 65
 Windows 2000 Professional, 151
 Windows Backup, System State restoration, 298
 Windows Exchange, 549
 Windows Installer packages, 643, 649–650, 654
 Windows Management Instrumentation (WMI) filters, 586, 607
 Windows NT
 administrative tools, 366
 creating computer accounts, 153
 external, one-way trusts, 396–397
 NTDS. *See* NTDS
 upgrading to Windows Server 2003, 274
 Windows Server 2003

 administrative tools, 366
 Backup utility, 721
 domain functional levels, 261–265
 forest functional level, 269
 forest functional levels, 265–268
 installing Active Directory on, 6
 new features (Active Directory), 65–71
 promoting to DC, 29
 restore options, 734
 Support Tools, 461
 version support for Active Directory, 9
 W32Time service, 254
 Windows Server 2003 interim, 65, 261–265, 268
 WinInstall, 654–656
 wizards
 Active Directory Installation, 15–18, 275, 484–492
 Backup or Restore, 721–732
 Configure Your Server, 478
 Configure Your Server Wizard, 484
 Delegation of Control, 106, 410
 Group Policy, 596
 New Trust, 317, 395
 Restore, 298
 Resultant Set of Policy, 584–588
 RSoP, 631

X

X.500
 and Directory Access Protocol (DAP), 10
 standard described, 559
 X.509 certificates, 63

Z

.zap setup files, 656–659, 681
 zones, DNS, 367–368



MCSE WORLD

If you've read the book, and you're looking for more of the best
MCSA and MCSE certification tips and tricks, go to

<http://www.mcseworld.com/>

Available Now:

- ▲ Discussion Forums
- ▲ InfoCenter Library
- ▲ Arcade
- ▲ Newsletters
- ▲ Questions of the Day
- ▲ Links
- ▲ eShop
- ▲ Polls

Coming Soon:

- ▲ Chat Rooms
- ▲ Practice Exams
- ▲ Study Guides

**Find more great
MCSA and MCSE
certification titles from
Syngress Publishing at
MCSE World!**


<http://www.mcseworld.com/>

MCSE World is brought you to by Area 51 Partners, Inc. and RS Networks
<http://www.area51partners.com/> <http://www.rsnetworks.net/>

MCSE WINDOWS 2003 FOUR CORE EXAM STUDY GUIDE & DVD TRAINING



Syngress' 100% Certified Study Guide & DVD Training System are a fully integrated learning system (Study Guide/Online Exams/DVD) guaranteed to deliver 100% coverage of Microsoft's learning objectives for MCSE Windows 2003 Server certification.



Exam 70-290: Managing and Maintaining a Microsoft Windows Server 2003 Environment


ISBN: 1-931866-60-7

Price: \$59.95 US

Exam 70-291: Implementing, Managing, and Maintaining a Microsoft Windows Server 2003

ISBN: 1-931836-92-2

Price: \$59.95 US



Exam 70-293: Planning and Maintaining a Microsoft Windows Server 2003 Network Infrastructure


ISBN: 1-931836-93-0

Price: \$59.95 US

Exam 70-294: Planning, Implementing, and Maintaining a Microsoft Windows Server 2003 Active Directory Infrastructure

ISBN: 1-931836-94-9

Price: \$59.95 US



MCSE Windows Server 2003 Boxed Set Study Guide & DVD Training System

ISBN: 1-931836-96-5

Price: \$199.95 US

MCSA Windows Server 2003 Boxed Set Study Guide & DVD Training System

ISBN: 1-932266-44-5

Price: \$99.95 US

MCSE 2003 Certification Upgrade KIT (Exams 70-292 and 70-296)

ISBN: 1-932266-61-5

Price: \$99.95 US

Career Advancement Through Skill Enhancement®

www.syngress.com/certification

SYNGRESS®