

SYNGRESS
SHINDERBOOKS

**INCLUDES FREE
WEB-BASED TESTING!**

COVERS ALL
**100%
CERTIFIED**

EXAM OBJECTIVES

MCSSE

**Exam 70-293: Planning and Maintaining a
Windows Server 2003 Network 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

Planning and Maintaining a Windows Server
2003 Network Infrastructure: Exam 70-293

STUDY GUIDE & DVD TRAINING SYSTEM

Martin Grasdal

Laura E. Hunter

Michael Cross

Laura Hunter 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 and Maintaining a Windows Server 2003 Network Infrastructure: Exam 70-293
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: 1-931836-93-0

Technical Editors: Debra Littlejohn Shinder
Dr. Thomas W. Shinder
Technical Reviewer: Laura E. Hunter
Acquisitions Editor: Jonathan Babcock
DVD Production: Michael Donovan

Cover Designer: Michael Kavish
Page Layout and Art by: John Vickers
Copy Editor: Michelle Melani and Marilyn Smith
Indexer: Nara Wood
DVD Presenter: Laura 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 book network in a highly professional manner and under severe time constraints, but still keeps a good sense of humor.



Contributors

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 sit. 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).

Van Varnell (Master CNE, MCSE, MCDBA) is a Senior Network Analyst for Appleton, Inc. His areas of expertise are development and maintenance of high-availability systems, storage area networks and storage platforms, performance monitoring systems, and data center operations. Van has held high-level positions in the industry over the 15 years of his career including that of Windows Systems Architect for Motorola and Senior Consultant for Integrated Information Systems. Van holds a bachelor's degree in Computer Information Systems and currently resides in Wisconsin with his wife Lisa and five children (Brennan, Kyle, Katelyn, Kelsey, and Kevin). He wishes to thank his wife and kids for *being* his wife and kids, and Jon Babcock of Syngress for his patience and assistance.

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, and network administration. 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 has been 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.

Paul M. Summitt (MCSE, CCNA, MCP+I, MCP) has a Masters degree in Mass Communication. Currently the IT Director for the Missouri County Employees' Retirement Fund, Paul has served as network, exchange, and database administrator as well as Web and application developer. Paul has written previously on virtual reality and Web development and has served as technical editor for several books on Microsoft technologies. Paul lives in Columbia, Missouri with his life and writing partner Mary. To the Syngress editorial staff, my thanks for letting me be a part of this project. To my kids, adulthood is just the beginning of all the fun you can have.

Rob Amini (MCSE, MCDBA, MCT) is currently a systems manager for Marriott International in Salt Lake City, Utah. He has a Bachelor's degree in computer science and has been breaking and fixing machines since the Atari 800 was considered state of the art. In 1993 he began his professional career by fixing IBM mainframes and various unix-flavored boxes. After a long stint as a technician and systems admin, he gained fabled notoriety as a

pun-wielding Microsoft trainer. Rob has continued as an instructor for more than three years and although teaching is his first love, he tends to enjoy technical writing more than a well-adjusted person should. When actually not working with and programming a variety of electronic gizmos, Rob enjoys spending every minute he can with his beautiful wife Amy and the rest of his supportive family.

Dan Douglass (MCSE+I, MCDBA, MCSD, MCT) is a software developer and trainer with a cutting edge medical software company in Dallas, Texas. He currently provides software development skills, internal training and integration solutions, as well as peer guidance for technical skills development. His specialties include enterprise application integration and design, HL7, XML, XSL, Visual Basic, database design and administration, Back Office and .NET Server platforms, network design, Microsoft operating systems, and FreeBSD. Dan is a former US Navy Submariner and lives in Plano, TX with his very supportive and understanding wife, Tavish.

Jada Brock-Soldavini is a MCSE and holds a degree in Computer Information Systems. She has worked in the Information Technology Industry for over 7 years. She is working on her Cisco certification track currently and has contributed to over a dozen books and testing software for the Microsoft exam curriculum. She works for the State of Georgia as a Network Services Administrator. When she is not working on her technical skills she enjoys playing the violin. Jada is married and lives in the suburbs of Atlanta with her husband and children.

Michael Moncur is an MCSE and CNE. He is the author of several best-selling books about networking and the Internet, including *MCSE In a Nutshell: The Windows 2000 Exams* (O'Reilly and Associates). Michael lives in Salt Lake City with his wife, Laura.



Technical Reviewer, DVD Presenter, and Contributor

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 websites.

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.



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.

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.

MCSE 70-293 Exam Objectives Map and Table of Contents



All of Microsoft's published objectives for the MCSE 70-293 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 70-293 Exam objectives.

Exam Objective Map

Objective Number	Objective	Chapter Number
1	Planning and Implementing Server Roles and Server Security	2
1.1	Configure security for servers that are assigned specific roles.	2
1.2	Plan a secure baseline installation.	2
1.2.1	Plan a strategy to enforce system default security settings on new systems.	2
1.2.2	Identify client operating system default security settings.	2
1.2.3	Identify all server operating system default security settings.	2
1.3	Plan security for servers that are assigned specific roles. Roles might include domain controllers, Web servers, database servers, and mail servers.	2
1.3.1	Deploy the security configuration for servers that are assigned specific roles.	2

Objective Number	Objective	Chapter Number
1.3.2	Create custom security templates based on server roles.	2
1.4	Evaluate and select the operating system to install on computers in an enterprise.	2
1.4.1	Identify the minimum configuration to satisfy security requirements.	2
2	Planning, Implementing, and Maintaining a Network Infrastructure	3, 4, 5
2.1	Plan a TCP/IP network infrastructure strategy.	3
2.1.1	Analyze IP addressing requirements.	3
2.1.2	Plan an IP routing solution.	3, 4
2.1.3	Create an IP subnet scheme.	3
2.2	Plan and modify a network topology.	3
2.2.1	Plan the physical placement of network resources.	3
2.2.2	Identify network protocols to be used.	3
2.3	Plan an Internet connectivity strategy.	5
2.4	Plan network traffic monitoring. Tools might include Network Monitor and System Monitor.	3
2.5	Troubleshoot connectivity to the Internet.	5
2.5.1	Diagnose and resolve issues related to Network Address Translation (NAT).	5
2.5.2	Diagnose and resolve issues related to name resolution cache information.	6
2.5.3	Diagnose and resolve issues related to client configuration.	4
2.6	Troubleshoot TCP/IP addressing.	3
2.6.1	Diagnose and resolve issues related to client computer configuration.	3
2.6.2	Diagnose and resolve issues related to DHCP server address assignment.	3
2.7	Plan a host name resolution strategy.	6
2.7.1	Plan a DNS namespace design.	6
2.7.2	Plan zone replication requirements.	6
2.7.3	Plan a forwarding configuration.	6

Objective Number	Objective	Chapter Number
2.7.4	Plan for DNS security.	6
2.7.5	Examine the interoperability of DNS with third-party DNS solutions.	6
2.8	Plan a NetBIOS name resolution strategy.	6
2.8.1	Plan a WINS replication strategy.	6
2.8.2	Plan NetBIOS name resolution by using the Lmhosts file.	6
2.9	Troubleshoot host name resolution.	6
2.9.1	Diagnose and resolve issues related to DNS services.	6
2.9.2	Diagnose and resolve issues related to client computer configuration.	6
3	Planning, Implementing, and Maintaining Routing and Remote Access	4, 7
3.1	Plan a routing strategy.	4
3.1.1	Identify routing protocols to use in a specified environment.	4
3.1.2	Plan routing for IP multicast traffic.	4
3.2	Plan security for remote access users.	7
3.2.1	Plan remote access policies.	7
3.2.2	Analyze protocol security requirements.	7
3.2.3	Plan authentication methods for remote access clients.	7
3.3	Implement secure access between private networks.	7
3.3.1	Create and implement an IPSec policy.	10
3.4	Troubleshoot TCP/IP routing. Tools might include the route, tracert, ping, pathping, and netsh commands and Network Monitor.	4
4	Planning, Implementing, and Maintaining Server Availability	8
4.1	Plan services for high availability.	8
4.1.1	Plan a high availability solution that uses clustering services.	9

Objective Number	Objective	Chapter Number
4.1.2	Plan a high availability solution that uses Network Load Balancing.	9
4.2	Identify system bottlenecks, including memory, processor, disk, and network related bottlenecks.	8
4.2.1	Identify system bottlenecks by using System Monitor.	8
4.3	Implement a cluster server.	9
4.3.1	Recover from cluster node failure.	9
4.4	Manage Network Load Balancing. Tools might include the Network Load Balancing Monitor Microsoft Management Console (MMC) snap-in and the WLBS cluster control utility.	9
4.5	Plan a backup and recovery strategy.	8
4.5.1	Identify appropriate backup types. Methods include full, incremental, and differential.	8
4.5.2	Plan a backup strategy that uses volume shadow copy.	8
4.5.3	Plan system recovery that uses Automated System Recovery (ASR).	8
5	Planning and Maintaining Network Security	10, 11
5.1	Configure network protocol security.	10
5.1.1	Configure protocol security in a heterogeneous client computer environment.	10
5.1.2	Configure protocol security by using IPSec policies.	10
5.2	Configure security for data transmission.	10
5.2.1	Configure IPSec policy settings.	10
5.3	Plan for network protocol security.	10
5.3.1	Specify the required ports and protocols for specified services.	4
5.3.2	Plan an IPSec policy for secure network communications.	10
5.4	Plan secure network administration methods.	11
5.4.1	Create a plan to offer Remote Assistance to client computers.	7

Objective Number	Objective	Chapter Number
5.4.2	Plan for remote administration by using Terminal Services.	7
5.5	Plan security for wireless networks.	11
5.6	Plan security for data transmission.	10
5.6.1	Secure data transmission between client computers to meet security requirements.	10
5.6.2	Secure data transmission by using IPSec.	10
5.7	Troubleshoot security for data transmission. Tools might include the IP Security Monitor MMC snap-in and the Resultant Set of Policy (RSoP) MMC snap-in.	10
6	Planning, Implementing, and Maintaining Security Infrastructure.	11, 12
6.1	Configure Active Directory directory service for certificate publication.	12
6.2	Plan a public key infrastructure (PKI) that uses Certificate Services.	12
6.2.1	Identify the appropriate type of certificate authority to support certificate issuance requirements.	12
6.2.2	Plan the enrollment and distribution of certificates.	12
6.2.3	Plan for the use of smart cards for authentication.	12
6.3	Plan a framework for planning and implementing security.	11
6.3.1	Plan for security monitoring.	11
6.3.2	Plan a change and configuration management framework for security.	11
6.4	Plan a security update infrastructure. Tools might include Microsoft Baseline Security Analyzer and Microsoft Software Update Services.	11

Contents

Foreword	xxxvii
Chapter 1 Using Windows Server 2003 Planning Tools and Documentation	1
Introduction	2
Overview of Network Infrastructure Planning	2
Planning Strategies	3
Using Planning Tools	3
Fundamentals of Network Design	9
Analyzing Organizational Needs	11
Information Flow Factors	11
Management Model and Organizational Structure	12
Centralization versus Decentralization	13
Management Priorities	14
Availability/Fault Tolerance	15
Security	15
Scalability	16
Performance	16
Cost	16
User Priorities	17
Electronic Communications	17
Scheduling/Task Management	18
Project Collaboration	19
Data Storage and Retrieval	21
Internet Research	23
Application Services	23
Print Services	24
Graphics/Audio/Video Services	26
Reviewing Legal and Regulatory Considerations	26
Calculating TCO	27

Planning for Growth	28
Developing a Test Network Environment	29
Planning the Test Network	30
Implementing the Test Network	34
Documenting the Planning and Network Design Process	36
Importance of Documentation	37
Creating the Planning and Design Document	37
Summary of Exam Objectives	39
Exam Objectives Fast Track	40
Exam Objectives Frequently Asked Questions	41
Self Test	43
Self Test Quick Answer Key	51
Chapter 2 Planning Server Roles and Server Security	53
Introduction	54
1.1.1 Understanding Server Roles	54
Domain Controllers (Authentication Servers)	58
Active Directory	58
Operations Master Roles	59
File and Print Servers	62
Print Servers	62
File Servers	62
DHCP, DNS, and WINS Servers	63
DHCP Servers	63
DNS Servers	64
WINS Servers	65
Web Servers	65
Web Server Protocols	66
Web Server Configuration	67
Database Servers	68
Mail Servers	68
Certificate Authorities	69
PKI	69
Certificates	70
Certificate Services	71
Application Servers and Terminal Servers	75
Application Servers	75

	Terminal Servers	78
1.1	Planning a Server Security Strategy	78
1.4	Choosing the Operating System	79
	Security Features	81
	Functional Levels	83
1.4.1	Identifying Minimum Security Requirements for Your Organization	91
	Identifying Configurations to Satisfy Security Requirements	93
1/1.2	Planning Baseline Security	94
	Security Templates and Tools	94
	Predefined Templates	95
	Security Configuration and Analysis	98
	Group Policy Object Editor	99
	Secedit	100
	Planning Secure Baseline Installation Parameters	103
	Using Security Configuration and Analysis to Analyze a Computer	103
1.2.1/1.2.2	Enforcing Default Security Settings on New Computers	109
1.2.3		
	Using Security Configuration and Analysis to Apply Templates a Local Computer	109
	Using Group Policy Object Editor to Apply Templates	109
1	Customizing Server Security	113
1.3/1.3.1	Securing Servers According to Server Roles	113
	Security Issues Related to All Server Roles	113
	Securing Domain Controllers	121
	Securing File and Print Servers	122
	Securing DHCP, DNS, and WINS Servers	125
	Securing Web Servers	126
	Securing Database Servers	127
	Securing Mail Servers	128
	Securing CAs	129
	Securing Application and Terminal Servers	130
1.3.2	Creating Custom Security Templates	131
	Deploying Security Configurations	134

	Summary of Exam Objectives	137
	Exam Objectives Fast Track	137
	Exam Objectives Frequently Asked Questions	139
	Self Test	140
	Self Test Quick Answer Key	146
	Chapter 3 Planning, Implementing, and Maintaining the TCP/IP Infrastructure	147
2/2.1/2.1.2	Introduction	148
	Understanding Windows 2003 Server Network Protocols	148
2.2.2	Identifying Protocols to Be Used	149
	Advantages of the TCP/IP Protocol Suite	151
	The Multiprotocol Network Environment	153
	Reviewing TCP/IP Basics	160
	What's New in TCP/IP for Windows Server 2003	164
	IGMPv3	165
	IPv6	165
	Alternate Configuration	166
	Automatic Determination of Interface Metric	167
2/2.1/2.1.2	Planning an IP Addressing Strategy	171
2.1.1	Analyzing Addressing Requirements	171
2.1.3	Creating a Subnetting Scheme	173
	Classful Addressing	173
	Understanding ANDing and Binary Numbering	175
	Subnetting Networks	177
	Classless Inter-Domain Routing (CIDR)	180
2.6	Troubleshooting IP Addressing	181
2.6.1	Client Configuration Issues	181
2.6.2	DHCP Issues	182
	Transitioning to IPv6	183
	IPv6 Utilities	184
	6to4 Tunneling	192
	IPv6 Helper Service	192
	The 6bone	193
	Teredo (IPv6 with NAT)	193
2/2.1	Planning the Network Topology	193
2.1.2/2.2		
	Analyzing Hardware Requirements	193
2.2.1	Planning the Placement of Physical Resources	194

2/2.1/2.1.1/	Planning Network Traffic Management	194
2.4	Monitoring Network Traffic and Network Devices	195
	Using Network Monitor	195
	Using System Monitor	196
	Determining Bandwidth Requirements	198
	Optimizing Network Performance	198
	Summary of Exam Objectives	200
	Exam Objectives Fast Track	200
	Exam Objectives Frequently Asked Questions	202
	Self Test	204
	Self Test Quick Answer Key	209
	Chapter 4 Planning, Implementing, and Maintaining a Routing Strategy	211
	Introduction	212
2/2.1.2/3	Understanding IP Routing	212
	Reviewing Routing Basics	213
	Routing Tables	216
	Static versus Dynamic Routing	220
	Gateways	222
3.1.2	Planning a Routing Strategy for IP Multicast Traffic	223
	Routing Protocols	225
	Using Netsh Commands	233
	Evaluating Routing Options	236
	Selecting Connectivity Devices	236
	Switches	242
	Routers	245
	Windows Server 2003 As a Router	245
2/2.1.2/3/	Security Considerations for Routing	257
3.1/5.3.1	Analyzing Requirements for Routing Components	259
	Simplifying Network Topology to Provide Fewer Attack Points	259
	Minimizing the Number of Network Interfaces and Routes	260
	Minimizing the Number of Routing Protocols	260
	Router-to-Router VPNs	263
	Packet Filtering and Firewalls	268
	Logging Level	269

2/2.1.2/3	Troubleshooting IP Routing	270
3.4	Identifying Troubleshooting Tools	271
	Common Routing Problems	274
	Interface Configuration Problems	274
	RRAS Configuration Problems	274
	Routing Protocol Problems	275
2.5.3	TCP/IP Configuration Problems	276
	Routing Table Configuration Problems	276
	Summary of Exam Objectives	277
	Exam Objectives Fast Track	277
	Exam Objectives Frequently Asked Questions	279
	Self Test	280
	Self Test Quick Answer Key	285
	Chapter 5 Planning, Implementing, and Maintaining an Internet Connectivity Strategy	287
	Introduction	288
2/2.3/2.5	Connecting the LAN to the Internet	289
	Routed Connections	289
	Advantages of Routed Connections	289
	Hardware and Software Routers	289
	IP Addressing for Routed Connections	290
	Translated Connections	290
2.5	Network Address Translation (NAT)	291
	Internet Connection Sharing (ICS)	297
2/2.3	Implementing Virtual Private Networks (VPNs)	300
	Internet-based VPNs	301
	How Internet-based VPNs Work	301
	Configuring Internet-based VPNs	302
	Router-to-Router VPNs	303
	On Demand/Demand-Dial Connections	304
	One-Way versus Two-Way Initiation	306
	Persistent Connections	306
	Remote-Access Policies	306
	VPN Protocols	306
	PPTP	307
	L2TP	307

	VPN Security	307
	MPPE	307
	IPSec	307
2/2.3	Using Internet Authentication Service (IAS)	308
	Advantages of IAS	308
	Centralized User Authentication and Authorization	308
	Centralized Auditing and Accounting	309
	RRAS Integration	309
	Control via Remote-Access Policies	309
	Extensibility and Scalability	309
	IAS Management	309
	Activating IAS Authentication	310
	Using the IAS MMC Snap-in	312
	IAS Monitoring	313
	IAS SDK	313
	Authentication Methods	314
	PPP-based Protocols	314
	EAP	314
	Authorization Methods	317
	Dial Number Identification Service (DNIS)	317
	Automatic Number Identification (ANI) and Calling Line Identification (CLI)	317
	Guest Authorization	317
	Access Server Support	318
	Outsourced Dialing	318
2/2.3	Using Connection Manager	318
	Using CMAK	319
	Installing and Running CMAK	319
	Service Profiles	323
	Custom Actions	323
	Custom Help	324
	VPN Support	324
	Connection Manager Security Issues	324
	Preventing Editing of Service Profile Files	324
	Client Operating System, File System, and Configuration	324
	Preventing Users from Saving Passwords	325

	Secure Distribution of Service Profiles	325
	Summary of Exam Objectives	326
	Exam Objectives Fast Track	326
	Exam Objectives Frequently Asked Questions	328
	Self Test	330
	Self Test Quick Answer Key	334
	Chapter 6 Planning, Implementing, and Maintaining a Name Resolution Strategy	335
	Introduction	336
2.7	Planning for Host Name Resolution	337
	Understanding Host Naming	337
	NetBIOS over TCP/IP	338
	Host Names	338
	Understanding the Hosts File	339
	Understanding DNS	341
2.7.1	Designing a DNS Namespace	357
	Choosing the Parent Domain Name	358
	Host Naming Conventions and Limitations	359
	DNS and Active Directory (AD)	361
	Supporting Multiple Namespaces	363
	Planning DNS Server Deployment	369
	Planning the Number of DNS Servers	369
	Planning for DNS Server Capacity	371
	Planning DNS Server Placement	372
	Planning DNS Server Roles	373
2.7.2	Planning for Zone Replication	377
	Active Directory-integrated Zone Replication Scope	379
	Security for Zone Replication	382
	General Guidelines for Planning for Zone Replication	382
2.7.3	Planning for Forwarding	383
	Conditional Forwarding	384
	General Guidelines for Using Forwarders	386
	DNS/DHCP Interaction	387
	Security Considerations for DDNS and DHCP	389
	Aging and Scavenging of DNS Records	391
2.7.5	Windows Server 2003 DNS Interoperability	392

	BIND and Other DNS Server Implementations	393
	Zone Transfers with BIND	395
	Supporting AD with BIND	397
	Split DNS Configuration	398
	Interoperability with WINS	399
2.7.4	DNS Security Issues	404
	Common DNS Threats	406
	Securing DNS Deployment	407
	DNS Security Levels	408
	General DNS Security Guidelines	410
	Monitoring DNS Servers	412
	Testing DNS Server Configuration with the DNS	
	Console Monitoring Tab	413
	Debug Logging	414
	Event Logging	415
	Monitoring DNS Server Using the Performance Console ..	415
	Command-line Tools for Maintaining and	
	Monitoring DNS Servers	416
2.8	Planning for NetBIOS Name Resolution	417
	Understanding NETBIOS Naming	418
	NetBIOS Name Resolution Process	418
2.8.2	Understanding the LMHOSTS File	420
	Understanding WINS	421
	What's New for WINS in Windows Server 2003	424
	Planning WINS Server Deployment	424
	Server Number and Placement	424
2.8.1	Planning for WINS Replication	427
	Replication Partnership Configuration	428
	Replication Models	434
	WINS Issues	437
	Static WINS Entries	438
	Multihomed WINS Servers	439
	Client Configuration	440
	Preventing Split WINS Registrations	444
	Performance Issues	444
	Security Issues	449

	Planning for WINS Database Backup and Restoration	451
2.5.2	Troubleshooting Name Resolution Issues	452
2.9	Troubleshooting Host Name Resolution	453
	Issues Related to Client Computer Configuration	454
2.9.1	Issues Related to DNS Services	455
	Troubleshooting NetBIOS Name Resolution	457
	Issues Related to Client Computer Configuration	457
	Issues Related to WINS Servers	458
	Summary of Exam Objectives	461
	Exam Objectives Fast Track	469
	Exam Objectives Frequently Asked Questions	472
	Self Test	474
	Self Test Quick Answer Key	483
	Chapter 7 Planning, Implementing, and Maintaining a Remote Access Strategy	485
	Introduction	486
3	Planning the Remote Access Strategy	486
	Analyzing Organizational Needs	487
	Analyzing User Needs	487
	Selecting Remote Access Types To Allow	487
	Dial-In	488
	VPN	488
	Wireless Remote Access	489
3	Addressing Dial-In Access Design Considerations	489
	Allocating IP Addresses	490
	Static Address Pools	490
	Using DHCP for Addressing	490
	Using APIPA	491
	Determining Incoming Port Needs	491
	Multilink and BAP	491
	Selecting an Administrative Model	492
	Access by User	493
	Access by Policy	494
3/3.3	Addressing VPN Design Considerations	495
	Selecting VPN Protocols	496
	Client Support	496

	Data Integrity and Sender Authentication	496
	PKI Requirements	497
	Installing Machine Certificates	497
	Configuring Firewall Filters	499
	Creating Access Policies	500
3	Addressing Wireless Remote Access Design Considerations	500
	The 802.11 Wireless Standards	501
	Using IAS for Wireless Connections	501
	Configuring Remote Access Policies for Wireless Connections	502
	Multiple Wireless Access Points	503
	Placing CA on VLAN for New Wireless Clients	503
	Configuring WAPs as RADIUS Clients	503
	Wireless Encryption and Security	504
	WEP (Wired Equivalent Privacy)	504
	802.1X	504
	WPA	505
3.2.2/3/3.2/	Planning Remote Access Security	505
3.2.1		
	Domain Functional Level	505
	Determining the Function Level	506
	Raising the Domain Functional Level	507
3.2.3	Selecting Authentication Methods	508
	Disallowing Password-Based Connections (PAP, SPAP, CHAP, MS-CHAP v1)	509
	Using MS-CHAP v2	511
	Using EAP	511
	Using RADIUS/IAS vs. Windows Authentication	512
	Selecting the Data Encryption Level	512
	Using Callback Security	513
	Managed Connections	513
	Mandating Operating System/File System	514
	Using Smart Cards for Remote Access	514
3	Creating Remote Access Policies	515
	Policies and Profiles	515
	Authorizing Remote Access	516
	Authorizing Access By User	516

	Authorizing Access By Group	518
	Restricting Remote Access	520
	Restricting by User/Group Membership	521
	Restricting by Type of Connection	521
	Restricting by Time	523
	Restricting by Client Configuration	524
	Restricting Authentication Methods	524
	Restricting by Phone Numbers of MAC Addresses	525
	Controlling Remote Connections	525
	Controlling Idle Timeout	525
	Controlling Maximum Session Time	525
	Controlling Encryption Strength	527
	Controlling IP packet Filters.....	528
	Controlling IP addresses for PPP Connections.....	528
3/5.4	Creating a Plan to Offer Remote Assistance to Client Computers ..	529
	How Remote Assistance Works	529
	Using Remote Assistance	530
	Configuring Remote Assistance for Use	530
	Asking for Assistance	532
	Completing the Connection	537
	Managing Open Invitations	540
	Offering Remote Assistance to your Clients	542
	Remote Assistance Security Issues	543
3/5.4.2	Planning for Remote Administration by Using Terminal Services ..	545
	Using Remote Desktop for Administration	545
	Configuring RDA	545
	Setting Up Authentication	546
	Advantages of RDA Over Other Remote	
	Administration Methods	546
	Remote Desktop Security Issues	547
	Summary of Exam Objectives	549
	Exam Objectives Fast Track	550
	Exam Objectives Frequently Asked Questions	552
	Self Test	553
	Self Test Quick Answer Key	558

	Chapter 8 Planning, Implementing, and Maintaining a High-Availability Strategy	559
	Introduction	560
4/4.1/4.2	Understanding Performance Bottlenecks	560
	Identifying System Bottlenecks	561
	Memory	561
	Processor	563
	Disk	564
	Network Components	568
4.2.1	Using the System Monitor Tool to Monitor Servers	570
	Using Event Viewer to Monitor Servers	584
	Using Service Logs to Monitor Servers	593
4/4.1/4.5	Planning a Backup and Recovery Strategy	593
4.5.1	Understanding Windows Backup	594
	Types of Backups	596
	Determining What to Back Up	600
	Using Backup Tools	602
	Using the Windows Backup Utility	602
	Using the Command-Line Tools	604
	Selecting Backup Media	604
	Scheduling Backups	605
	Restoring from Backup	606
4.5.3/4/4.1	Planning System Recovery with ASR	612
	What Is ASR?	613
	How ASR Works	613
	Alternatives to ASR	614
	Safe Mode Boot	614
	Last Known Good Boot Mode	614
	ASR As a Last Resort	615
	Using the ASR Wizard	615
	Performing an ASR Restore	617
	Planning for Fault Tolerance	618
	Network Fault-Tolerance Solutions	619
	Internet Fault-Tolerance Solutions	619
	Disk Fault-Tolerance Solutions	620
	RAID	620
	Hot Spare Drives	624
	Server Fault-Tolerance Solutions	624

	Summary of Exam Objectives	626
	Exam Objectives Fast Track	627
	Exam Objectives Frequently Asked Questions	630
	Self Test	631
	Self Test Quick Answer Key	638
	Chapter 9 Implementing Windows Cluster Services and Network Load Balancing	639
	Introduction	640
4.1.1	Making Server Clustering Part of Your High-Availability Plan	641
	Terminology and Concepts	641
	Cluster Nodes	641
	Cluster Groups	642
	Failover and Failback	643
	Cluster Services and Name Resolution	643
	How Clustering Works	643
	Cluster Models	644
	Single Node	644
	Single Quorum Device	645
	Majority Node Set	646
4.3	Server Cluster Deployment Options	647
	N-Node Failover Pairs	648
	Hot-Standby Server/N+1	649
	Failover Ring	651
	Random	652
	Server Cluster Administration	653
	Using the Cluster Administrator Tool	653
	Using Command-Line Tools	654
4.3.2	Recovering from Cluster Node Failure	657
	Server Clustering Best Practices	657
	Hardware Issues	658
4.3	Cluster Network Configuration	662
	Security	667
4.1.2	Making Network Load Balancing Part of Your High-Availability Plan	678
	Terminology and Concepts	678
	Hosts/Default Host	678
	Load Weight	679

	Traffic Distribution	679
	Convergence and Heartbeats	680
	How NLB Works	681
	Relationship of NLB to Clustering	681
4.4	Managing NLB Clusters	682
	Using the NLB Manager Tool	682
	Remote Management	683
	Command-Line Tools	684
	NLB Error Detection and Handling	687
	Summary of Exam Objectives	699
	Exam Objectives Fast Track	699
	Exam Objectives Frequently Asked Questions	701
	Self Test	702
	Self Test Quick Answer Key	708
	Chapter 10 Planning, Implementing, and Maintaining Internet Protocol Security	709
	Introduction	710
3.3.1/5/5.3	Understanding IP Security (IPSec)	710
5.6/5.6.1/5.6.2	Terminology and Concepts	712
	How IPSec Works	713
	Securing Data in Transit	714
	Purposes of Encryption	715
	IPSec Modes	717
	Tunnel Mode	717
	Transport Mode	718
	IPSec Protocols	718
	Primary IPSec Protocols	719
	Additional Protocols	722
	IPSec Components	724
	IPSec Policy Agent	724
	IPSec Driver	725
	IPSec and IPv6	726
3.3.1/5/5.3	Deploying IPSec	726
5.6/5.6.1/5.6.2/5.1	Determining Organizational Needs	727

	Security Levels	727
3.3.1/5/5.6.2	Managing IPsec	728
	Using the IP Security Policy Management MMC Snap-in	728
	Using the netsh Command-line Utility	731
	Default IPsec Policies	732
	Client (Respond Only)	732
	Server (Request Security)	733
	Secure Server (Require Security)	733
	Custom Policies	734
	Using the IP Security Policy Wizard	735
	Defining Key Exchange Settings	743
	Managing Filter Lists and Filter Actions	744
	Assigning and Applying Policies in Group Policy	746
	Active Directory Based IPsec Policies	747
	IPsec Monitoring	749
	Using the netsh Utility for Monitoring	749
	Using the IP Security Monitor MMC Snap-in	750
5.7	Troubleshooting IPsec	751
	Using netdiag for Troubleshooting Windows Server 2003 IPsec	751
	Viewing Policy Assignment Information	752
	Viewing IPsec Statistics	753
	Using Packet Event Logging to Troubleshoot IPsec	755
	Using IKE Detailed Tracing to Troubleshoot IPsec	757
	Using the Network Monitor to Troubleshoot IPsec	759
	Disabling TCP/IP and IPsec Hardware Acceleration to Solve IPsec Problems	760
3.3.1/5/ 5.2/5.7	Addressing IPsec Security Considerations	761
	Strong Encryption Algorithm (3DES)	761
	Firewall Packet Filtering	762
	Diffie-Hellman Groups	762
	Pre-shared Keys	763
	Advantages and Disadvantages of Pre-shared Keys	764
	Considerations when Choosing a Pre-shared Key	764
	Soft Associations	764
3.3.1/5/5.7	Using RSoP for IPsec Planning	765
	Using the RSoP Wizard	766

	Security and RSoP	766
	Selecting the RSoP Mode for IPSec-related Queries	766
	Logging Mode Queries	767
	Planning Mode Queries	768
	Summary	769
	Exam Objectives Fast Track	770
	Exam Objectives Frequently Asked Questions	772
	Self Test	772
	Self Test Quick Answer Key	779
	Chapter 11 Planning, Implementing, and Maintaining a Security Framework	781
	Introduction	782
5/5.4/6/6.3	Planning and Implementing Active Directory Security	782
	Understanding Permission Types	787
	Active Directory Permissions	787
	NTFS Permissions	788
	Share Permissions	789
	Physically Securing Domain Controllers	790
	Securing the Schema	790
	Managing Cross-domain and Cross-forest Security Relationships	791
	Cross-domain Relationships	791
	Cross-forest Relationships	793
	Account Security	795
5/5.4/5.5/ 6/6.3	Planning and Implementing Wireless Security	801
	Understanding Wireless Networking	803
	Wireless Network Types	803
	EAP Authentication	804
	How Wireless Networking Works	806
	Authentication for Wireless Networks	806
	Authentication Protocols	810
	Wireless Security Issues	812
	Default Settings	813
	WEP Weaknesses	815
	Making Wireless More Secure	815

5/6/6.3/6.3.1	Monitoring and Optimizing Security	817
	Wireless Monitor	817
	Object-based Access Control	818
	Auditing	818
	Auditing Registry Keys	821
	Auditing Files or Folders	822
	Viewing the Results of Auditing	823
	Security Log Settings	823
	Security Policies	823
	Password Policies	824
	Kerberos Policies	825
	Account Lockout Policies	826
	User Rights	826
	Security Templates	827
5/6/6.3/6.3.1	Planning a Change and Configuration Management Framework ..	830
5.4		
5/6/6.3/6.3.1	Planning a Security Update Infrastructure	830
5.4		
	Understanding the Importance of Regular	
	Security Updates	831
	Using Microsoft Baseline Security Analyzer (MBSA)	831
	Installing the Microsoft Baseline Security Analyzer	832
	Using Microsoft Software Update Services (SUS)	837
	Summary of Exam Objectives	848
	Exam Objectives Fast Track	851
	Exam Objectives Frequently Asked Questions	852
	Self Test	853
	Self Test Quick Answer Key	859
Chapter 12 Planning, Implementing, and Maintaining		
a Public Key Infrastructure		861
	Introduction	862
6/6.2	Planning a Windows Server 2003 Certificate-Based PKI	862
	Understanding Public Key Infrastructure	863
	Public Key Cryptography	864
	The Function of the PKI	867
	Components of the PKI	867
	Understanding Digital Certificates	868
	User Certificates	870

	Machine Certificates	870
	Application Certificates	870
6.2.1	Understanding Certification Authorities	870
6.2.1	CA Hierarchy	871
	How Microsoft Certificate Services Works	872
6/6.1/6.2.1	Implementing Certification Authorities	875
	Analyzing Certificate Needs within the Organization	881
	Determining Appropriate CA Type(s)	881
	Enterprise CAs	882
	Stand-Alone CAs	882
	Planning the CA Hierarchy	883
	Planning CA Security	885
	Certificate Revocation	886
6/6.1/6.2.2	Planning Enrollment and Distribution of Certificates	887
	Certificate Templates	887
	Certificate Requests	892
	Auto-Enrollment Deployment	895
	Role-Based Administration	896
6/6.2.3	Implementing Smart Card Authentication in the PKI	897
	What Are Smart Cards?	897
	How Smart Card Authentication Works	898
	Deploying Smart Card Logon	898
	Smart Card Readers	899
	Smart Card Enrollment Station	899
	Using Smart Cards To Log On to Windows	899
	Using Smart Cards for Remote Access VPNs	903
	Using Smart Cards To Log On to a Terminal Server	906
	Summary of Exam Objectives	907
	Exam Objectives Fast Track	908
	Exam Objectives Frequently Asked Questions	910
	Self Test	912
	Self Test Quick Answer Key	918
	Self Test Appendix	919
	Index	1025

Foreword

This book's primary goal is to help you prepare to take and pass Microsoft's exam number 70-293, *Planning and Maintaining a Microsoft Windows Server 2003 Network Infrastructure*. 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 70-293?

Exam 70-293 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 70-293 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 70-293 covers the basics of managing and maintaining the network infrastructure in a network environment that is built around Microsoft's Windows Server 2003. Objectives are task-oriented, and include the following:

- **Planning a secure baseline installation**, including planning a strategy to enforce system default security settings on new systems, identifying client operating system default security settings, and identifying all server operating system default security settings.

- **Planning and configuring security for servers that are assigned specific roles**, including domain controllers, Web servers, database servers, and mail servers. This includes deploying the security configuration for servers assigned to these specific roles and creating custom security templates based on server roles.
- **Evaluating and selecting the operating system to install on computers in an enterprise**, including identifying the minimum configuration to satisfy security requirements.
- **Planning a TCP/IP network infrastructure strategy**, including analyzing IP addressing requirements, planning an IP routing solution, and creating an IP subnetting scheme.
- **Planning and modifying a network topology**, including planning the physical placement of network resources and identifying network protocols to be used.
- **Planning an Internet connectivity strategy**.
- **Planning network traffic monitoring**, using tools such as Network Monitor and System Monitor.
- **Troubleshooting connectivity to the Internet**, including diagnosing and resolving issues related to Network Address Translation (NAT), name resolution cache information, and client configuration.
- **Troubleshooting TCP/IP addressing**, including diagnosing and resolving issues related to client computer configuration and DHCP server address assignment.
- **Planning a host name resolution strategy**, including planning the DNS namespace design, planning zone replication requirements, planning a forwarding configuration, planning for DNS security, and examining the interoperability of DNS with third-party DNS solutions.
- **Planning a NetBIOS name resolution strategy**, including planning a WINS replication strategy and planning NetBIOS name resolution by using the Lmhosts file.
- **Troubleshooting host name resolution**, including diagnosing and resolving issues related to DNS services and client computer configuration.
- **Planning a routing strategy**, including identifying routing protocols to use in a specified environment and planning routing for IP multicast traffic.
- **Planning security for remote access users**, including planning remote access policies, analyzing protocol security requirements and planning authentication methods for remote access clients, offering remote assistance to client computer, and performing remote administration using terminal services.

- **Implementing secure access between private networks**, including creating and implementing an IPSec policy.
- **Troubleshooting TCP/IP routing**, using tools such as ROUTE, TRACERT, PING, PATHPING, and NETSH, as well as the Network Monitor.
- **Planning services for high availability**, including planning high availability solutions that use clustering services and Network Load Balancing (NLB).
- **Identifying system bottlenecks**, including memory, processor, disk and network related bottlenecks, using System Monitor.
- **Implementing a cluster server and recovering from cluster node failure.**
- **Monitoring Network Load Balancing**, using tools such as the NLB Monitor MMC snap-in and the WLBS cluster control utility.
- **Monitoring servers that provide network services**, using tools such as System Monitor, Event Viewer, and service logs.
- **Planning a backup and recovery strategy**, including identifying appropriate backup types such as full, incremental and differential, planning a backup strategy that uses volume shadow copies, and planning system recovery that uses Automated System Recovery (ASR).
- **Configuring network protocol security**, including configuring protocol security in a heterogeneous client computer environment and configuring protocol security by using IPSec policies.
- **Configuring security for data transmission**, including configuring IPSec policy settings.
- **Planning for network protocol security**, including specifying the required ports and protocols for specified services and planning an IPSec policy for secure network communications.
- **Planning secure network administration methods**, including creating a plan to offer Remote Assistance to client computers and planning for remote administration by using terminal services.
- **Planning security for wireless networks.**
- **Planning security for data transmission**, including securing data transmissions between client computers to meet security requirements and securing data transmissions by using IPSec.
- **Troubleshooting security for data transmission**, using tools such as the IPSec Monitor MMC snap-in and the Resultant Set of Policies (RSOP) MMC snap-in.

- **Configuring the Active Directory directory service for certificate publication.**
- **Planning a public key infrastructure (PKI) that uses Certificate Services,** including identifying the appropriate type of certificate authority to support certificate issuance requirements, planning the enrollment and distribution of certificates, and planning for the use of smart cards for authentication.
- **Planning a framework for planning and implementing security,** including planning for security monitoring and planning a change and configuration management framework for security.
- **Planning a security update infrastructure,** using tools such as the Microsoft Baseline Security Analyzer and Microsoft Software Update Services.

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/70-293.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 70-293 will earn you the MCP certification (if it is the first Microsoft exam you've passed). Exam 70-293 also counts toward the MCSE. Exam 70-293 is *not* a requirement or elective for the MCSA.



TIP

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 70-293 does not apply to any of these specialty and developer certifications.

Prerequisites and Preparation

There are no mandatory prerequisites for taking Exam 70-293, although Microsoft recommends that you meet the target audience profile described earlier, and many candidates will first take Exams 70-290 and 70-291 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/70-293.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 the book:

- **Planning tools and documentation** We begin with an overview of network infrastructure planning, introducing you to planning strategies and how to use planning tools. We will review the fundamentals of network design, including analysis of organizational needs. This includes such factors as information flow, management model and organizational structure, and centralization vs. decentralization issues. We discuss management priorities, including availability and fault tolerance, security, scalability, performance and cost. Next, we address user priorities, which include email communications, scheduling and task management, project collaboration, data storage and retrieval, Internet research, application services, print services and graphics/audio/video services. This chapter also looks at legal and regulatory considerations, how to calculate Total Cost of Ownership (TCO) and how to plan for future growth. We discuss how to develop a test network environment, and how to document the planning and network design process.
- **Planning server roles and server security** You will first review server roles and ensure that you have an understanding of the many roles a Windows Server 2003 server can play on the network. We discuss domain controllers, file and print servers, DHCP, DNS and WINS servers, Web servers, database servers, mail servers, certification authorities and terminal services application servers. Then we delve into how to plan a server security strategy. Here we examine how to choose the right operating system according to security needs, how to identify minimum security requirements for your organization and how to identify the correct configurations to satisfy those security requirements. You will learn how to plan baseline security, first planning the secure baseline installation parameters and then enforcing default security settings on new computers, both client and server machines. We'll show you how to customize server security, securing your servers according to their roles. Then we'll walk you through the process of creating custom security templates and show you how to deploy security configurations.

- **Planning, Implementing and Maintaining the TCP/IP infrastructure** We then examine the TCP/IP infrastructure, and you will learn all about the network protocols supported by Windows Server 2003 and how to identify the protocols to be used in your network environment. We discuss the advantages of the TCP/IP protocol suite and we also address the multi-protocol environment that is increasingly common in today's business organizations. We will review TCP/IP basics, and then get into what's new in TCP/IP for Server 2003. Specifically, we'll discuss IGMP v3, IPv6 support, the alternate configuration feature, and automatic determination of interface metric. You'll find out how to plan an IP addressing strategy, including how to analyze your addressing requirements and how to create an effective subnetting scheme. Then we will address methods for troubleshooting IP addressing problems, both those related to client configuration and those related to DHCP server issues. You'll learn about transitioning to the next generation of IP, IPv6, and we'll introduce IPv6 utilities such as Netsh commands, Ipsec6.exe, and the IPv6 PING and TRACERT parameters. We discuss 6to4 tunneling, the IPv6 Helper service, and connecting to the 6bone. Next, we'll discuss the planning of the network topology. This includes analysis of hardware requirements and how to plan for the placement of physical resources. You'll learn to plan network traffic management, and how to monitor network traffic and devices using Network Monitor and System Monitor. We'll show you how to determine bandwidth requirements and how to optimize your network's performance.
- **Planning, implementing and maintaining a routing strategy** We first review the basics of IP routing, including the role of routing tables, static and dynamic routing, and routing protocols such as RIP and OSPF. You'll learn to use the netsh commands related to routing, and then we'll show you how to evaluate routing options. This includes selecting the proper connectivity devices, and we'll discuss hubs, bridges, switches (layer 2, 3 and 4 varieties), and routers. We will look at how you can use a Windows Server 2003 machine as a router, and how to configure the Routing and Remote Access Service (RRAS) to do so. Next, we look at security considerations related to routing. We'll show you how to analyze requirements for routing components from a security-conscious point of view, and discuss methods of simplifying the network topology to provide fewer attack points. This includes minimizing the number of network interfaces, the number of routes, and the number of routing protocols. We will also discuss router to router VPNs and packet filtering and firewalls, as well as setting the logging level. Finally, we cover how to troubleshoot IP routing issues. We'll identify troubleshooting tools and take a look at some common routing problems, including those related to interface configuration, to RRAS configuration, to routing protocols, to TCP/IP configuration and to routing table configuration.

- **Planning, implementing and maintaining an Internet connectivity strategy** We then turn to how to develop the best strategy for connecting your company's Windows Server 2003 network to the Internet. We discuss connecting the LAN to the Internet using routed connections or translated connections (via Internet Connection Sharing or the RRAS Network Address Translation component). You'll learn about virtual private networking, and how to use both Internet-based VPNs and router-to-router VPNs to provide connectivity to the company's LAN from remote locations or connect two branch offices. We discuss the intricacies of demand-dial/on-demand connections and persistent connections, and explain the difference between one-way and two-way initiation. We also show you how to use remote access policies to control VPN connections, and we discuss VPN protocols supported by Windows Server 2003 and how to make VPN connections using either the Point to Point Tunneling Protocol (PPTP) or the Layer 2 Tunneling Protocol (L2TP). You'll learn about VPN security and the authentication and encryption protocols that make your virtual network private. Next, we take a look at the Internet Authentication Service (IAS), and how it can provide centralized user authentication and authorization, centralized auditing and accounting, and extensibility and scalability. You'll learn about IAS integration with Server 2003 RRAS and how to control authentication via remote access policies. We show you how to use the IAS MMC snap-in and how to implement monitoring of IAS, and we discuss the use of the IAS Software Developers' Kit (SDK). Then we delve a little deeper into the IAS authentication methods, and discuss RADIUS access server support, wireless access points and authenticating switches. In the next section, we walk you through the process of using the Connection Manager Administration Kit (CMAK) to create service profiles, custom actions and custom Help, as well as VPN support, to make it easier for non-technical users to connect remotely without having to do complex configuration. We'll talk about security issues pertaining to Connection Manager, and show you how to prevent editing of service profile files, how to prevent users from saving their passwords, and how to distribute service profiles securely.
- **Planning, implementing and maintaining a name resolution strategy** You will learn how to plan for the best way of resolving host names on your network. We'll present an overview of host naming, and how host names are resolved using the hosts file and using DNS. We'll discuss issues involved in designing a DNS namespace, such as choosing the parent domain name, the conventions and limitations that govern host names, the relationship of DNS and the Active Directory, and how to support multiple namespaces. Then we move on to planning DNS server deployment. You'll find out how to factor in such things as number of servers, server roles, server capacity and server placement. We'll also show you how to plan for zone replication between your DNS servers, and we'll address planning

for forwarding and how DNS interacts with DHCP on a Server 2003 network. We'll discuss Server 2003 DNS server interoperability with BIND and other non-Windows DNS implementations. You'll learn about zone transfers between Server 2003 DNS servers and BIND servers, and we'll discuss supporting Active Directory with BIND. You'll learn about split DNS configurations and how interoperability relates to other services such as WINS and DHCP. Next, we address DNS security issues, including common DNS threats such as footprinting, redirection and DNS DoS attacks. You'll learn how to best secure your DNS deployment, using a split namespace and using packet filtering. We'll discuss how to determine the best DNS security level for your network. Next, we look at DNS performance issues. We show you how to monitor DNS server performance and how to analyze DNS server tests. In the next section, we'll address NetBIOS name resolution and provide an overview of how NetBIOS names are resolved using lmhosts files and NetBIOS Name Servers such as WINS servers. You'll find out what's new for WINS in Server 2003, and we'll show you how to plan WINS server deployment and how to plan for WINS replication. We'll walk you through the process of configuring WINS replication partnerships, including Push Only, Pull Only and Push/Pull configurations. We'll also discuss common WINS issues, including configuration issues, performance issues and security issues. We'll show you how to plan for WINS database backup, and how to troubleshoot name resolution problems related to both host names and NetBIOS names.

- **Planning, implementing and maintaining a remote access strategy** We examine the issues and procedures involved in devising a remote access strategy, including planning tasks such as analyzing organizational needs, analyzing user needs, and selecting the remote access types that will be allowed (dial-in, VPN, and/or wireless). We'll discuss design considerations related to dial-in access, such as the allocation of IP addresses, how to determine incoming port needs, and how to select the best administrative model based on your organizational needs and the functional level of your domain. Next, we'll talk about design considerations related to VPN access. You'll learn how to select the VPN protocols to be allowed, based on client support, PKI requirements and the need for data integrity and sender authentication. You'll learn how to install machine certificates, how to configure firewall filters, and how to create access policies governing VPN connections. In the next section, you'll learn about the design considerations that relate to wireless remote access. We'll discuss the use of IAS for wireless connections, and how to configure remote access policies for wireless connections. We'll address the use of multiple wireless access points, and the advantages of placing a certification authority on a Virtual LAN (VLAN) for new wireless clients. We'll also show you how to configure wire access points (WAPs) as RADIUS clients. Next, we move on to planning overall security strategies for remote access connections. We'll dis-

cuss the best practices in selecting authentication methods that will be allowed, and the benefits of disallowing insecure password based connections such as PAP, SPAP, CHAP and MS-CHAPv1). We'll then look at the more secure methods such as MS-CHAPv2 and EAP, and discuss the advantages of using RADIUS/IAS rather than Windows authentication. We'll also address the selection of the data encryption level, and other security measures such as requiring callback, mandating operating system and file system choices, using managed connections and using smart cards for remote access. We'll delve deeply into the subject of remote access policies, and show you how to authorize remote access by user or group, how to restrict remote access in various ways, and how to control remote connections.

- **Planning, implementing and maintaining a high availability strategy** We then look at the concept of high availability and how it can be attained. We'll provide an overview of performance bottlenecks and what causes them, and show you how to identify such common system bottlenecks as memory, processor, disk and network components. We'll walk you through the steps of using the System Monitor to monitor server performance, and show you how to use Event Viewer and service logs to monitor server issues, as well. Next, we show you how to plan a backup and recovery strategy. We'll introduce you to the Windows Backup utility, and ensure that you understand the differences between full, incremental and differential backups. We'll also discuss the use of volume shadow copies as a backup option. You'll learn how to decide what information should be backed up, and we'll show you how to back up user data, system state data, the DHCP, WINS and DNS databases and cluster disk signatures and partition layouts. We'll walk you through the process of using the Windows Backup administrative tool, including the Backup and Restore Wizard feature and the Advanced Mode feature. We'll also discuss the use of command line tools. Next, we'll talk about how to select your backup media, and you'll learn about scheduling backups and how to restore data from backup when necessary. In the next section, we'll address how to plan for system recovery using the Automated System Recovery (ASR). You'll learn about system services, how to make an ASR backup and how to do an ASR restore. We'll explain how ASR works, and discuss alternatives to ASR such as Safe Mode boot and Last Known Good. Finally, we'll discuss the importance of planning for fault tolerance, including solutions aimed at providing fault tolerance for local network connectivity, for Internet connectivity, for data on disk, and for mission-critical servers.
- **Windows Cluster Services and Network Load Balancing** We will look at the ultimate in fault tolerance: server clustering, and shows you how you can make clustering services part of your enterprise-level organization's high availability plan. We'll start by introducing you to the terminology and concepts involved in understanding clustering; you'll learn about cluster nodes, cluster groups, failover and fail-back, name resolution as it pertains to cluster services, and how server clustering

works. We'll discuss three cluster models: single node, single quorum device and majority node set. Then we'll talk about cluster deployment options, including N-node failover pairs, hot standby server/N+1, failover ring and random. You'll learn about cluster administration and we'll show you how to use the cluster administrator tool as well as provided command line tools. Next, we'll discuss best practices for deploying server clusters. You'll learn about hardware issues, especially those related to network interface controllers, storage devices, power saving features and general compatibility issues. We'll discuss cluster network configuration and you'll learn about multiple interconnections and node-to-node communication. We'll talk about the importance of binding order, adapter settings, and TCP/IP settings, and we'll discuss the default cluster group. Next, we'll move on to the subject of security for server clusters. This includes physical security, public/mixed networks, private networks, secure remote administration of cluster nodes, security issues involving the cluster service account and how to limit client access. We'll also talk about how to secure data in a cluster, how to secure disk resources, and how to secure cluster configuration log files. The next section addresses how to make Network Load Balancing (NLB) part of your high availability plan. We introduce you to NLB concepts such as hosts/default host, load weight, traffic distribution and convergence and heartbeats. You'll learn how NLB works, and the relationship of NLB to clustering. We'll show you how to manage NLB clusters using the NLB Manager tool, remote management and the command line tools. We'll also discuss NLB error detection and handling. Next, we'll move on to monitoring NLB using the NLB Monitor MMC snap-in or using the Windows Load Balancing Service (WLBS) cluster control utility. We discuss best practices for implementing and managing NLB, including issues such as multiple network adapters, protocols and IP addressing, and NLB Manager logging. Finally, we address NLB security.

- **Planning, implementing and maintaining Internet Protocol Security** We then turn to Windows Server 2003's implementation of the Internet Protocol Security protocol (IPSec). We start by introducing IPSec terminology and concepts and explaining how IPSec works "under the hood" to secure data in transit over the network. We discuss the purposes of IPSec encryption: authentication, integrity and confidentiality. You'll learn about how IPSec operates in either of two modes: tunnel or transport. You'll also learn about the protocols used by IPSec. These include the two primary protocols: the Authentication Header (AH) protocol and the Encapsulating Security Payload (ESP) protocol. We'll also discuss the roles of additional protocols used by IPSec, including the Internet Security and Key Management Protocol (ISAKMP), Internet Key Exchange (IKE), the Oakley key determination protocol and the Diffie-Hellman key agreement protocol. You'll also learn about Server 2003's IPSec components such as the IPSec driver and we'll discuss the relationship of IPSec to IPv6. Next, we'll show you how to deploy IPSec

on your network, taking into consideration organizational needs and security levels, and help you determine the appropriate authentication methods. You'll learn about managing IPsec and we'll walk you through the process of using the IPsec MMC snap-in as well as the command line tools. We'll discuss the role of IPsec policies, including default and custom policies, and we'll show you how to assign and apply policies. We'll also talk about IPsec security considerations and issues, including the use of a strong encryption algorithm (3DES), authentication methods, firewall packet filtering, unprotected traffic, Diffie-Hellman groups and the use of pre-shared keys. We'll show you how to use RSoP and the RSoP MMC snap-in to view policy assignments and to simulate policy assignments for deployment planning.

- **Planning, implementing and maintaining a security framework** We look at several aspects of creating an effective security framework for your organization's network. First, we look at how to plan and implement Active Directory security. This includes such measures as physically securing domain controllers, securing the schema, managing cross-forest security relationships, account security and implementing Active Directory access controls. Next, we discuss the issues and procedures involved in planning and implementing wireless security. We'll provide an overview of the terminology and concepts relating to 802.11 wireless technologies and you'll learn about authenticators and supplicants, as well as how wireless networking works "under the hood." We'll discuss authentication methods for wireless networks, including such authentication subtypes as open system and shared key. You'll learn about the protocols generally used for wireless authentication, including the Extensible Authentication Protocol (EAP), EAP-Transport Layer Security (EAP-TLS), EAP-MS-CHAPv2, and the Protected Extensible Authentication Protocol (PEAP). We'll also talk about using IAS with wireless. We'll address wireless security issues such as common insecure default settings (administrative password, SSID, and WEP settings) and the weaknesses of Wired Equivalent Privacy protocol (WEP) encryption, as well as how WEP can be made more secure. Next, we'll move on to discuss security monitoring, and we'll address object based access control and security policies, including password policies, Kerberos policies, account lockout policies, user rights and the use of security templates. We'll also talk about security auditing, and you'll learn to set the auditing policy, modify the security log settings and audit objects such as files or folders. In the next section, you'll learn about planning a Change and Configuration Management framework. We'll walk you through the steps of using the Security Configuration Manager tool as well as command line tools included with Windows Server 2003. We'll also discuss Security Analysis and Configuration best practices. Finally, we take you through the process of planning a security update infrastructure. You'll understand the importance of regular security updates and you'll learn

to use the Microsoft Baseline Security Analyzer (MBSA) and the Microsoft Software Update Services to ensure that your Server 2003's security features are always current.

- **Planning, implementing and maintaining a public key infrastructure** We will examine the complex issues involved in planning a certificate based PKI. We'll provide an overview of the basic terminology and concepts relating to the public key infrastructure, and you'll learn about public key cryptography and how it is used to authenticate the identity of users, computers, and applications/services. We'll discuss the role of digital certificates and the different types of certificates (user, machine and application certificates). You'll learn about certification authorities (CAs), the servers that issue certificates, including both public CAs and private CAs such as the ones you can implement on your own network using Server 2003's certificate services. Next, we'll discuss the CA hierarchy, and how root CAs and subordinate CAs act together to provide for your organization's certificate needs. You'll find out how the Microsoft certificate services work, and we'll walk you through the steps involved in implementing one or more certification authorities based on the needs of the organization. You'll learn to determine the appropriate CA type—enterprise or standalone CA—for a given situation, and how to plan the CA hierarchy and provide for security of your CAs. We'll show you how to plan for enrollment and distribution of certificates, including the use of certificate requests, role based administration and autoenrollment deployment. Next, we'll discuss how to implement the use of smart cards for authentication within the PKI. You'll learn what smart cards are and how smart card authentication works, and we'll show you how to deploy smart card logon on your network. We'll discuss smart card readers and show you how to set up a smart card enrollment station. Finally, we'll discuss the procedures for using smart cards to log onto Windows, for remote access and VPNs and to log onto a terminal server.

Exam Day Experience

Taking the exam is a relatively straightforward process. Both Vue and Prometric testing centers administer the Microsoft 70-293 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.

MCSE 70-293

Using Windows Server 2003 Planning Tools and Documentation

Solutions in this chapter:

- ☑ Overview of Network Infrastructure Planning
 - ☑ Analyzing Organizational Needs
 - ☑ Developing a Test Network Environment
 - ☑ Documenting the Planning and Network Design Process
-
- ☑ Summary of Exam Objectives
 - ☑ Exam Objectives Fast Track
 - ☑ Exam Objectives Frequently Asked Questions
 - ☑ Self Test
 - ☑ Self Test Quick Answer Key

Introduction

Planning is the first step in building a reliable, secure, high-performance and highly available Windows Server 2003-based network. In this chapter, we begin with an overview of network infrastructure planning, introducing you to planning strategies and how to use planning tools.

We will review the fundamentals of network design, including analysis of organizational needs. These include factors such as information flow, management model, organizational structure, and issues of centralization versus decentralization. We discuss management priorities, including availability and fault tolerance, security, scalability, performance, and cost. Next, we address user priorities, which include e-mail communications, scheduling and task management, project collaboration, data storage and retrieval, Internet research, application services, print services, and graphics/audio/video services.

This chapter also looks at legal and regulatory considerations, how to calculate total cost of ownership (TCO), and how to plan for future growth. We discuss how to develop a test network environment, and how to document the planning and network design process.

Overview of Network Infrastructure Planning

Proper planning of a network infrastructure is essential to ensuring high performance, availability, and overall user satisfaction with your network operations. In order to create a viable network design, you'll need an understanding of both the business requirements of your organization as well as current and emerging networking technologies. Accurate network planning will allow your organization to maximize the efficiency of its computer operations, lower costs, and enhance your overall business processes.

When planning for a new infrastructure or upgrading an existing network, you should take some or all of the following steps:

- Document the business requirements of your client or organization.
- Create a baseline of the performance of any existing hardware and network utilization.
- Determine the necessary capacity for the physical network installation, including client and server hardware, as well as allocating network and Internet bandwidth for network services and applications.
- Select an appropriate network protocol and create an addressing scheme that will provide for the existing size of the network and that will allocate room for any foreseeable expansions, mergers, or acquisitions.
- Specify and implement technologies that will meet the existing needs of your network, while allowing room for future growth.

- Plan to upgrade and/or migrate any existing technologies, including server operating systems and routing protocols.

In this section, we'll discuss best practices and strategies for planning your network implementation. We'll then look at the various tools that you can use for network planning, both from Microsoft and from other vendors. We'll conclude with some fundamentals of network design that will provide you with a good starting point for designing a network that will best meet the needs of your organization and its users.

Planning Strategies

When designing a new network, you should first use the business requirements of your organization as the primary source of planning information. You'll need to create a network infrastructure that addresses the needs of your management structure, such as fault tolerance, security, scalability, performance, and cost. You'll need to balance these requirements with the types of services that your users and clients will expect from a modern network, including e-mail, calendaring, project collaboration, Internet access, file, print, and application services.

After you've determined the business requirements of your network, you should then analyze the technical requirements of your organization. These requirements may apply to any applications that are already in use or that you plan to implement, as well as to the associated hardware and operating system. You should carefully note all of these requirements so that you won't create any difficulties later on during the implementation process. Be sure to analyze and document the existing network, including any hardware, software, and network services that are already in place. This will make it easier to take the existing configuration into account when planning the new or upgraded network.

Finally, any well-formed network plan should make allowances for future changes to the organization, including support for new technologies and operating systems, as well as additional hardware and users. Your organization's business requirements can change—through a merger, an acquisition, or simple growth and expansion. Although it is impossible to foresee all possible changes of this nature, a good network design will be flexible enough to accommodate as many adjustments as possible.

Using Planning Tools

There are a number of tools available to assist you in developing a plan for your network infrastructure. The first and best of these, however, might be the simplest: pencil and paper. As we discussed in the previous section, you should begin your planning by determining the requirements of the business that will be using the network. The best way to do this is through face-to-face interactions, by interviewing relevant managers and staff members of each department, branch, or business unit. Not only does this allow you to construct a complete picture of your network requirements, but it also involves stakeholders from the various departments. This sort of involvement is critical in ensuring the successful deployment of any new or upgraded technology.

After you have a high-level understanding of your company's organizational structure and computing needs, you should inventory the hardware and software that is already in place. In a small office environment, you can accomplish this by simply taking a walk to determine the physical layout of network cables, routers, and the like. In a medium- to large-sized enterprise network, you will probably want to rely on automated inventory tools such as Microsoft's Systems Management Server (SMS) or a third-party equivalent. Take as detailed of an inventory as possible, including the hardware configuration of server and workstation machines as well as vendor names and the version numbers of the operating system and business applications the systems are running.

You can use a network analyzer, such as the Network Monitor utility built into the Windows Server 2003 operating system or the more full-featured version of Network Monitor included in SMS, to create a baseline of the current utilization of your network bandwidth. If this utilization is already near capacity, you can use this baseline to justify and plan upgrades to your network infrastructure (moving from 10MB Ethernet to 100MB Ethernet, for example).



EXAM WARNING

The version of Network Monitor that ships with Windows Server 2003 can analyze only traffic addressed to the network interface card (NIC) on the server itself or that is sent by the server on which it is running. The SMS version of Network Monitor operates in *promiscuous mode*, enabling it to capture all network traffic on a given segment, even if the traffic isn't addressed to or from the local server.

Windows Server 2003 has introduced new management features that will assist you in planning your network configuration, especially in the areas of user and computer management. The Resultant Set of Policy (RSOP) Microsoft Management Console (MMC) snap-in contains a Group Policy modeling function that will allow you to simulate changes to Group Policy Objects (GPOs) in an Active Directory (AD) environment before actually applying them to a production network. For example, if you want to apply a new GPO to a departmental Organizational Unit (OU), the modeling report will indicate how the new GPO will affect the objects within the OU to which it's being applied. The Group Policy Management Console (GPMC) can also provide detailed configuration reports on existing GPO settings in place on a Windows 2000 or Windows Server 2003 AD installation.

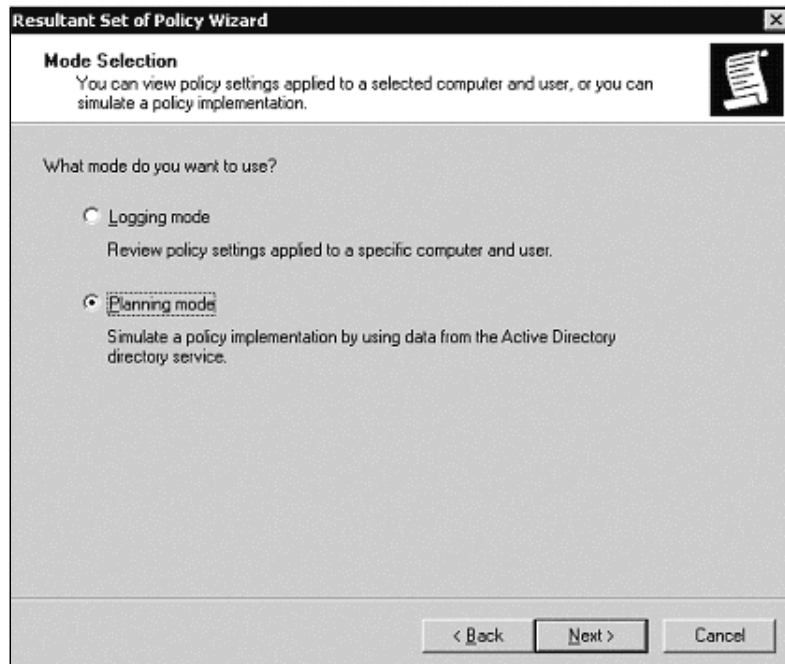
EXERCISE 1.01

GENERATING A GROUP POLICY MODELING REPORT

In this exercise, we'll take a look at a GPMC modeling report for a Windows Server 2003 domain.

1. Click **Start | Run**, type **mmc**, and click **OK**.
2. Click **File | Add/Remove Snap-in**, and then select the **Resultant Set of Policy** snap-in. Click **Add**, and then click **Close**.
3. Right-click **Resultant Set of Policy**, and then click **Generate RSoP Data**. Click **Next** to bypass the initial Welcome screen.
4. On the **Mode Selection** page, select **Planning mode** as shown in Figure 1.1, and then click **Next**.

Figure 1.1 Selecting the RSoP Report Mode



5. On the **User and Computer Selection** page, shown in Figure 1.2, specify the name of the user and computer that you wish to analyze, and then click **Next**. Alternatively, you can select an entire user and/or computer container (such as a site, domain, or OU) to analyze.

Figure 1.2 Specifying the User and Computer Information

The screenshot shows the 'Resultant Set of Policy Wizard' dialog box, specifically the 'User and Computer Selection' page. The title bar reads 'Resultant Set of Policy Wizard'. Below the title bar, the page title is 'User and Computer Selection' with a subtitle: 'You can view simulated policy settings for a selected user (or a container with user information) and computer (or a container with computer information)'. There is a help icon in the top right corner. The main area contains the following text: 'Example container name: CN=Users,DC=airplanes,DC=com' and 'Example user or computer: AIRPLANES\Administrator'. Below this, it says 'Simulate policy settings for the following:'. There are two sections: 'User information' and 'Computer information'. Each section has a radio button for 'Container:' (selected) and a text box containing 'DC=airplanes,DC=com', along with a 'Browse...' button. The 'Computer information' section also has a radio button for 'Computer:' which is unselected. At the bottom of the main area, there is a checkbox labeled 'Skip to the final page of this wizard without collecting additional data' which is unselected. At the very bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

6. From the **Advanced Simulation Options** page, shown in Figure 1.3, you can choose to modify a number of reporting options, such as simulating a slow network connection or the use of loopback processing. Click **Next** when you're ready to continue.

Figure 1.3 Advanced Simulation Options

The screenshot shows the 'Resultant Set of Policy Wizard' dialog box, specifically the 'Advanced Simulation Options' page. The title bar reads 'Resultant Set of Policy Wizard'. Below the title bar, the page title is 'Advanced Simulation Options' with a subtitle: 'You can select additional options for your simulation.' There is a help icon in the top right corner. The main area contains the following text: 'Simulate policy implementation for the following:'. There are two checkboxes: 'Slow network connection (for example, a dial-up connection)' and 'Loopback processing', both of which are unselected. Under 'Loopback processing', there are two radio buttons: 'Replace' and 'Merge', both of which are unselected. Below this, there is a 'Site:' label and a dropdown menu currently showing '(None)'. At the bottom of the main area, there is a checkbox labeled 'Skip to the final page of this wizard without collecting additional data' which is unselected. At the very bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

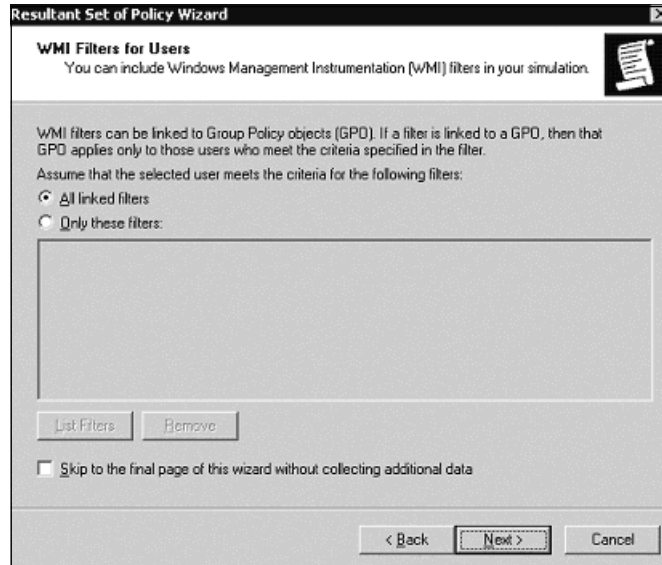
7. On the **User Security Groups** page, shown in Figure 1.4, you'll see the security groups to which the specified user belongs. You can use the **Add** or **Remove** buttons to specify different security group memberships to simulate. (If you make a mistake, you can click **Restore Defaults** to return to the user's actual group membership.) Click **Next** when you're ready to continue.

Figure 1.4 Simulating User Security Group Membership



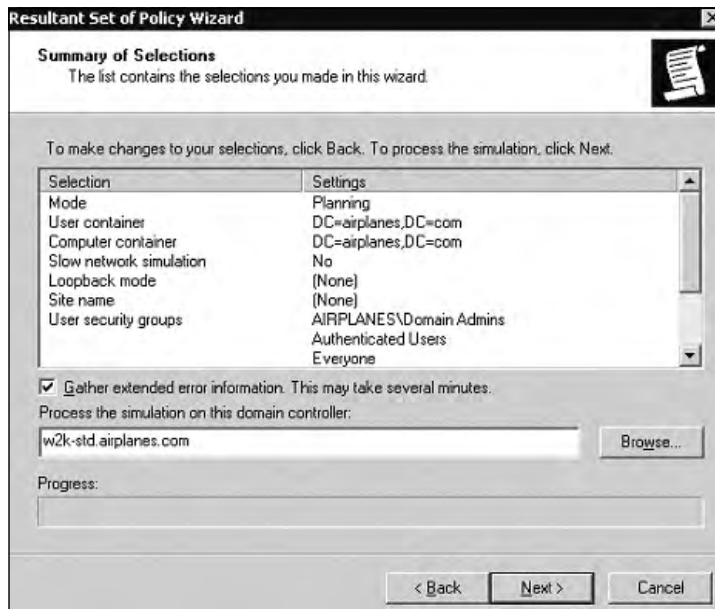
8. The next page lists the security groups to which the specified *computer* belongs. As in Step 7, you can use the **Add** or **Remove** buttons to change the contents of the RSoP report. Click **Next** to continue.
9. By default, the report will include all possible Windows Management Instrumentation (WMI) filters, as shown in Figure 1.5. (WMI filters allow you to apply GPOs to users or computers based on hardware and software attributes such as operating system, free hard drive space, and the like.) If you've created any WMI filters that would cause the computer you've specified to *not* be subject to Group Policy, you should remove them by clicking the **Only these filters** radio button and selecting **Remove**. Click **Next** to repeat the process for any computer-specific WMI filters.

Figure 1.5 Selecting WMI Filters



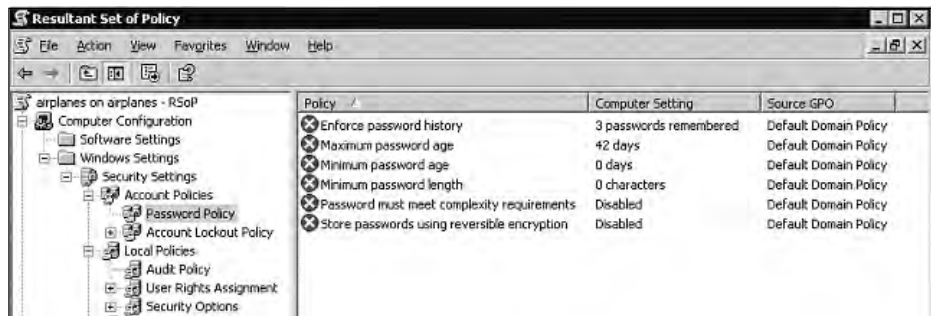
10. Click **Next** again. You'll see a summary of your choices, as shown in Figure 1.6. If you are satisfied with the selections you've made, click **Next** again to run the simulation.

Figure 1.6 RSOP Summary Screen



11. When the simulation has completed, click **Finish**. In the console tree, click the RSOP query to view the data. You'll see the output in a screen similar to the one shown in Figure 1.7.

Figure 1.7 A Completed RSOP Simulation



As you can see, Group Policy modeling will allow you to perform “what-if?” analyses to simulate the creation of new security groups or OUs. You can also use simulated WMI filters to see how GPO settings and inheritance would change if you upgraded a workstation from Windows NT to Windows XP Professional, for example. GPMC modeling is definitely a useful tool to have in your arsenal as you begin developing your Windows 2003 Server network design.

Fundamentals of Network Design

When you design a network, the most important question is unfortunately the most often overlooked: Why are you building the network to begin with? It's easy to become so excited about the new technologies available to you that you can overlook the business requirements of your organization. Even if you eventually configure the resultant network to meet your needs, it can become a far more complicated (and expensive) process than if you had begun by fully detailing business requirements in the first place. This can be even more hazardous when you are working as a consultant for an independent company, because you need to be very specific in obtaining the appropriate information from your clients. Too often, you'll hear, “We need a Frame Relay network” or, “We need you to install a Check Point firewall.” These statements give you a solution without telling you about the problem or need that the company is attempting to address. (Imagine walking into your doctor's office for the first time and telling her that you need your foot amputated, rather than simply reporting that you have an ingrown toenail.) It is important to use available technologies to meet business requirements, rather than implementing them for their own sake.

A company's business requirements can include a number of factors that you need to keep in mind. An obvious issue is that of *cost*, whether you are interested in improving user efficiency to save money, or pumping cash into high-powered server farms to increase sales revenue on an e-commerce site. You need to decide how much money your company is willing to spend, or how much money you expect a new technology to save the company. Either way, if your network design costs more than it ends up making (or saving) for a company, you've failed to meet this critical requirement. This will come up later in this chapter in the "Calculating TCO" section.

After you've determined the budget for your new network, you should take stock of the current state of your company's computing technology. Ask the following questions:

- What resources are already in place?
- How much needs to be upgraded or replaced?
- What can be reused in the new or upgraded network?

Plan Now or Pay Later

Although completely new network installations are becoming a rarity except when dealing with new construction, they do present their own unique challenges. When planning a new network installation, don't take even the most basic configuration items for granted. Here's a real-world example: A medical supply firm was moving from an environment consisting exclusively of mainframes and dumb terminals to an installation of networked PCs and servers. Part of the physical installation included running pipes under the flooring to allow the network cabling to run throughout the building. Unfortunately, the construction manager received his specifications from the mainframe administrator, who was relatively unfamiliar with PC technology.

The mainframe manager assumed that the PCs would use the same type of cable to connect to the routers and hubs that was used by the existing dumb terminals. He did not consult with the new LAN administrator, or he would have known that the new networked PCs would be using Category 5 (CAT5) Ethernet cabling, which proved to be roughly three times the diameter of the mainframe access terminal cabling. This error wasn't discovered until after the subfloor piping had already been laid; the LAN administrator quickly discovered that there wasn't enough physical room to run all the necessary cable drops through the too-small piping.

Rather than incur the increased cost of running the piping all over again, management tasked the LAN administrator with installing network connectors that would use the smaller network cabling. This created an excess of performance bottlenecks until the subfloor piping was rerun two years later. Remember this true tale of how a seemingly insignificant detail can escalate into a much larger problem when you're establishing the particulars of your network design plan.

There might be existing technologies that will need to be maintained and supported even after the new design is in place. Be sure to include budget information for performing all necessary upgrades and providing ongoing support for your legacy systems.

The next step in designing your network is to understand where your users are located. Understanding the physical geography of your company and its employees is critical in designing a cost-effective local area network (LAN) or wide area network (WAN). You'll not only need to determine where your users are located, but also the location of the services that they need to access. A geographically diverse user base can easily necessitate the installation of dedicated WAN links or a virtual private network (VPN). Understanding where your users and resources are located will also help you to determine the amount of network bandwidth that your design will require. Network planning tools such as a network traffic analyzer will help you to determine the amount of traffic generated by your users and clients. To determine bandwidth requirements, you must consider current traffic levels while always leaving room for growth.

Analyzing Organizational Needs

Understanding the needs of a business or other organization is a fundamental step in creating a well-designed network. In this section, we'll take a look at information flow—recognizing where data originates in your network and how it should be disseminated to the users and customers who require it. Next, we'll discuss the importance of understanding an organization's management structure and how you can use that information to design appropriate network services. We'll also discuss some common priorities for an organization's management group, as well as its more task- and project-oriented users. These range from factors such as performance and availability that affect an entire network, to more specific services and applications such as e-mail, file sharing, and audio/video services. All of these issues should be taken into account to ensure the overall success of your network design.

Information Flow Factors

If the “Information Age” moniker is to be believed, it only stands to reason that access to a company's information needs to be a top priority of any network design. This means that all necessary personnel need on-demand access to their critical data in order to understand how their company's profits and losses are occurring, to call up a customer's account information at a moment's notice, and to collate information from multiple sources to allow for effective decision making. The most successful organizations are those whose front-line employees have instant access to the information they need, rather than waiting for managers or central “gatekeepers” to disseminate scheduled or ad hoc reports.

Understanding information flow requires you to determine *where* your users are located, *what* data they need to access, *when* they need it, and *how* they need to access it to best perform their jobs—whether that job is running a quarterly sales report or a high-school fundraiser. Providing appropriate information flow can involve physical considerations such as sufficient bandwidth allocation, along with logical controls within the

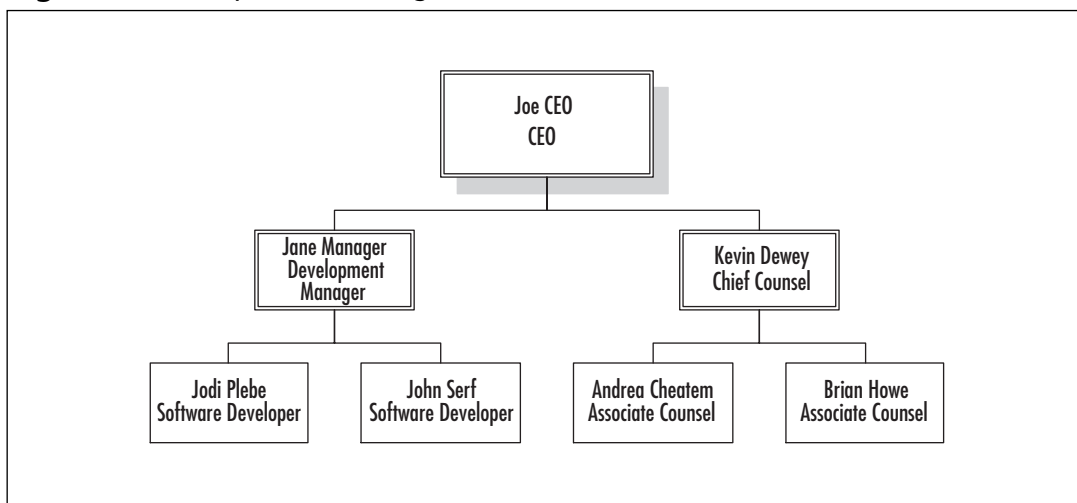
computer operating system. Remote and traveling users introduce their own unique challenges, because you will likely need to provide data access from varied and ever-changing locations around the globe. Whatever steps end up being necessary for your own network implementation, information flow can make or break a modern organization.

Management Model and Organizational Structure

Understanding a company's organizational structure is imperative in designing a network to meet its needs. You should begin by becoming familiar with the high-level divisions within an enterprise and how they related to one another. Large divisions usually have their own organizational structure, and they might be broken into several smaller departments or workgroups. For example, the Division of Finance might encompass separate Payroll, Accounts Payable, and Collections departments. Most companies have developed an organizational chart to provide a graphical illustration of this overall structure.

Once you have an understanding of the organizational structure, you can take a closer look at the individual departments themselves. Does the management structure of your organization have many levels, with Assistant Directors reporting to Directors, who report to Senior Directors, and so forth? (You can see an example of this sort of structure in Figure 1.8.) Or is the management model more flat in design, with a single manager taking responsibility for an entire department? This information will greatly benefit you when designing network functions such as user groups and AD OUs, as well as when you are determining appropriate delegation of network management responsibilities.

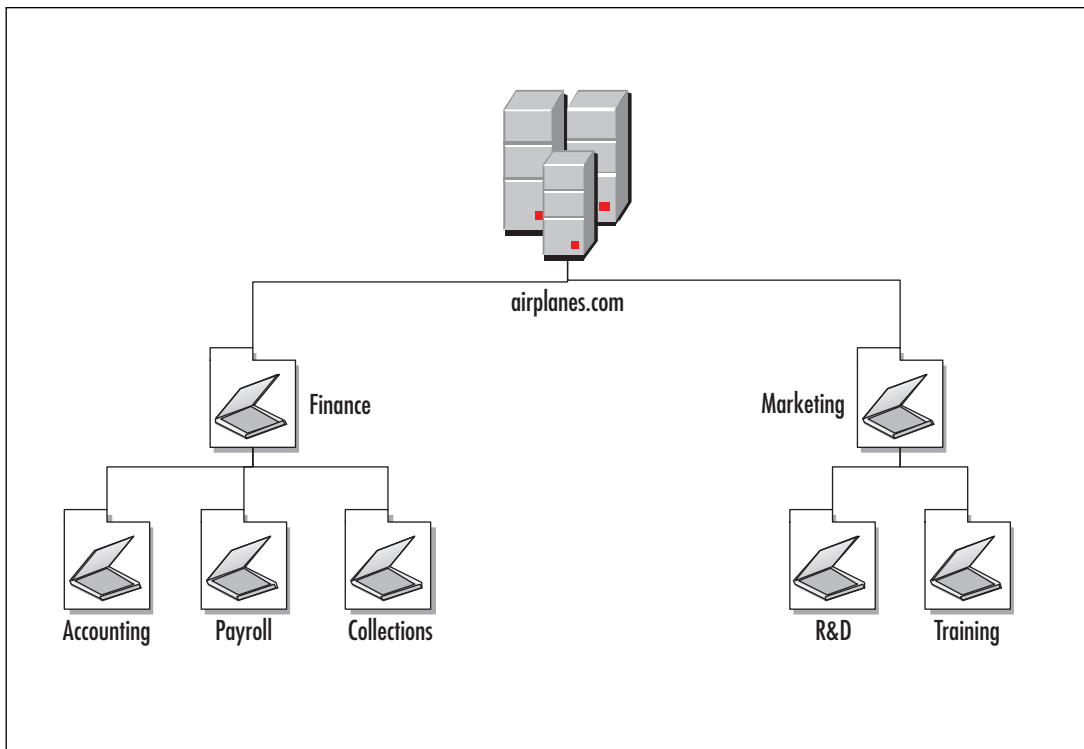
Figure 1.8 A Departmental Organizational Chart



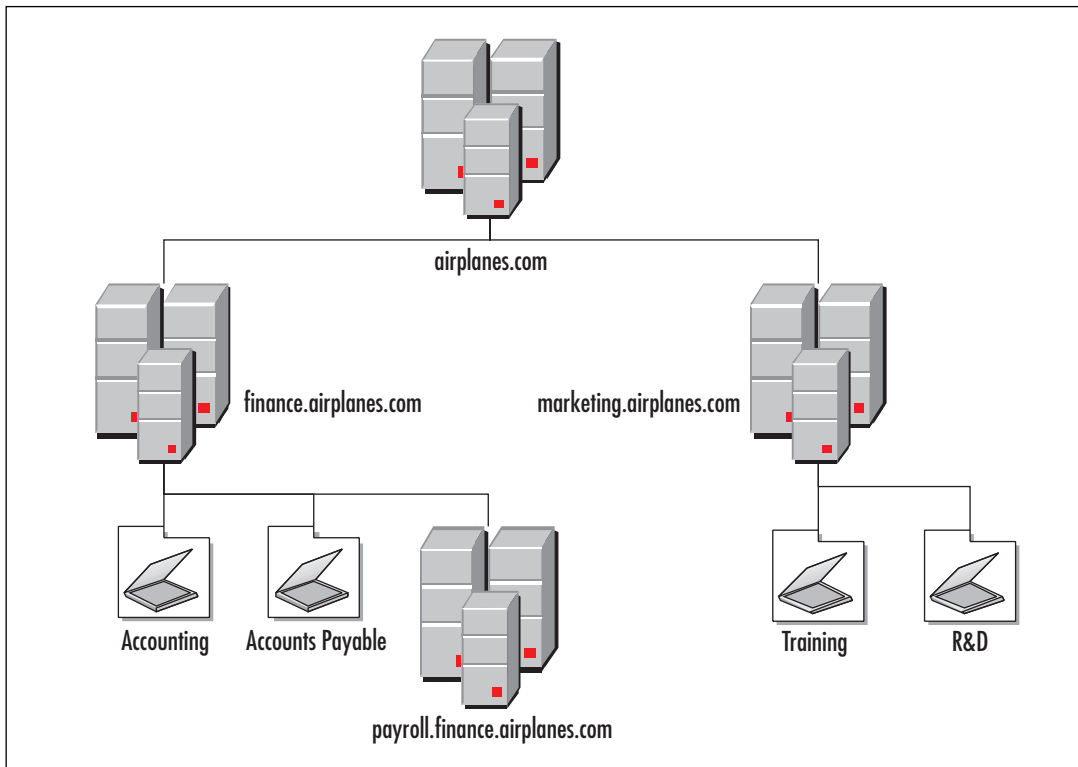
Centralization versus Decentralization

Once you've determined the organizational structure of your client or company, you should also recognize whether that structure is a centralized or decentralized one. Some companies adhere to a strictly hierarchical reporting structure in which the organizational chart resembles a family tree, with each sublevel reporting to a subsequently higher level and a single individual or group at the top of the hierarchy. In an AD environment, this type of structure lends itself to a system of nested OUs like the ones shown in Figure 1.9.

Figure 1.9 A Centralized Organizational Structure



Other organizational structures allow for greater autonomy within their business units, where various departments or project teams can function more independently. You might create an AD environment consisting of multiple domains, allowing each to maintain its own security requirements. And you can certainly mix-and-match these models to meet the unique requirements of your organization, as illustrated in Figure 1.10.

Figure 1.10 A Combination of Centralization and Decentralization

Your network design should also consider the Information Technology (IT) management structure of the organization. A company with a decentralized management structure can still handle network management centrally and vice versa. The transitive trust relationships built into Windows Server 2003 can allow centralized management of a multidomain or multiforest environment, or for tasks to be split among departmental IT administrators. The IT management structure of your organization can help you to decide how tasks such as user and group management should be structured and delegated.

Management Priorities

The management perspective of network design can be more conceptual, or high level, than the end-user priorities that we'll discuss in the next section. Rather than focusing on specific tasks and applications, a company's management structure should focus on design attributes that are common to and can benefit the entire organization, not just specific departments or workers. These include network availability, security, scalability, performance, and cost. When designing a network for an enterprise organization, be sure to address as many of these concerns as possible.

Availability/Fault Tolerance

As people and companies have become more reliant on computer technology to function and perform personal and business tasks, network designers have needed to contend with increasing expectations for “always on” availability. A sales manager traveling in Europe or Asia will not be pleased to find that although she can access her e-mail client, the data on the server itself is available only during business hours in the Eastern Standard time zone, or that a hardware failure will prevent her from accessing sales figures for eight hours while the server is being repaired. To avoid such difficulties, business-critical applications such as database and e-mail servers should be placed on systems that are designed for high availability whenever possible. This rationale applies even more to retail Web sites (e-commerce sites) and other Web-based businesses. Planning for high availability and fault tolerance will help you to minimize the downtime experienced by your end users and customers. Windows Server 2003 offers two separate but related clustering technologies—server clustering and Network Load Balancing—that can provide the high availability required by most enterprises.

Fault tolerance specifically refers to the ability of a piece of hardware or software to withstand the failure of a key component. This can be implemented at the hardware level using redundant power supplies or a Redundant Array of Inexpensive Disks (RAID) hard drive array. Advanced fault-tolerance technologies will even allow an administrator to replace individual components within a server without powering down the server. Clustering provides the ultimate in fault tolerance: completely redundant systems.



TEST DAY TIP

The ability to replace hardware on the fly, without powering down or rebooting the server, is referred to as *hot-swapping*.

Security

To create an effective network design, you must perform a juggling act between providing easy access to data for those who require it and, at the same time, protecting the data against unauthorized or illicit access. Accessibility and security are always at opposite ends of a continuum—more of one results in less of the other. Establishing an information security strategy is critical in ensuring that your network design is prepared to address security concerns *when*, not *if*, they arise.

A well-developed network security policy is as much a business concern as a technological one; consequently, you should involve key decision-makers from all parts of an organization, including Risk Management, Legal, Human Resources, and so on. Your security policy will provide a common baseline of security procedures based on your company's security requirements.

When addressing security concerns within your network design, your three primary concerns are the *confidentiality*, *integrity*, and *availability* of your data. These three security objectives answer the following key questions:

- Who has access to your data?
- Has the data been corrupted or altered in any way?
- Will your users be able to access their data when they need to?

All technologies and practices within information security will ultimately address one or more of these key concepts.

Scalability

When planning a network design, *scalability* refers to how well a service or application can grow to meet client performance demands that will inevitably increase over time. It can refer to increasing system resources such as processors, memory, disk drives, and network adapters to an existing piece of hardware, or being able to seamlessly replace existing hardware with more powerful equipment. It can also refer to adding new servers to meet increased demands.

A scalable network is one that can expand over time to address network growth and improve (or at least maintain) client response time. Server clustering, mentioned earlier as a technology to ensure availability, can also be used to address scalability issues by allowing you to add nodes to a cluster when your network encounters a period of growth.

Performance

Network performance—good or bad—is one of the most noticeable outcomes of any network design plan. Performance has a direct impact on all aspects of end-user productivity and customer satisfaction. If your e-commerce Web servers are overloaded, you will probably lose customers who abandon their shopping carts out of impatience. This translates directly into lost customers and lost income for your company. Likewise, providing adequate performance on a corporate LAN will allow your corporate employees to focus less on waiting for their workstations to reboot and more on productivity, thus creating revenue for the company.

Cost

There is an old joke among software developers that goes something like this: “Cheap, fast, right... pick two.” Monetary considerations can make or break a network design. An improperly budgeted network installation can create any number of long-term difficulties and end up costing even *more* money to correct problems that cropped up during the initial installation. Almost everyone embraces the goal of cutting costs, but remember that it is almost always less expensive to do something right the first time than it is to correct or upgrade an insufficient installation.

User Priorities

No matter what network infrastructure your organization uses, you can be certain that it won't be deployed in a vacuum. Whether your users are internal employees of your corporation or external customers paying for the services that your company provides, your network installation must provide for their needs if it is to be cost-effective and successful. You must create an environment that provides for the current needs of your users, as well as allowing room for future growth and changing requirements. We'll describe some of the more common network services in use today: e-mail and other communications, scheduling and task management, project collaboration, data storage and retrieval, Internet research, application services, print services, and graphics/video/audio services. (Of course, a complete list of network services is limited only by the imaginations of your customers, clients, and users.)

Electronic Communications

Electronic communication, specifically e-mail, has become the de facto means of communication in the modern business world. Whether a company manages its own e-mail storage, using a technology like Lotus Notes or Microsoft Exchange, or outsources its e-mail to an external Internet Service Provider (ISP), modern computer users have come to expect a great deal from their e-mail service in terms of performance and availability. The outage of an e-mail server is now perceived to be just as disruptive as the loss of telephone service. In designing e-mail services for your network, you should make allowances for high performance and availability to meet the expectations of your network users.

You can provide high availability and performance for your e-mail services by making sure that you've allocated enough server resources to support all of your current clients, as well as planning for the growth of your user base. As with most other network services, fault tolerance can be achieved through the use of redundant hardware within an individual server, like redundant power supplies and NICs, as well as RAID arrays for your hard drives. Also, you can use server clustering to create two or more physical e-mail servers that your clients will see as one logical server; if one physical node of the cluster fails, the other will take over, usually without your clients noticing more than a few seconds' outage.

Another common issue with e-mail servers relates more to *how* e-mail is used within your organization. Unsolicited commercial e-mail, commonly referred to as *spam*, can clog the inboxes of your client workstations, decreasing productivity as users sift through pages of junk mail looking for relevant messages. This can also lead to sexual harassment questions if the spam includes messages with adult content or graphics. As an e-mail administrator, you can implement spam-filtering centrally at the server level or install client-level tools for your users to configure according to their own tastes. Spam-filtering uses a number of different technologies, including blocking e-mail from lists of known spammers and filtering messages based on keywords such as "get rich."

Along with deciding how to address unsolicited e-mail, you should create a policy describing how e-mail and other computing resources can and cannot be used within your

environment. You'll often hear this referred to as an Acceptable Use Policy (AUP). An AUP essentially provides a road map for your users to make decisions about what is and is not appropriate to do with their office computers.

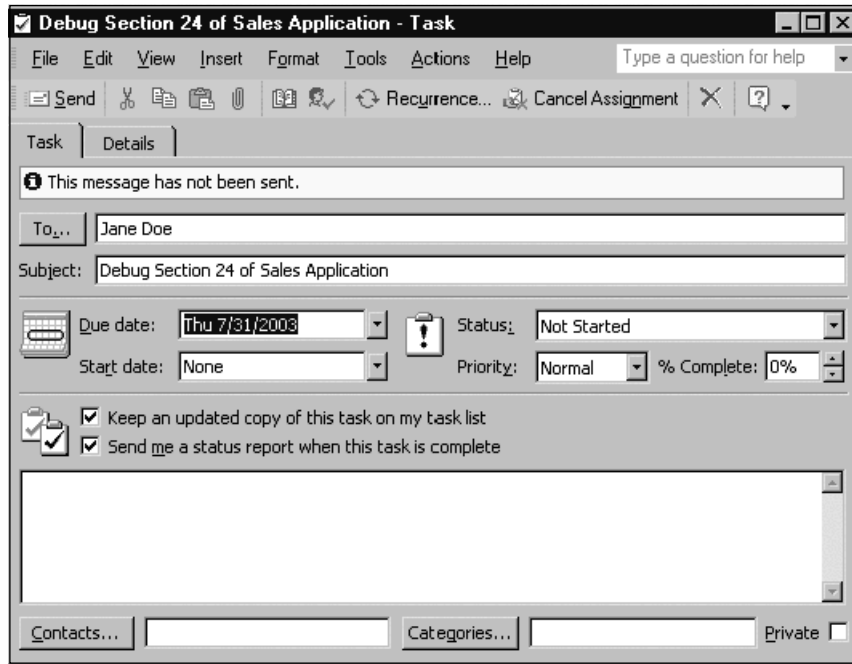
Some organizations have a strict zero-tolerance policy, where there can be no personal use of any company resource, including e-mail. More often, though, you'll see a phrase in an AUP that allows for "reasonable personal use" of computing resources. The purely financial argument might say, "If we can keep each of our 80,000 employees from spending one minute a day sending a personal e-mail, then we've saved the company X number of dollars." But at the same time, you need to consider the potential for *added* productivity for the account manager who is happier being able to send a quick note to his daughter who lives halfway across the country. You need to carefully consider what type of policy will best suit your organization.

Scheduling/Task Management

Fully featured e-mail clients such as Microsoft Outlook and Lotus Notes can extend e-mail functionality to include a wide range of calendaring and task-management functions. Users can manage appointments for anything from small project teams to entire departments and offices. This can improve the efficiency of users' time management by providing automatic meeting and resource scheduling, including notifications of appointments and time conflicts. Supervisors can manage schedules for an entire group of individuals, tracking meeting attendance, scheduled appointments, and vacations. Administrative assistants can even create, move, and delete appointments on their managers' behalf.

Centralized task management can also assist managers or team leaders in directing the projects under their supervision. Managers can assign specific tasks and track their progress and completion date from a single location. As you can see in Figure 1.11, you can keep copies of tasks you've assigned on your personal task list, as well as receive status reports when an assigned task has been marked as complete. When integrated with e-mail and calendaring functionality, task-management functions can greatly streamline work processes for project teams and departments of any size.

Figure 1.11 Assigning Tasks in Microsoft Outlook 2002



Along with using network resources to schedule and assign tasks for users and employees, you'll also want to allow for scheduling of computer-based tasks. This can include scheduling recurring events such as nightly backups of user data, or the ability to run tasks on an as-needed basis to create user accounts, reset a forgotten password, and the like. Windows Server 2003 has a graphical Task Scheduler interface that allows you to schedule tasks on a daily, weekly, or custom basis. You can also integrate many Windows commands and utilities into scripted batch files or custom applications. For example, the administrator for a university department might want to automate the process of creating user accounts for incoming freshmen every year, rather than spending time creating each individual account manually. Well-developed scheduling and task-management functions will allow the administrator to accomplish this in an efficient and timesaving manner.

Project Collaboration

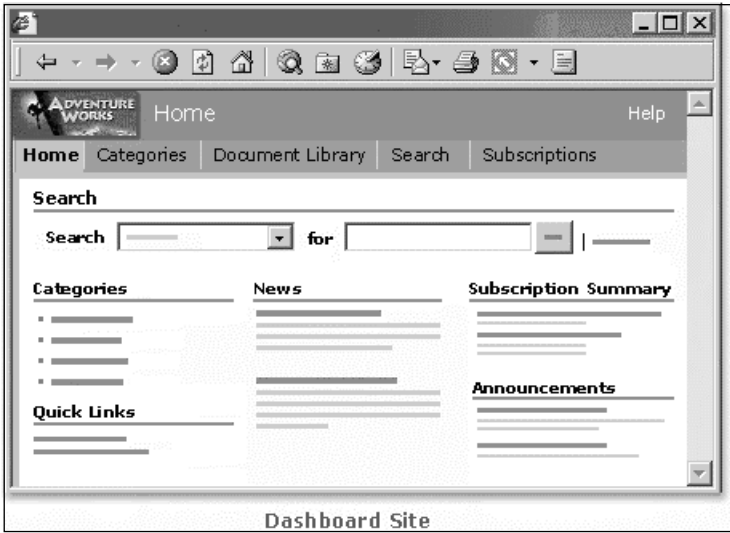
No matter the size of an organization, sharing information within an organization and with outside parties is vital to increasing productivity and creativity on projects of all kinds. In this case, a "project" can refer to any situation where people need to share information, from a formal business research project to a high school marching band. Project collaboration technologies must provide an intuitive and easy-to-use means of sharing documents,

deadlines, and other key pieces of information among people working from multiple locations.

Packages such as Microsoft SharePoint offer users the ability to organize and access information through a Web browser or another familiar Microsoft Office environment. Figure 1.12 (from the SharePoint homepage on www.microsoft.com) illustrates the kind of information that a project collaboration technology can gather at a user's fingertips.

Microsoft SharePoint comes in two varieties: SharePoint Team Services, and the more

Figure 1.12 A Microsoft SharePoint Project Collaboration Web Page



full-featured SharePoint Portal Server. SharePoint Team Services is actually integrated directly into the Windows Server 2003 operating system, and provides the ability for small or ad hoc project teams to share information. The full-blown SharePoint Portal Server is designed to work in an enterprise installation, allowing users to share and manage documents among multiple servers. The key differences between the two versions of SharePoint are listed in Table 1.1.

Table 1.1 Comparing SharePoint Portal Server and SharePoint Team Services

Feature	Team Services	Portal Server
Core function	Ad hoc team collaboration	Enterprise portal and search
Search capabilities	Documents within team Web site and subsites	Across multiple servers and data types
Discussion and notifications	Discussions, notifications, and surveys	Discussions and notifications

Continued

Table 1.1 Comparing SharePoint Portal Server and SharePoint Team Services

Feature	Team Services	Portal Server
Customization	Browser-based, Microsoft FrontPage 2002, and SDK	Web Parts and SDK
Document management options	Publishing	Check-in and check-out, versioning, routing, and publishing
Client applications	Browser, Microsoft Office XP, and FrontPage 2002	Browser, Microsoft Windows Explorer, Office 2000, and Office XP
Security options	Customizable roles: Administrator, Advanced Author, Author, Contributor, and Browser	Administrator, Coordinator, Author, and Reader roles
Licensing requirements	One FrontPage 2002 server license, no separate client access license (CAL)	Server license and CALs

Data Storage and Retrieval

Providing a central location for users to store and access files is one of the oldest and most common uses for a network file server. This provides your users with the ability to access shared data within a department, an organization, or an enterprise. The Windows operating system has provided the means to share files and folders since the release of Windows 95. The Windows Server operating systems allow an administrator to add management, security, and scalability functions to their users' ability to share information. When planning file services for your network, you should keep the following objectives in mind:

- Simplify user access to files in a large organization, especially when those resources are located on multiple servers and shares. This can include the ability to retrieve data stored on multiple servers from a single access point.
- Provide efficient data access for users accessing information from multiple locations. For example, if a sales manager in Chicago needs frequent access to reporting data from remote servers, he should be able to access that data without using an expensive leased line to do so.
- You should be able to migrate data to various servers without affecting the way that users access that data. If you must visit each user's workstation whenever you reconfigure a share or a server, it will greatly restrict the flexibility of your network infrastructure.
- Minimize any delays that can occur when accessing a frequently used file or folder.

Windows Server 2003 has introduced new features (and improved on existing Windows Server functions) to improve file sharing services, including the following:

- **Volume Shadow Copy** This allows network backups to take place while users are still accessing files and folders, increasing the availability of shared documents on the network.
- **Distributed File Service (DFS)** Like its predecessor in Windows 2000, DFS allows you to take shared folders located on multiple physical servers and group them using a single namespace. With this feature, you can add or remove physical folders, drives, and even entire servers without affecting how your users access the resources they need. DFS can also be used to provide fault tolerance and load balancing for the file sharing services on your network.
- **NTFS permissions** As in previous versions of Windows, file permissions prevent unauthorized access to the resources on your network. Windows Server 2003 also has continued support for file compression to save space used by infrequently accessed files on your hard drives.
- **Disk quotas** As with Windows 2000, you can use the disk quota function of Windows Server 2003 to passively monitor or actively control disk usage on your file servers. Disk quotas can be enabled on a per-user basis on any of your server volumes. Properly implemented disk quotas will increase the availability of your file sharing services by preventing drive space from filling up without warning.
- **Removable storage** Windows Server 2003 provides enhanced support for removable storage devices such as Zip drives, FireWire devices, and Universal Serial Bus (USB) storage devices.
- **Offline files** Like Windows 2000, Windows Server 2003 will allow users to “check out” a network file and make changes to it on their local machine before the file is checked back into the network storage location. You can use this to improve performance, especially when accessing files over a WAN link or when you’re dealing with remote and traveling users who may need to work on network files while they are disconnected from the network.
- **Encrypted File System (EFS)** This feature uses Public Key Infrastructure (PKI) certificates to digitally encrypt user files stored on a server or a local hard drive. This feature is largely unchanged from Windows 2000. It relies on users’ private keys to provide encryption for their stored files.
- **Indexing Service** Another feature found in previous versions of Windows, the Indexing Service in Windows Server 2003 creates indexes of the contents of a server or workstation hard drive, as well as indexing the properties, or *metadata*, for various document files. This allows you to index files not only by name and location, but also by such properties as author, category, timestamp, and so on. You

can create multiple indexes on a single machine to exert granular control over how the Indexing Service operates.

Internet Research

The Internet and the World Wide Web have created instant access to a wealth of information on countless topics for both personal and business use. This instant access to information has become crucial to the modern workplace, allowing access to a wide variety of resources (some of which we've already discussed), including e-mail, file transfers, business and personal collaboration, access to multimedia information, and more. The various resources available on the Internet allow users to research vast amounts of material and information.

Internet research differs greatly from traditional "paper" library research because information is not centrally catalogued in a single location, and it can move and change from day to day and week to week. Addresses of Internet sites can change and sometimes disappear altogether, creating a fluid and somewhat volatile environment. Information found on the Internet can also vary widely in terms of accuracy, credibility, and attention to detail, making it crucial to evaluate not only the information, but also the source of that information.

When designing a network for any setting, whether for corporate, educational, or personal use, it's almost a given that you will be making some allowance for access to the Internet and the World Wide Web. Whether this access is universal to all users or restricted to only those who need to perform Internet research as part of their job functions, a good network design will provide secure access that will permit access to necessary resources while protecting the security of the internal network resources. You can accomplish this through the use of firewalls, proxy servers, and other hardware and software-based technologies.

Application Services

A well-designed network can allow you to host client applications from a central location, thus reducing deployment time and management costs as well as providing for centralized security. Centralized application management addresses some of the following user needs:

- Central storage of application data so that users can access needed files from anywhere on the network
- Centralized deployment, upgrading, and patching of applications without requiring user intervention (or sometimes even user knowledge)
- Enabling offline access to network applications so that users can perform their tasks while disconnected from the network

Using a central application server such as Windows Terminal Services can enable you to deploy an application one time only to the server itself, rather than installing it on each

user's desktop. This can greatly improve both user and administrator efficiency in the case of custom applications that require frequent updates or applications that need to be deployed to users in geographically remote locations. Centralized application hosting can also increase security by maintaining sensitive data in a centralized location, rather than allowing it to traverse insecure network connections.



NOTE

When considering deployment costs of Terminal Services, remember to take into account Terminal Services licensing fees. Each client must have not only a CAL for the client operating system it is running, but also a Terminal Services license. See www.microsoft.com/windowsserver2003/techinfo/overview/termservlic.msp for a white paper that discusses all the intricacies of the Terminal Services licensing structure in Windows Server 2003.

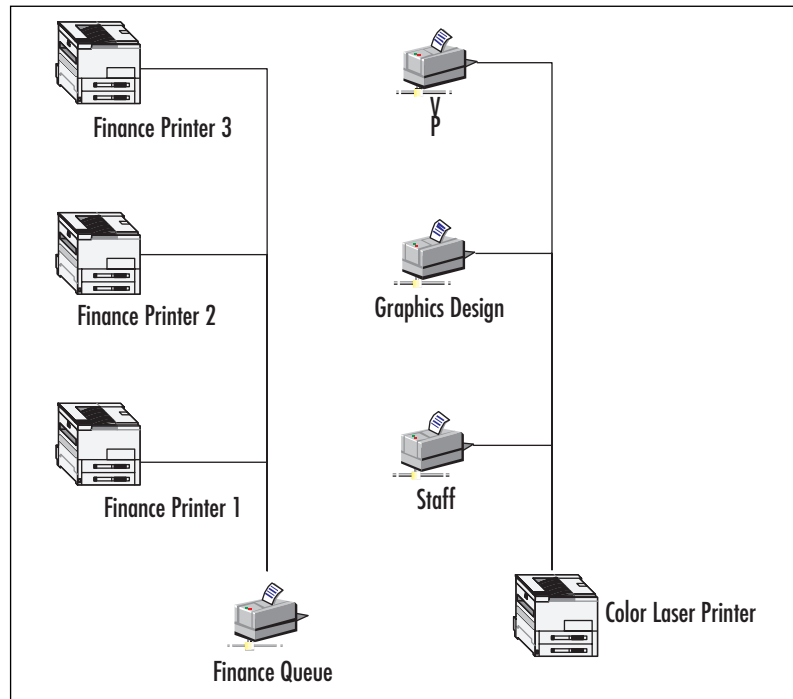
Print Services

Almost every environment relies on highly available printing services to produce all forms of paper output. A properly designed network will create shared printing resources across the network, allowing workstations to submit print jobs to printers that are attached to local servers or that are accessed across the Internet. Network operating systems such as Windows Server 2003 allow you to cluster network printers for high availability, and to automatically deploy printer drivers to clients of many different operating systems. Well-designed print services will also enable users to easily locate the printers they require. Administrators should be able to centrally manage and configure printers from any location.

The Windows printing architecture consists of two components:

- **Physical printer** The *printer* is exactly what you think it is: the physical print device that is attached to a workstation or server's parallel or USB port, or plugged directly into the network.
- **Logical print queue** The *print queue* is the software piece that translates between the physical printer and the software application from which that the user is printing.

To improve printing efficiency, you can have a single print queue submit jobs to multiple printers, referred to as a *printer pool*. In the example shown in Figure 1.13, the Finance queue feeds to three separate printers. This is useful if a department produces a large amount of paper output, since you can manage the three physical printers as a single logical unit.

Figure 1.13 Printer Pools and Prioritized Queues**NOTE**

You should implement a printer pool only if the printers themselves are physically close to one another; otherwise, your users will be running from printer to printer looking for their output.

On the opposite end of the spectrum, you can have multiple logical print queues feed to a single physical printer in order to prioritize your users' print jobs. You can assign a priority to a print queue between 1 and 99. Print jobs from higher-priority print queues will be processed before jobs submitted from lower-priority print queues.

You can also establish schedules in which printing to a certain queue may not be available at *all*. In Figure 1.13, there are three print queues set up for a single color laser printer. Let's say that you want your graphics designers to have first priority when printing to this device, followed by any of your vice presidents. You can assign a priority of 99 to the Graphics Design print queue and a priority of 1 to the VP queue. Furthermore, you've recently discovered that some staff members have stayed after business hours to print personal material to the color laser printer. In order to keep from wasting the expensive color laser toner, you can establish a third Staff queue that can only be printed to between the hours of 9:00 A.M. to 5:00 P.M., Monday through Friday.

Graphics/Audio/Video Services

The increasing prevalence of high-speed Internet connectivity has created a market for high-quality streaming media services, ranging from streaming audio services offered by online radio stations to full-fledged audio/video streams used for training and conferences. Windows Server 2003 includes the latest version of Windows Media Services, allowing companies of any size to create and host powerful streaming media capabilities.

As a network manager, you need to be aware of the hardware, software, and network bandwidth considerations created by your organization's current and potential future use of streaming media capabilities. Planning for the requirements of this technology is essential in creating an efficient network design.

Reviewing Legal and Regulatory Considerations

Depending on the business in which you are involved, your network design plan should address the legal issues associated with your industry, geographic location, and so on. Backup schedules and offsite data availability have become federally regulated matters, especially in the financial arena. Consult your Legal department during the design process, because like everything else in this venture, it's certainly best to get it right the first time.

Don't forget to include your client workstations when making allowances for legal and regulatory matters. For example, if your corporate data-retention policy calls for maintaining e-mail data for twelve months, but some users have copies of every item they've sent or received in the last five years, that fact could come back to haunt you in a legal proceeding.

Some fields of business are subject to very detailed governmental regulations regarding data security. For example, healthcare providers now fall under strict laws regarding electronic patient information since the Health Insurance Portability and Accountability Act (HIPAA) went into effect in 2003. Regardless of your field, if you work on government projects, your network might be required to meet specified security criteria.

Network communications can also subject your company to legal liability when employees misuse the network. For example, pornographic material on the company network can subject the company to charges of the "hostile workplace" definition of sexual harassment under Title VII of the federal Civil Rights Act of 1964 and various state laws. You should also consider intellectual property (copyright, trademark, and patent) laws in establishing your network policies.

Common factors that also need to be reviewed for legal compliance are any Service Level Agreements (SLAs) in place on your network. An SLA attempts to define the scope of a service provider's responsibilities in maintaining applications or services on a network. This provider can be an external vendor to whom you've outsourced a critical service (your ISP, for example), or the SLA can be an internal document detailing the IT department's duties in maintaining network availability. The following are the major components of an external SLA, using an ISP as a real-world example:

- **Scope of services** This spells out exactly which service or application that an SLA is referring to and the level of responsibility that the internal IT department will have in maintaining this service versus the external vendor. This includes outlining the hardware, software, and resources that comprise the particular service, such as the modems, network connectivity equipment, ISP help desk, and engineering personnel in the case of an ISP.
- **Roles and responsibilities** Your ISP should establish a coverage schedule so that at least one primary and one backup support avenue is available to report any service outages. You'll also need to establish a system to escalate support calls if the scheduled support person is unavailable or cannot correct the problem. You can use this information to inform your users of the turnaround time they can anticipate in responding to and resolving any problems.

These are only a few of the legal considerations that are important in a corporate network environment. You should always include a legal advisor as a member of your network planning team.

Calculating TCO

“These upgrade proposals look interesting, but how will they impact our company’s TCO?” TCO is a calculation that was designed to assist consumers and corporate managers in assessing the direct and indirect costs and benefits associated with the implementation of new or upgraded computer technology. The purpose of TCO is to quantify the financial bottom line associated with a computer or technology purchase decision.

TCO calculations do not rely on a single formula. For example, a high-end computer will have a higher initial purchase price, but will probably incur fewer repair bills during its active life cycle. TCO is balanced against the benefits created by the technology purchase, such as improved user efficiency or perceived happiness with improved performance, in attempting to make a final purchase decision.

The first part of calculating TCO is relatively simple: What is the initial purchase price of the new technology? Include the cost of hardware, software licensing, networking equipment, installation charges, and so on. Don't forget to factor in the necessary time to train your end users and IT staff in the use and administration of the new technology. Next, determine the ongoing costs for maintenance and support. These costs can include charges for vendor support, as well as in-house labor expended on interoperability issues with third-party and legacy software support. Try to estimate the total costs for the full anticipated life cycle of the proposed technology.

Determining the soft costs associated with a new technology is a bit more complicated. How much money will your company save by reducing the number of times your users are forced to reboot their computers each day? Conversely, how much money is lost when an account manager cannot access the order-entry application for 20 minutes, for an hour, and for a day? These costs are fairly difficult to quantify, but they can be critical when deter-

mining the total benefits afforded by a network upgrade. You can start investigating soft costs by talking to your users and reviewing TCO models from network analysts.

Your users can certainly tell you how much it aggravates them when their e-mail or order database is “running too slowly,” even if they can’t tell you what “too slowly” means in terms of actual response time. This can also point out performance bottlenecks that you may not have known about before. For example, a real estate lending office for a well-known bank shared a T1 line with the bank branch in the lobby of the office building. The real estate lenders encountered severe network performance degradation every day at around 4:30 P.M. Further investigation revealed that this time frame coincided with the bank tellers transmitting their daily totals to the bank’s main headquarters when the branch closed each day.

Preconfigured TCO models from organizations like the Gartner Group, IDC, or other independent network analysts can walk you step-by-step through plugging in various budget figures to arrive at the TCO of a specific technology, hardware, or software package. However, remember that these models are not set in stone, and they should be modified as needed to meet the specific needs of your organization. These models will rely more on actual calculations, such as dividing a help desk analyst’s salary by the number of support calls he or she is able to process in a day, or determining the “cost per e-mail message” of an e-mail server upgrade that increases the number of messages it can transmit in a day, week, or hour. You can then take these numbers and factor in the soft costs already mentioned. Using a combination of calculations and judgment calls will typically lead you to the most accurate assessment of TCO within your organization.

Planning for Growth

If there is one nearly universal truth to network design, it is that networks and their resource requirements always eventually grow. Your network design needs to account for not only what your users require today, but also what they are likely to require in the future. Even if your users or clients have not thought about future growth, you should provision your network design to accommodate for a reasonable increase in user population and bandwidth usage as time goes on.

One of the best ways to ensure that your design will support the future needs of your network is to implement well-known, standards-based technologies, rather than those that are proprietary or experimental. Expanding your network’s router core, for example, will be much simpler if the new hardware you purchase is compatible with the initial installation. (Otherwise, you might need to scrap the initial installation entirely and install all new hardware, greatly increasing your costs and overall headaches.) You should also deploy hardware and software in as consistent and well-documented a manner as possible, so that you can perform maintenance and upgrades as quickly as possible.

Examine the feasibility of allocating items like high-capacity network cabling and other infrastructure components at the initial installation of your network. For example, it may cost an extra 25 cents per foot to run 100MB Ethernet cable instead of 10MB Ethernet

cable when you're initially wiring your building, but it will cost significantly more if you find you need to rip out the cabling and redo it later.

In planning for network growth, you should again consult with your users, especially those in strategic planning and decision-making capacities. Although no one can accurately predict what will or will not happen to a company over months and years, these decision-makers will be able to give you some idea of the overall vision of the company. Are they hoping to expand dramatically through mergers and acquisitions? Or are they satisfied with their specific market niche and anticipate adding personnel and equipment in only smaller increments as production increases?

Finally, when considering desktop computers, laptops, and servers, keep in mind that most current hardware will come with a one- to three-year warranty, sometimes with an option to purchase an extended warranty at the time that you buy the equipment. It's not necessarily true that your computer hardware will immediately break down the day after the warranty expires; however, the length of your warranty and/or service contract *should* factor into your projections regarding how often you plan to replace your equipment. For this reason, many organizations adopt a three- to four-year replacement cycle, budgeting sufficient funds to replace one-third or one-quarter of the installed computer base every year, or setting aside money to replace all of the equipment *en masse* when it reaches the end of its warranty cycle.

Developing a Test Network Environment

When implementing a new network or computer solution, you should perform a thorough battery of testing before deploying it into production. You'll begin the test process in an isolated lab where new technologies will have no chance of adversely affecting the existing computing environment. After you are satisfied with the new technology's performance in the test lab, you can expand testing into a pilot deployment involving a few actual users, analyzing their input and reactions to make any necessary adjustments to your design. Only after you are satisfied with the pilot deployment should you perform a full-scale deployment in your production environment.



TEST DAY TIP

Depending on the total number of users you have, you might want to split your full-scale deployment schedule into stages. After each stage, you can verify that your system is accommodating the increased processing load from the additional users as expected before you begin deploying the next group of users.

The success of any network deployment depends heavily on your ability to develop an effective test environment. This test lab can consist of a single lab or several labs, each of which can test various pieces of the overall design without risking the integrity of your production environment. Working in the test lab will allow you to verify the effectiveness

of your design, discover any potential deployment problems, and increase your staff's familiarity with the new technology before it "goes live." In short, a well-developed test environment will reduce the risk of errors during the deployment of a new technology, thus minimizing any potential downtime for your clients and users.

Planning the Test Network

Before you begin testing your network design, you need to plan the test network itself. The first step is to determine the hardware resources required to set up the lab. This involves identifying the standard configurations of your existing or new client computers. (If you support diverse workstations, do your best to include a representative workstation from each supported configuration.) Be sure to include all components and peripherals, including the following:

- BIOS versions
- USB adapters
- CD and DVD drives
- Sound cards
- Video cards
- Network adapters
- Smart card readers
- Removable storage devices, such as Zip drives or external hard drives
- Small Computer System Interface (SCSI) adapters
- Removable storage devices
- Mouse or trackball devices
- Keyboards

Although using separate hardware devices for your test lab is the ideal, many small and medium-sized businesses simply cannot afford to buy dozens of computers for the test lab. Using a third-party product such as VMware (www.vmware.com) will allow you to simulate a multiple server/domain environment, as well as multiple desktop operations systems, without the expense of multiple individual machines. VMware can run multiple operating systems—such as Microsoft Windows, Linux, and Novell NetWare—simultaneously on a single PC, including all networking and connectivity that you would need to perform your testing.

In addition to purchasing hardware or virtual PC environments for the test lab, you need to secure appropriate licensing for all necessary software, including operating systems, service packs, management utilities, and business applications. Make sure that you can obtain or duplicate the following configuration and information when creating a test lab for Windows Server 2003:

- **Network services** Install the same services on a test server that will be used in the actual deployment. This can include Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), Windows Internet Name Service (WINS), or any other Windows service.
- **User accounts** Create a domain controller in your test environment to effectively simulate any upgrade procedures.



TEST DAY TIP

You can use the Clone Principal tool (Clonepr.dll) utility, included in the Windows Server 2003 Resource Kit, to copy production users into a test domain.

- **Domain structure** Simulate the domain hierarchy of your proposed environment, including forests, trees, parent and child domains, and all necessary trust relationships. Configure sites as necessary to simulate any WAN testing considerations.
- **Network protocols and topology** Re-create the network technologies that will be used in your production environment as completely as possible. For example, if your production environment will be using 100MB cabling, using Gigabit Ethernet when doing performance testing will provide erroneous results. You should also include routers to test for performance latency as well as replication across WAN links.
- **Domain authentication** Use the appropriate authentication to mimic the desired production environment, including mixed mode versus native mode, and NTLM versus Kerberos client authentication. Selecting the appropriate authentication model will allow you to compare apples to apples during testing and avoid any unexpected behavior later.



EXAM WARNING

Remember that Windows NT 4.0 workstations or servers cannot use Kerberos authentication. You will need to rely on either NTLM authentication or its stronger successor, NTLM version 2.

- **Group Policy Object (GPO) settings** Create GPOs with the settings that you wish to deploy in your production environment. You can use the GPMC (discussed earlier) to test the potential behavior of any policy objects on user and group objects.

Test Lab Domain Structure

Although you usually want your test lab to mimic your production environment as closely as possible, there are exceptions to every rule. Some tests that you might wish to perform will affect an entire domain or forest, rather than a single machine. If you are testing this type of functionality, you might wish to create a separate domain within the test lab so that the remainder of the lab environment will not be adversely affected.

Some of the tests for which you might wish to create a separate, isolated domain or forest are as follows:

- **Switching from mixed mode to native mode** Changing from mixed mode to native mode will allow for much tighter security in a Windows 2000 or Windows Server 2003 environment, but it assumes that you have no Windows NT 4.0 backup domain controllers (BDCs) remaining in your domain. (After the switch to native mode, Windows NT 4.0 BDCs will no longer be able to replicate with Windows 2000 or Windows Server 2003 domain controllers.) This change will affect an entire domain and cannot be reversed.
- **Upgrading the domain or forest functional level** This feature was introduced in Windows 2000, where you had the ability to run a domain in mixed mode for backward compatibility or native mode for increased security and functionality. Windows Server 2003 expands on this by creating several levels of both forest *and* domain functionality that can expose different features of the operating system for your use. For example, raising the functional level of a domain to Windows Server 2003 native will prevent any existing Windows NT 4.0 or Windows 2000 Server domain controllers from participating in domain replication. Like the switch from mixed to native mode, this will affect the entire domain and/or forest in question and cannot be undone.
- **DNS settings** Changes to a DNS server will affect all clients who use that server for name resolution. Although this does not involve the kinds of one-way changes described above, you should still proceed with caution before making changes that can affect other tests that might be running simultaneously in the lab environment.

One important (but often overlooked) step in the planning process is that of carefully selecting a location for your test lab. Too often, the test lab is relegated to a corner of a server room or whatever room is available in a file or storage area. However, if you will be performing tests for an extended period of time, you should consider allocating a permanent or semipermanent location for the lab. Be sure to locate the test lab in an area with enough space for all necessary equipment and personnel. If you will be testing network equipment that will be deployed to multiple locations, you should consider deploying a test

lab at each site to test WAN links, replication, and site configurations. Also, identify the personnel you'll need to perform testing, as well as whatever training they will need.

Finally, be sure to provide both physical and technological security measures for the equipment and resources of the test lab. This includes isolating the test lab topology from your corporate network using routers, switches, or firewalls, as appropriate. If you need to provide a connection from the test lab to the corporate network, decide in advance how you will control and monitor that connection, and be sure to devise a way to quickly terminate the connection if something unexpected or adverse occurs.

Building a Test Lab

How you create your test lab depends on your specific requirements. Here, we present the basic steps for building a test lab, which you can alter as necessary to meet the needs of your organization:

1. Begin by acquiring the necessary hardware and software, including the following:
 - Routers, switches, cabling, and other network infrastructure devices
 - Computer hardware for servers and workstations
 - Operating system software and any administrative tools
 - Line-of-business software applications
2. Install and configure the necessary routers and switches to provide network connectivity for the test lab. Label all devices and network cables.
3. Install and configure server hardware. Try to use the same random access memory (RAM) and central processing unit (CPU) configuration that you plan to deploy in your production environment. Configure the hard drive arrays, partitions, and drive letters to match the intended production environment.
4. Defragment all hard drives and install up-to-date antivirus software.
5. Install the appropriate operating system for the test environment.
6. If you will be deploying new Windows Server 2003 servers, perform a clean installation of Windows Server 2003.
7. If you will be upgrading an existing server, install a copy of your existing network operating system (NOS) to test the upgrade process.
8. If you will be repeating the test process several times, consider using a disk-imaging utility to save time when re-creating the test environment.
9. Test the network connectivity in the lab environment. Testing the network connectivity first allows you to isolate any problems more easily.

Continued

10. Install all application software that will be present in the production environment. Include all server-based applications and administrative tools such as SMS. If you are using Terminal Services in your production environment, install all applications that will be present in the production environment.
11. Install and configure all client computers.
12. Secure the test lab using physical measures, such as a card-reader on the entrance door, and technical measures, such as a dual-homed router to segregate network traffic originating in the test lab from traffic passing on the production network.

Implementing the Test Network

After you've finished designing your test lab, you can finally get down to the actual business of testing. The steps needed to create test procedures can be broken down into two conceptual halves: *What* do we want to test and *how* should the tests be performed? You'll often hear the former referred to as a *feature test description*, which lists all features or aspects of a technology that need to be tested. For example, the feature test description when assessing how trust relationships behave during an operating system upgrade might read something like this:

“All trust relationships between the Windows NT 4.0 domain PRODUCTION and the Windows NT 4.0 domain SALES should continue to function normally when the SALES domain is upgraded to Windows Server 2003.”

You should design tests that will measure the functionality of each feature included in your design plan. Additionally, you need to test how your new network will function in conjunction with any existing systems in the production environment. For Windows systems, you need to test hardware, driver, and application compatibility on every hardware configuration that will be running a Windows operating system.



TEST DAY TIP

Be sure to test functionality with existing technologies even if they are going to be upgraded or replaced as part of the new installation. Although the new technology may need to coexist with the old technology for only a short period of time, you need to know in advance how the interoperability will behave.

Another key factor in using a test lab is creating a schedule of when testing should be performed. Especially if you have many different individuals or teams performing various tests, this scheduling should be formalized, rather than handled on an ad hoc basis. Ideally, you should designate someone to act as a lab manager to maintain and upgrade the lab

schedule, as well as review testing plans to ensure that all necessary equipment and software can be made available at the requested time. Even if the lab manager is not dedicated solely to this function, he or she should be responsible for the following:

- Establish and enforce test lab policies, procedures, budget, and inventory control.
- Oversee scheduling of required configuration changes and communicate these changes to other test lab users.
- Develop and manage an incident reporting and tracking system.
- Monitor the change-control process.
- Maintain test lab documentation (we'll discuss documentation in the next section).
- Manage hardware and software configurations, updates, and preventative maintenance.
- Establish and maintain physical security.

New & Noteworthy...

Exploring the Group Policy Management Console (GPMC)

A prominent new feature of Windows Server 2003 is the GPMC, which allows administrators to monitor, troubleshoot, and plan Group Policy settings across an entire enterprise from a single management console. Along with a console window that provides a graphical representation of GPO settings, the GPMC also includes a collection of scripts that you can run from the command line to streamline administration and planning tasks. You can download and install the GPMC from Microsoft's Web site. Once it's installed, you'll have a shortcut to it in the Administrative Tools folder, and it will be available as an MMC snap-in.

The scripts that are included with GPMC can greatly simplify your life when you attempt to take stock of an existing network environment (for example, when you begin to plan for an upgrade). Using GPMC, you can quickly perform the following tasks using its automated scripting function:

- List all GPOs that are present in a given domain
- List any disabled GPOs
- List GPOs at a backup location
- List GPOs by policy extension or security group
- List any orphaned GPOs (GPOs that are no longer linked to any AD object) that are still present in the SYSVOL directory
- List unlinked GPOs in a domain

Continued

- List GPOs with duplicate names
- List GPOs without security filtering

GPMC's reporting functions will also generate HTML-formatted reports in an easy-to-read format, which is always a hit when you're presenting the upgrade proposal to management or a budget committee. Additionally, the GPMC includes the Resultant Set of Policy Planning function to allow you to simulate changes to GPO settings for a user, computer, or container object. Both of these functions will greatly assist you with the administrative and technical aspects of a network design project.

Documenting the Planning and Network Design Process

After you've determined *what* needs to be tested to ensure that your network design is working correctly, you need to create detailed descriptions of *how* each test should be performed. This is crucial in order to ensure that all necessary functions have been properly tested. Documenting the test process becomes even more vital when multiple individuals or teams are using the test lab resources, because you need to keep track of how one test may have affected others being run concurrently or later. For each test that you wish to perform in the lab, be sure to identify the following:

- The prerequisites for the test to function: how to prepare the lab for the specific test
- The specific action or change that you will be testing
- The individual steps required to implement the installation, change, or troubleshooting step
- The expected result of the test
- What rollback actions to take (if any) if the test fails
- What subsequent actions to take if the test succeeds

You should document the layout and initial configuration of the test lab itself, using both text documents and diagrams where applicable. Both of these, but especially the diagrams, should be posted in a prominent location so that users of the lab will be aware of any design changes that you have implemented. This information will also improve the efficiency of the test process itself, since the testers will know where to locate each component or server to which they require access.

Finally, remember to periodically test your lab equipment to determine what effects testing has had on it. A computer that has undergone many changes and upgrades throughout the course of the testing process will certainly behave differently than a com-

puter that has been newly installed, even with an identical configuration. You should refresh the disk images on your lab machines periodically using disk-imaging software such as Symantec Ghost or Microsoft's Remote Installation Services (RIS) to be certain that your test machines are offering a fair representation of the systems that you are attempting to analyze.

Importance of Documentation

The importance of documenting your computing environment after you have deployed a new network design such as Windows Server 2003 cannot be overemphasized. As you move through the network design and testing processes, you should also keep detailed documentation of each design, product, or vendor decision that you make, including your reasons for choosing one alternative over another. Personnel changes can occur without warning, and a well-maintained design document will quickly answer the question of "Why did we choose Vendor X over Vendor Y?" when it is posed by the new Vice President of IT who just started last week. Knowing that Vendor Y's product proved incompatible after several hours of troubleshooting will save you from needing to waste time by repeating portions of the design process.

Because of the effects that ongoing changes can have in a production environment, many organizations use test equipment to test every patch and service pack that is released by their product vendors, so that any potential problems or bugs can be intercepted before the patch is applied globally. Whatever method you use to roll out ongoing updates and changes, you should include detailed documentation, not only of *what* update was rolled out on a given date, but also of *how* the change was applied to client machines or other devices on your network.

Creating the Planning and Design Document

When documenting both your test lab and your overall network design, there are a number of items that need to be discussed. Although maintaining network documentation is often relegated to a backseat behind the numerous fires that we must put out on a daily basis as network administrators, comprehensive records in this area will actually help you in whatever troubleshooting issues come up after the new network is placed into production. Include configuration information about the following components of your final network design (although a complete list is limited only by the amount of time you have in the day!):

- Windows Server 2003 domain structure information, including DNS hierarchy and replication information, AD hierarchy information (site configuration, forest, domains, and OUs), and GPO settings and where they are applied within the AD hierarchy

**EXAM WARNING**

Be sure to include information about **Enforce** and **Block Inheritance** flags in Group Policy implementation. These affect how GPOs are inherited throughout the AD infrastructure.

- Trust relationships, both transitive and explicitly defined
- Network connectivity hardware (switches, routers, firewalls, and other LAN and WAN connectivity devices)
- Client computer configuration, both hardware and software
- Line-of-business application inventory and configuration
- Backup, restore, and disaster recovery procedures

Summary of Exam Objectives

The 70-293 MCSE exam measures skills related to the planning and maintenance of a Windows Server 2003 infrastructure. This exam covers tasks relating to all aspects of network design and planning, including making provisions for network security, performance, and availability. This chapter has introduced you to these topics; subsequent chapters will examine the tasks introduced here in far greater detail. The upcoming chapters in this guide will take you through all the necessary steps to prepare for the 70-293 exam.

The first skill set measured by this exam involves the ability to plan roles for installed servers in your network. We'll discuss how to evaluate existing technologies and hardware to select the appropriate function for each machine in your network, including Web servers, database servers, and domain controllers. You'll also learn how to plan and configure your physical network infrastructure, including TCP/IP addressing schemes, traffic monitoring, and planning for Internet connectivity.

Windows Server 2003 includes features to provide fault tolerance and increased availability for your network environment. Network Load Balancing and server clustering will enable you to configure logical groups of servers that will function as a single entity, allowing you to continue providing network services to your users and clients in the event of a hardware or another type of system failure. We'll also examine the steps needed to set up an effective security infrastructure, including the use of Internet Protocol Security (IPSec) and PKI.

To begin our look at the exam objectives here, we started with an overview of the network design process. As you can tell, this process is as much interpersonal as it is technical; in order to develop a useful network, you need to understand what your users, clients, and their managers are expecting the network to do in the first place. Before you can get down to the specifics of choosing server operating systems, software, and hardware, it's critical to develop a high-level perspective on your organization's overall makeup, managerial structure, and business requirements. This can include specific functions like e-mail, Internet availability, and printer sharing, along with overall organizational requirements like fault tolerance, growth capacity, and information security. You'll use this information to design a network that will meet the needs of all members of the organization and make their work as smooth and efficient as possible.

Once you've developed a design that you're satisfied with, you should test the design plan rather than immediately implementing it in a production environment. This will allow you to work out any quirks in the design or to spot something that doesn't work in reality quite as well as it seemed to on paper. We covered the various options available in creating a test lab, including using temporary equipment, creating a permanent site for testing, and using third-party tools to simulate multiple operating systems when time, space, or money are too tight for a full-blown test lab. We also talked about design considerations for the test lab itself and how best to secure the test lab so that any changes you make there won't affect your production equipment and environment. Finally, we looked at the importance of network documentation, during the planning stages as well as throughout the life of your network.

Exam Objectives Fast Track

Overview of Network Infrastructure Planning

- ☑ Proper planning of a new or upgraded network infrastructure must provide high performance, availability, and fault tolerance, as well as security and user satisfaction in as economically efficient a manner as possible.
- ☑ Creating a workable network design requires an understanding of both the business requirements of your organization and the new and existing technologies that can help to fulfill those goals.
- ☑ An effective network plan balances overall network concerns such as security and fault tolerance with the ability to provide specific applications and services that will improve the efficiency of your users' daily lives.

Analyzing Organizational Needs

- ☑ Identifying management priorities such as security, fault tolerance, and capacity for growth will help you in addressing high-level network considerations to adhere to management priorities that apply to an entire enterprise.
- ☑ Information flow involves an understanding of where your users are located in relation to the data that they need in order to perform their job functions and designing your network to ease their access to that information.
- ☑ Total cost of ownership (TCO) is a useful but elusive figure that includes the actual cost of purchasing new equipment, combined with less tangible items such as money saved by increasing user efficiency, improving customer accessibility, and reducing downtime.

Developing a Test Network Environment

- ☑ Your test network should mimic as closely as possible your existing or proposed production environment. This includes client and server hardware configurations, network services, and network connectivity hardware such as routers, hubs, and cabling.
- ☑ When performing testing, you should try to isolate the test environment from your production equipment so that nothing that occurs during the testing process will create downtime or unexpected results for the users and computers on your working network.

- ☑ Create well-defined procedures for each test you wish to perform, including detailed instructions for how to perform each test, the expected results of each test, and what steps you should take to recover the test network if the testing fails.

Documenting the Planning and Network Design Process

- ☑ Comprehensive documentation of the network design and testing process, while it involves an upfront time commitment, can save an administrator or help desk countless hours of frustration during later upgrades or troubleshooting.
- ☑ Document the initial configuration of all hardware and software that is implemented in a new network, and keep detailed records of all configuration changes, patches, and updates that are performed throughout the lifetime of the equipment.
- ☑ Network documentation should include a diagram and inventory of all network equipment, along with procedural instructions for backup and recovery procedures, security policies, any internal or external Service Level Agreements (SLAs), and any other legal or regulatory documentation.

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 support a very small network with only a handful of PCs and a single server. Do I still need to set up a test lab to upgrade from Windows NT to Windows Server 2003?
- A:** Yes, even small environments will benefit greatly from testing any new technologies before they are implemented into the production environment. In some respects, thorough testing is even *more* critical for a small network. Consider that if one node in your four-way SQL Server cluster fails during the upgrade process, you have three more servers in the cluster to handle user requests until you correct the failed upgrade. If you have only one server to provide file and print shares, applications, database, and e-mail, and *that* server fails, you can imagine the kind of chaos that would commence.

- Q:** What are the advantages of deploying an AD structure consisting of multiple domains, rather than a single domain with a separate OU for each department?
- A:** The chief difference between these two deployments is that of security requirements. Some security settings—such as auditing, password complexity requirements, and account lockout policies—can be implemented only at the domain level. If you have a group of users who require a substantially different set of security mechanisms than the rest of your network, you might wish to create a child domain for that group. Features such as two-way transitive trusts will still enable you to manage multiple domains centrally.
- Q:** I have recently begun a new position as a network administrator for a Windows Server 2003 forest containing many domains and child domains. The previous administrator created a number of GPOs, and it seems as if each network user has different policy settings applied to their accounts. I would like to simplify the GPO implementation on the network and wish to begin by creating a “baseline” report of exactly which GPOs are in effect for the various users on the network. What is the most efficient means of accomplishing this?
- A:** You can use the GPRresult command-line utility in the Windows Server 2003 Resource Kit. GPRresult provides the same functionality as the Resultant Set of Policy Logging mode, but you can run it from the command line, during each user’s logon script.
- Q:** What happens to Windows NT trust relationships when you upgrade to Windows Server 2003?
- A:** When you upgrade a Windows NT domain to a Windows Server 2003 domain, all of your existing Windows NT trusts will be preserved as-is. Remember that trust relationships between Windows Server 2003 domains and Windows NT domains are nontransitive.
- Q:** My company is working on a limited budget for its Windows Server 2003 upgrade. Do I need to provide separate licenses for the equipment in my training lab?
- A:** If the training lab machines will be either decommissioned or transferred from the test environment into production, you should not need a separate license than what you’ve budgeted for the machine upgrades. If, however, the test lab will be a permanent or semipermanent installed base of equipment, you do need to provide separate licensing for the software in the test lab.

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.

Overview of Network Infrastructure Planning

1. You are proposing the purchase of a new e-mail server for your corporate network. You have specified a new server from a major OEM manufacturer that is configured with a powerful quad-processor configuration, hot-swappable hard drives, and redundant power supplies and network adapters, with a three-year onsite warranty. Due to a budget crunch, the chairperson of the budget committee has suggested that the company can make do with a less powerful workgroup server from a local computer store. This server has only a single processor and no redundancy features, and a one-year onsite warranty. What reasons can you provide the budget committee members that might convince them to authorize the purchase of the server that you specified, even though it has a higher price tag?
 - A. A more powerful server will provide better performance and scalability as the company's needs grow over time.
 - B. Redundant hardware components will increase the server's availability to service the needs of the company's users and customers.
 - C. The extended warranty on the more powerful server will increase support costs over time, since you're paying to cover the machine under warranty for three times as long.
 - D. Windows Server 2003 requires at least a dual-processor configuration.
2. You are the network administrator of a Windows NT 4 domain for a shipping warehouse that operates 24 hours a day, 6 days a week. You perform a full nightly backup of all user files at 3:00 A.M. Users on the overnight shift are complaining that they are often locked out of files that they need access to while the backup process is running. You are proposing a network upgrade to Windows Server 2003 in the near future. What Windows Server 2003 feature will assist you in addressing this problem?
 - A. Disk quotas
 - B. NTFS file security
 - C. Volume Shadow Copy
 - D. Network Load Balancing

3. A portion of your company's organizational structure is shown in Figure 1.14. Third-level department managers report to the second-level department managers directly above them in the organizational chart. Second-level managers report to their corresponding vice presidents, who then report to the company CEO. Your company CEO would like a consistent security policy to be implemented across the entire network, but each subdepartment has specific desktop and application installation settings that you would like to be able to control and deploy centrally. What is the most efficient AD structure to design for this company?

Figure 1.14 Organizational Structure



- A. Configure a single domain for the organization, and configure a series of nested OUs for each second-level and third-level department. Configure the domain with a single security policy, and link a GPO to each OU to enable each specific department's desired settings.
- B. Configure a parent domain for each second-level department, and configure a child domain for each third-level department. Create and link a separate GPO to each domain to control security and application settings.
- C. Configure a single domain for the organization, and configure a global security group for each department. Configure the domain with a single security policy, and link a GPO to each global group to enable each specific department's desired settings.
- D. Create a separate forest for each second-level department, and create a child domain for each third-level department. Configure a security policy for each forest, and configure a domain GPO for each third-level department.

4. You are the administrator for a network that supports a mixture of Windows NT 4 Workstation, Windows 2000, and Windows XP Professional. You are preparing to upgrade your network servers from Windows NT Server to Windows Server 2003. What is the strongest level of network authentication that you can configure your Windows domain to use in its current configuration (without installing third-party software)?
- A. Kerberos
 - B. LM
 - C. NTLM
 - D. NTLM version 2

Analyzing Organizational Needs

5. You are the administrator of a Windows 2000 network and are planning an upgrade to Windows Server 2003. As part of the upgrade process, you are attempting to determine whether you need to upgrade your network cabling from Token Ring cabling to 100MB Ethernet. What is the best way to go about making this determination?
- A. Use Performance Monitor to capture a baseline of network utilization at several points during the day over the course of several weeks.
 - B. Use Network Monitor to capture network frames being sent to and from your domain controller's network adapter.
 - C. Use the IPSec Monitoring utility to view network traffic being sent between your domain controllers and your Windows 2000 Professional clients.
 - D. Use Performance Monitor to capture a single snapshot of network utilization when most users are in the office, such as mid-morning.
6. After returning from a two-day technology management seminar, your CEO tells you that he would like to create a fault-tolerant configuration for the company's heavily trafficked Web and database servers. Your network is currently running the Standard Edition of Windows NT 4.0. You have recently proposed an upgrade to Windows Server 2003. What features offered by this proposed upgrade would provide an attractive option to meet your CEO's request?
- A. SMP processing
 - B. Volume Shadow Copy
 - C. Network Load Balancing
 - D. Server clustering

7. You are the network administrator for a medium-sized company that consists of Sales, Customer Service, Accounting, Human Resources, and Data Entry departments. You have been receiving complaints that your company's e-mail server has been performing more slowly than usual over the past several weeks. Several users have mentioned that their e-mail clients have "frozen" in the middle of sending an e-mail message, forcing them to reboot their machines. Upon investigating, you find that one user's mailbox is roughly ten times the size of the second largest mailbox on the server, and this user is receiving approximately 1,000 messages per day, compared to a company average of 46. The user in question is a data-entry clerk who does not use e-mail for sales inquiries or other business-related contacts. When you ask the user about her e-mail usage, she reports that she has been surfing the Web signing up for Internet coupons and contests, and she has been deluged with spam as a result. Since the user does not require e-mail access to perform her job function, you disable her e-mail account, and server performance slowly returns to normal. What measures can you implement to prevent this sort of incident from recurring? (Select all that apply.)
- A. Implement disk quotas on the e-mail server so that users' inboxes cannot exceed a certain size.
 - B. Increase the level of authentication security so that only Kerberos-authenticated users can access the e-mail server.
 - C. Distribute an Acceptable Use Policy to your user base so that they understand what they can and cannot do while using their office PCs.
 - D. Use NTFS file permissions to restrict network access to personnel in your Sales and Customer Service department only.

Developing a Test Network Environment

8. You are the network administrator for a law firm that has multiple locations throughout the United States. Your firm has purchased a customer relationship management (CRM) application that will be hosted in the firm's main office in Key Biscayne, Florida, and accessed by other offices using dedicated WAN links. You would like to test the performance of this software over a WAN link before deploying it to the other offices in the firm. Unfortunately, you only have access to test equipment in the Key Biscayne office location. What is the best way to test the performance of this application?
- A. Use the average network bandwidth utilization in each office to estimate the performance of the application over the WAN.
 - B. Install routers within the test lab to simulate the latency of the dedicated WAN links between offices.

- C. Access the CRM application from your home computer using your high-speed Internet connection.
 - D. Test the application using production systems in each of the remote offices.
9. You are in the process of building a lab environment to test a new network application. You would like to isolate the test environment from your production equipment as much as possible to prevent any test changes from affecting your users' daily tasks. What can you do to protect your production environment from changes performed in your test lab? (Select all that apply.)
- A. Place a router or firewall between the network infrastructures connecting the test lab to your production machines.
 - B. Keep the network cabling for the test lab physically separated from the network hardware that provides connectivity to your production environment.
 - C. Contain the test lab in a separate OU.
 - D. Use 100MB Ethernet for your production machines, but only 10MB Ethernet for the test lab.
10. You are designing a lab environment to test a proposed upgrade to Windows Server 2003. You are in the process of creating a domain structure in the test lab to assess various features and functions of the upgrade process, including switching the domain from mixed mode to native mode and moving from a standard DNS zone to AD-integrated DNS. At the same time that the Windows Server 2003 testing is taking place, you would also like to use the test lab to evaluate a new accounting package that will be implemented on the production network before the Windows Server 2003 upgrade takes place. You do not want the two batteries of tests to interfere with each other. Which of the following would be good design choices for the domain structure of the test lab? (Select all that apply.)
- A. Create two separate domains: one to test the accounting software and one to test the domain mode and DNS functionality of Windows Server 2003.
 - B. Create a single domain in the test lab to encompass the entire test environment.
 - C. Create a separate OU to test the accounting software so that it will not be affected by the switch in domain mode.
 - D. Create two separate forests: one to test the DNS configuration and the switch from mixed mode to native mode and one to perform the tests on the accounting software package.

11. You have received a critical software update from the vendor of your accounting software suite. The software vendor has indicated that you should apply this patch as quickly as possible to correct a potential security breach. As the administrator for your network, what should you do when you receive this notice?
 - A. Install the patch on all production systems as quickly as possible.
 - B. Install the patch in your network's test lab to ensure that it functions properly and without any adverse side effects, and then apply it to all of your production systems as soon as possible.
 - C. Install the patch on a single workstation on your production environment to see if there are any bugs or malfunctions. When you are satisfied, apply the patch to the remainder of your workstations.
 - D. Send the software patch to Microsoft Product Support Services for testing before applying it to your network computers.

12. You are the network administrator for a small company that is considering purchasing a Windows Server 2003 machine to replace an aging Windows NT 4 Server machine. The client workstations run a mix of Windows 98, Windows NT Workstation, and Windows XP Professional. Each network client needs to be able to access the network server after it is upgraded, since the client workstations will be upgraded on a one-by-one basis over the course of several months. You have been informed that you will need to use the production server itself for testing, and that there is only sufficient budget to allot one representative workstation PC for test purposes. What is the best way for you to test client connectivity to Windows Server 2003?
 - A. Configure the test workstation with Windows Server 2003. Connect a production Windows 98, Windows NT 4, and Windows XP Professional workstation to the test server.
 - B. Use a utility like VMware to simulate how each operating system on your network will function with the new Windows Server 2003 server.
 - C. Check each client operating system one at a time, reformatting the test PC after you've finished testing each operating system.
 - D. Connect a production Windows 98 and Windows NT 4 Workstation to the Windows Server 2003 . Configure the test workstation to use Windows XP Professional.


Documenting the Planning and Network Design Process

13. You have recently started working as a network administrator for a company whose network consists of multiple Windows Server 2003 domains. The previous network administrator left you with little documentation detailing how the network is config-

ured, and you've discovered that many client workstations are behaving inconsistently—sometimes the Run line is unavailable, sometimes a user cannot access the Control Panel, and so on. You suspect that this is the result of Group Policy settings, and want to put together a list of all GPOs that are present within each domain on your network. What is the most efficient way of accomplishing this task?

- A. View each domain's settings within the Group Policy Management Console (GPMC) and take note of the values listed under the Group Policy node in each domain.
 - B. Use a GPMC script to list all GPO objects within each domain.
 - C. Load the Resultant Set of Policies (RSOP) snap-in to view the various GPOs that are causing client settings to be applied.
 - D. Examine the Group Policy tab of each domain's Properties sheet in Active Directory Users & Computers.
14. A portion of your network is shown in the Figure 1.15. You are using Network Monitor from WorkstationB to capture network traffic for analysis. You suspect that there is an Internet Relay Chat (IRC) connection between WorkstationA and WorkstationC, but the Network Monitor trace does not show any sign of that connection. What is the most likely reason for this?

Figure 1.15 Network Portion



- A. Network Monitor captures broadcast traffic only on a Windows network.
- B. Windows workstations do not support IRC connections.
- C. The version of Network Monitor that ships with Windows Server 2003 products does not operate in promiscuous mode.
- D. You need to use Performance Monitor to capture and analyze network traffic between machines on a Windows network.

15. Your company, airplanes.com, has recently undergone a merger with southern-airplanes.com, and you have taken over the network management of both halves of the newly formed company. Airplanes.com has a strict policy of desktop and software installation restrictions, while southern-airplanes.com has historically been more lenient with allowing users to customize their computers and install personal software. Several of the users from southern-airplanes.com have complained about the policy restrictions that have been placed on their desktops. You have been asked to present a report to the management group detailing which restrictions are in place on various OUs. What is the most efficient way to present this information to the management group in an easily readable format?
- A. Capture a screen shot of the Properties sheet of the various OUs' Group Policy settings and save the screen shot using a desktop publishing software package.
 - B. Export the GPO settings to a text file, then import the text file into an Excel spreadsheet.
 - C. Demonstrate the use of the Group Policy Editor to apply GPO settings during the meeting with the management group.
 - D. Use the Group Policy Management Console (GPMC) to present the various GPO settings in an organized HTML-formatted report.

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, B**

2. **C**

3. **A**

4. **D**

5. **A**

6. **C, D**

7. **A, C**

8. **B**

9. **A, B**

10. **A, D**

11. **B**

12. **B**

13. **B**

14. **C**

15. **D**

MCSE 70-293

Planning Server Roles and Server Security

Exam Objectives in this Chapter:


- 1 Planning and Implementing Server Roles and Server Security
 - 1.1 Configure security for servers that are assigned specific roles.
 - 1.4 Evaluate and select the operating system to install on computers in an enterprise.
 - 1.4.1 Identify the minimum configuration to satisfy security requirements.
 - 1.2 Plan a secure baseline installation.
 - 1.2.1 Plan a strategy to enforce system default security settings on new systems.
 - 1.2.2 Identify client operating system default security settings.
 - 1.2.3 Identify all server operating system default security settings.
 - 1.3 Plan security for servers that are assigned specific roles. Roles might include domain controllers, Web servers, database servers, and mail servers.
 - 1.3.1 Deploy the security configuration for servers that are assigned specific roles.
 - 1.3.2 Create custom security templates based on server roles.

Introduction

Planning an effective security strategy for Windows Server 2003 requires an understanding of the roles that different servers play on the network and the security needs of different types of servers based on the security requirements of your organization. Securing the servers is an important part of any network administrator's job.

In this chapter, we will first review server roles and ensure that you have an understanding of the many roles Windows Server 2003 can play on the network. We will discuss domain controllers; file and print servers; DHCP, DNS, and WINS servers; Web servers; database servers; mail servers; certification authorities; and terminal servers. Then we will delve into how to plan a server security strategy. We will examine how to choose the right operating system according to security needs, how to identify minimum security requirements for your organization, and how to identify the correct configurations to satisfy those security requirements.

Next, you will learn how to plan baseline security on both client and server machines. We will cover planning the secure baseline installation parameters and enforcing default security settings on new computers. We will show you how to customize server security, securing your servers according to their roles. Then we will walk you through the process of creating custom security templates and show you how to deploy security configurations.



Understanding Server Roles

When Windows Server 2003 is installed on a computer, it provides a wide variety of tools and functionality. However, additional features may still need to be installed on the server to bring clients the services they need. The server may need to supply file and print services, authenticate users, or support a local intranet Web site. Until Windows Server 2003 is configured to supply these services, clients will be unable to use the server in a manner that is required by the organization.

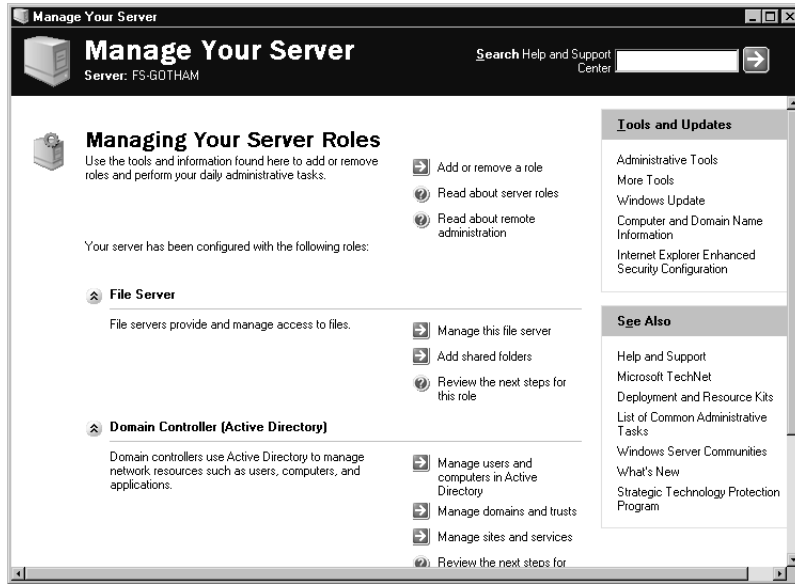
Server roles are profiles that are used to configure Windows Server 2003 to provide specific functionality to the network. When you set up a server to use a specific role, various services and tools are enabled or installed, and the server is configured to provide additional services and resources to network clients. Roles are applied to machines using the Configure Your Server Wizard and managed using the Manage Your Server tool.

As shown in Figure 2.1, Manage Your Server provides information about the roles that are currently configured for a server, and it provides the ability to add and remove roles from a server. Depending on your server's settings, this tool will start automatically upon logon. If you've checked the **Don't display this page at logon** check box at the bottom of this window, Manage Your Server will not start automatically. You can start it manually by selecting **Start | Administrative Tools | Manage Your Server**.

As shown in Figure 2.1, there are a variety of items in Manage Your Server's main window. The left side of the window lists the roles currently configured for the server. Beside each entry, there are buttons that relate to the corresponding role. These buttons

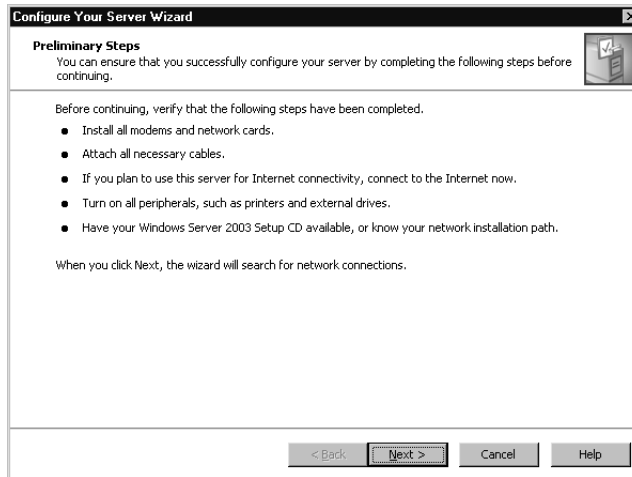
differ from role to role, and they are used to invoke other tools for managing the role or to view information on additional steps that can be taken to configure, administer, and maintain the role.

Figure 2.1 The Main Manage Your Server Window

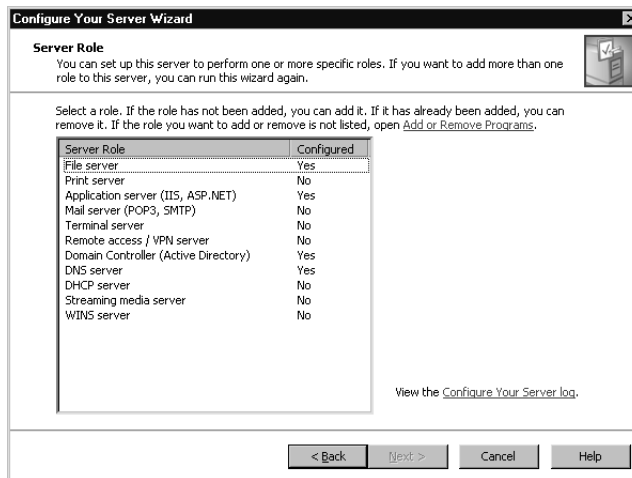


Near the top of the Manage Your Server window are three buttons. Two of these are used to obtain additional information about roles and remote administration. The other button, labeled **Add or remove a role**, is used to invoke the Configure Your Server Wizard. You can also start the Wizard by selecting **Start | Administrative Tools | Configure Your Server**.

When the Configure Your Server Wizard starts, it informs you of possible preliminary steps that need to be taken before a new role is added. As shown in Figure 2.2, these steps include ensuring that network and Internet connections have been set up and are active for the server, peripherals are turned on, and your Windows Server 2003 installation CD is available. When you finish reading this information, click the **Next** button to have the Wizard test network connections and continue to the next step.

Figure 2.2 Preliminary Steps of the Configure Your Server Wizard

In the next window, shown in Figure 2.3, roles that are available to add and remove through the Wizard are listed in the **Server Role** column; the **Configured** column indicates whether the role has been previously installed. If you want to install a role that isn't listed here, click the **Add or Remove Programs** link to open the Add or Remove Programs applet (in the Windows Control Panel), where you can configure additional services.

Figure 2.3 Configuring Server Roles

In Figure 2.3, you can see that there are 11 different roles that can be applied to Windows Server 2003 through the Configure Your Server Wizard. These roles are as follows:

- **Domain controller** This role is used for authentication and installs Active Directory on the server.
- **File server** This role is used to provide access to files stored on the server.
- **Print server** This role is used to provide network printing functionality.
- **DHCP server** This role allocates IP addresses and provides configuration information to clients.
- **DNS server** This role resolves IP addresses to domain names (and vice versa).
- **WINS server** This role resolves IP addresses to NetBIOS names (and vice versa).
- **Mail server** This role provides e-mail services.
- **Application server** This role makes distributed applications and Web applications available to clients.
- **Terminal server** This role provides Terminal Services for clients to access applications running on the server.
- **Remote access/VPN server** This role provides remote access to machines through dial-up connections and virtual private networks (VPNs).
- **Streaming media server** This role provides Windows Media Services so that clients can access streaming audio and video.

After you select the role to add to the server, click **Next** to step through the process of setting up that role. Each set of configuration windows is different for each server role. Also, although multiple roles can be installed on Windows Server 2003, only one role at a time can be configured using the Configure Your Server Wizard. To install additional roles, you need to run the Wizard again.

New & Noteworthy...

Manage Your Server

The Manage Your Server tool is new to Windows Server 2003. It is similar to the Configure Your Server utility in Windows 2000 and provides a centralized location for administrators to access tools, view information, and launch programs used to maintain specific roles. In addition, servers with Internet access can benefit from this tool, because it can be used to invoke Windows Update to apply security patches, service packs, new drivers, and other updates. Manage Your Server also provides links to Web pages located on Microsoft's site, which can assist administrators in understanding how to deal with specific problems and obtaining the latest information.

Manage Your Server also provides a way to launch the Configure Your Server Wizard, where you can add roles to a server or remove existing ones. Because the roles installed on a server can be modified at any time, administrators are able to change a server's role on the network as needs within the organization change.

Before setting up a server role (as we will do in Exercise 2.1, later in this chapter), it is important to understand each of the roles that can be applied to Windows Server 2003. In the sections that follow, we will discuss these roles in greater detail and examine how they are installed with the Configure Your Server Wizard and other tools.

Domain Controllers (Authentication Servers)

Domain controllers are a fundamental part of a Microsoft network because they are used to manage domains. A *domain* is a logical grouping of network elements, including computers, users, printers, and other components that make up the network and allow people to perform their jobs. When a server is configured to be a domain controller (DC), it can be used to manage these objects and provide other capabilities for configuring and controlling your network.

An important function of a domain controller is user authentication and access control. *Authentication* is used to verify the identity of an object such as a user, application, or computer. For example, when a user logs on to a domain, he or she will enter a username and password, which is compared to information that is stored on the domain controller. If the information provided by the user matches data in the user account, the domain controller considers the person to be authentic. The process continues by giving an appropriate level of access, so the user can utilize resources on the network. *Access control* manages which services and resources users (or other objects) are permitted to use and how they can use them. By combining authentication and access control, a user is permitted or denied access to network services and resources.

Active Directory

To perform these functions, the domain controller must have information about users and other objects in a domain. In Windows 2000 and Windows Server 2003, this data is stored in *Active Directory* (AD), which is a directory service that runs on domain controllers. A *directory* serves as a structured source of information, containing data on objects and their attributes. *Objects* in the directory represent elements of your network (including users, groups, and computers). *Attributes* are values that define an object (such as its name, location, security rights, and other features). Using tools that access AD, an administrator can manage an object's attributes to provide information that is accessible to users and control security at a granular level. By serving as a data store of information about a domain, AD is the means by which administrators achieve greater and more flexible control over a network.

When AD is installed, the server becomes a domain controller. Until this time, it is a member server that cannot be used for domain authentication and management of domain users or other domain-based objects. This does not mean, however, that AD can be installed on every version of Windows Server 2003. It can be installed on Standard Edition, Enterprise Edition, and Datacenter Edition, but servers running the Web Edition of Windows Server 2003 cannot be domain controllers. Web Edition servers can be only stand-alone or member servers that provide resources and services to the network.



EXAM WARNING

A server without AD installed on it can still deliver a variety of services, file storage, and access to other resources. However, until AD is installed, the server cannot authenticate domain users or provide the other functions of a domain controller. Once AD is installed, the member server ceases to be a member server and becomes a domain controller.

A Windows Server 2003 computer can be changed into a domain controller by using the Configure Your Server Wizard or by using the Active Directory Installation Wizard (DCPROMO). DCPROMO is a tool that promotes a member server to domain controller status. During the installation, a writable copy of the AD database is placed on the server's hard disk. The file used to store directory information is called NTDS.dit and, by default, is located in `%systemroot%\NTDS`. When changes are made to the directory, they are saved to this file.

Each domain controller retains its own copy of the directory, containing information about the domain in which it is located. If one domain controller becomes unavailable, users and computers can still access the AD data store on another domain controller in that domain. This allows users to continue logging on to the network, even though the domain controller that is 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 domain controller, changes made to the directory on one domain controller must be updated on others. The process of copying these updates is called *replication*, which is used to synchronize information in the directory. Without replication, features in AD would fail to function properly. For example, if you added a user on one domain controller, 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 could not log on to other domain controllers until the account was replicated. When a change is made on one domain controller, the changes need to be replicated, so that every domain controller continues to have an accurate copy of AD. This type of replication is called *multi-master*, because each domain controller contains a full read/write copy of the AD database.

Operations Master Roles

By default, all domain controllers are relatively equal. However, there are still some operations that need to be performed by a single domain controller in the domain or forest. To address these, Microsoft created the concept of *operations masters*. Operations masters serve many purposes. Some control where components of AD can be modified; others store specific information that is key to the healthy function of AD at the domain level. Because only one domain controller in a domain or forest fulfills a given role, these roles are also referred to as *Flexible Single Master of Operations* (FSMO) roles.

Some FSMO roles are unique to each domain; others are unique to the forest. A *forest* is one or more domain trees that share a common schema, Global Catalog, and configuration information. The *schema* is used to define which types of objects (classes) and attributes can be used in AD. Without it, AD would have no way of knowing what objects can exist in the directory or what attributes apply to each object. The *Global Catalog* is a subset of information from AD. It stores a copy of all objects in its host domain, as well as a partial copy of objects in all of the other domains in the forest.

There are five different types of master roles, each serving a specific purpose. Two of these master roles are applied at the forest level (forest-wide roles), and the others are applied at the domain level (domain-wide roles). The following are the forest-wide operations master roles:

- **Schema master** A domain controller that is in charge of all changes to the AD schema. As mentioned, the schema determines which object classes and attributes are used within the forest. If additional object classes or attributes need to be added, the schema is modified to accommodate these changes. The schema master is used to write to the directory's schema, which is then replicated to other domain controllers in the forest. Updates to the schema can be performed only on the domain controller acting in this role.
- **Domain naming master** A domain controller 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. This role prevents naming conflicts, because such changes can be performed only if the domain naming master is online.

In addition to the two forest-wide master roles, there are three domain-wide master roles: relative ID (RID) master, primary domain controller (PDC) emulator, and infrastructure master. These roles are described in the following sections.

Relative ID Master

The *relative ID master* is responsible for allocating sequences of numbers (called relative IDs, or RIDs) that are used in creating new security principles in the domain. Security principles are user, group, and computer accounts. These numbers are issued to all domain controllers in the domain. When an object is created, a number that uniquely identifies the object is assigned to it. This number consists of two parts: a domain security ID (or computer SID if a local user or group account is being created) and an RID. Together, the domain SID and RID combine to form the object's unique SID. The domain security ID is the same for all objects in that domain. The RID is unique to each object. Instead of using the name of a user, computer, or group, Windows uses the SID to identify and reference security principles. To avoid potential conflicts of domain controllers issuing the same number to an object, only one RID master exists in a domain. This controls the allocation of RID numbers to each domain controller. The domain controller can then assign the RIDs to objects when they are created.

PDC Emulator

The *primary domain Controller (PDC) emulator* is designed to act like a Windows NT PDC when the domain is in Windows 2000 mixed mode. This is necessary if Windows NT backup domain controllers (BDCs) still exist on the network. Clients earlier than Windows 2000 also use the PDC emulator for processing password changes, though installation of the AD client software on these systems enables them to change their password on any domain controller in the domain to which they authenticate. The PDC emulator also synchronizes the time on all domain controllers the domain. For replication accuracy, it is critical for all domain controllers to have synchronized time.

Even if you do not have any servers running as BDCs on the network, the PDC emulator still serves a critical purpose in each domain. The PDC emulator receives preferred replication of all password changes performed on other domain controllers within the domain. When a password is changed on a domain controller, it is sent to the PDC emulator. If a user changes his or her password on one domain controller, and then attempts to log on to another, the second domain controller may still have old password information. Because this domain controller considers it a bad password, it forwards the authentication request to the PDC emulator to determine whether the password is actually valid. In addition, the PDC emulator initiates urgent replication so that the password change can propagate as soon as possible. Urgent replication is also used for other security-sensitive replication traffic, such as account lockouts.

This operations master is by far the most critical at the domain level. Because of this, you should ensure that it is carefully placed on your network and housed on a high-availability, high-capacity server.

Infrastructure Master

The *infrastructure master* is in charge of updating changes that are made to group memberships. When a user moves to a different domain and his or her group membership changes, it may take time for these changes to be reflected in the group. To remedy this, the infrastructure master is used to update such changes in its domain. The domain controller in the infrastructure master role compares its data to the Global Catalog, which is a subset of directory information for all domains in the forest and contains information on groups. The Global Catalog stores information on universal group memberships, in which users from any domain can be added and allowed access to any domain, and maps the memberships users have to specific groups. When changes occur to group membership, the infrastructure master updates its group-to-user references and replicates these changes to other domain controllers in the domain.



TEST DAY TIP

FSMO roles are an important part of a domain controller's function on a network. FSMO roles that are unique to a forest affect all domains within that forest. FSMO roles that are unique to a domain apply only to that domain. There is only one schema master and one domain naming master in a forest. There is only one RID master, PDC emulator, and infrastructure master in a domain.

File and Print Servers

Two of the basic functions in a network are saving files in a central location on the network and printing the contents of files to shared printers. Each of these functions is vital to most environments. Most organizations require users to be able to save their work to a shared location on the network and to print hard copies of it for others to review and/or retain. When file server or print server roles are configured in Windows Server 2003, additional functions become available that make using and managing the server more effective.

Print Servers

Print servers are used provide access to printers across the network. A benefit of print servers for administrators is that they provide an added level of manageability for network printing. Print servers allow you to control when print devices can be used by allowing you to schedule the availability of printers, set priority for print jobs, and configure printer properties. Using a browser, an administrator can also view, pause, resume, and/or delete print jobs.

By configuring Windows Server 2003 in the role of a print server, you can manage printers remotely through the GUI and by using Windows Management Instrumentation (WMI). WMI is a management application program interface (API) that allows you to monitor and control printing. Using WMI, an administrator can manage components like print servers and print devices from a command line.

Print servers also provide alternative methods of printing to specific print devices. Users working at machines running Windows XP can print to specific printers by using a *Uniform Resource Locator* (URL). If you've used the Internet, you're probably already familiar with URLs. A URL is the address that is entered to access a Web site. Using URLs, other resources can also be accessed from remote locations, such as printers offered by Windows Server 2003 print servers.

File Servers

File servers are used to provide access to files that are stored on the server's hard disks. Users are able to store files in a centralized location, rather than to their local hard disks, and share them with other users. When a file is saved to a volume on a file server, clients who have access to the directory in which the file was saved can access it remotely from the server. This type of server is also important when multiple employees use network-accessible

applications. In such cases, data may need to be saved from the application to a shared database, spreadsheet, or other type of file.

Administrators benefit from file servers by being able to manage disk space, control access, and limit the amount of space that is made available to individual users. If NTFS volumes are used, disk quotas can be set to limit the amount of space available to each user. This prevents users from filling the hard disk with superfluous data or older information that may no longer be needed.

In addition to these features, a file server also provides other functionality that offers security and availability of data. File servers with NTFS volumes have the *Encrypted File System* (EFS) enabled, so that any data can be encrypted using a public key system. This makes it difficult for unauthorized users to access data, while being transparent to authorized users. To make it easier for users to access shared files, the *Distributed File Service* (DFS) can be used, which allows data that is located on servers throughout the enterprise to be accessible from a single shared folder. When DFS is used, files stored on different volumes, shares, or servers appear as if they reside in the same location. This makes it easier for users to find the data they need, because they do not need to search through multiple locations to access the files they are permitted to use.

DHCP, DNS, and WINS Servers

The roles of DHCP, DNS, and WINS servers are used for uniquely identifying computers and finding them on the network. A DHCP server issues a unique number called an IP address to a computer. DNS and WINS servers resolve this number to and from user-friendly names that are easier for users to deal with. With Windows Server 2003 acting as a DHCP, DNS, and/or WINS server, clients can be automatically issued a number that distinguishes them on the network, and find other machines and devices more effectively.

DHCP Servers

DHCP is the *Dynamic Host Configuration Protocol*, and it is used to issue IP addresses to clients on networks using the Transmission Control Protocol/Internet Protocol (TCP/IP). An *IP address* is a number that uniquely identifies a client when sending or receiving packets of data. When information is sent across the network, the data is broken up into smaller packets, which are reassembled by the receiver. Each packet contains the IP address of who is sending the data and who should receive it. This is similar to a letter with an address of who should receive the message and a return address of who sent it.

Because no two computers on a network can have the same IP address at the same time, assigning these addresses to clients is an important responsibility. IP addresses can be assigned statically, so that each computer always uses the same IP address. Allocating addresses in this way can result in mistakes and is difficult to consistently track. Many enterprises use static IP addresses only for their servers and network infrastructure equipment (switches, routers, and so on). Dynamic addresses are used for all clients. Dynamic addresses are assigned using DHCP. When an IP address is dynamically assigned, the client contacts

the DHCP server for an IP address. The DHCP server responds by issuing an IP address from a pool of available addresses stored in a database, as well as any configuration information (such as the IP addresses of the default gateway, DNS server, and WINS server) that is needed by the client.

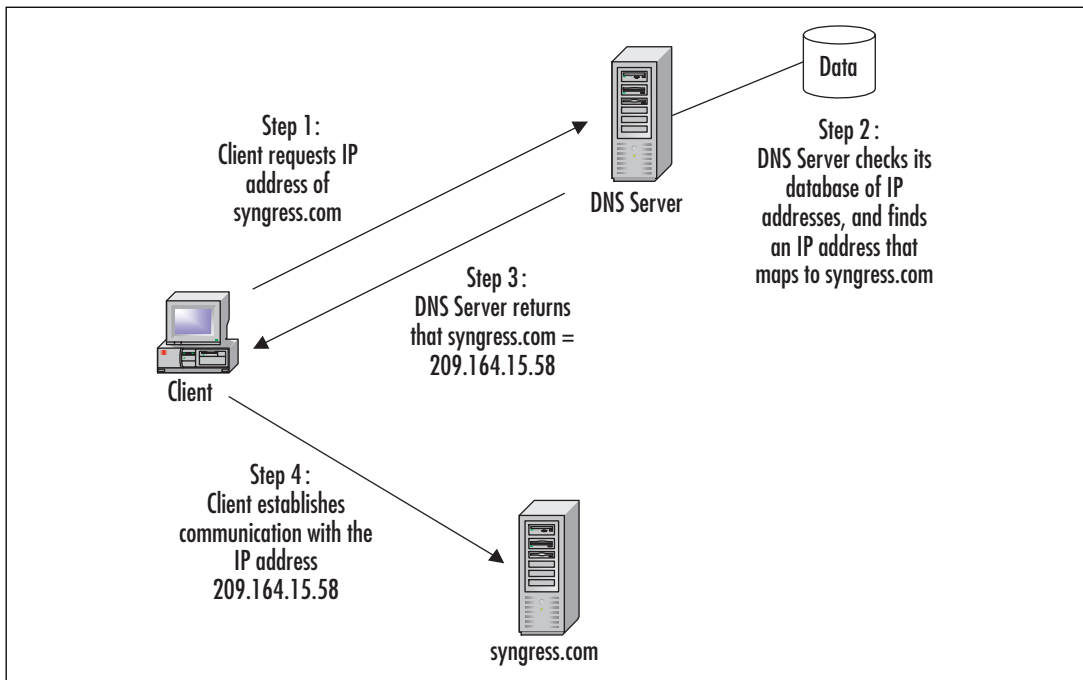
When a DHCP server allocates an IP address to the client, it is for a limited amount of time. Because there are only so many IP addresses available in a pool, they are often recycled between computers. This can happen if a client is shut off for an extended period of time, or if it is a laptop that is assigned to a user who is typically on the road and away from the office. For this reason, when a DHCP lease expires, the DHCP server is free to issue the IP address to other clients.

DNS Servers

Because remembering a series of numbers can be difficult, methods have been created to resolve IP addresses to user-friendly names and vice versa. Imagine trying to remember what Web site or computer the IP address 192.168.10.250 represented on a network, in addition to all the other IP addresses you would need to remember for other sites and computers. To remedy this situation, *name resolution* is used, so users can enter a name that is translated to a corresponding IP address.

The *Domain Name System* (DNS) is a popular method of name resolution that is used on the Internet and other TCP/IP networks. AD is integrated with DNS, and it uses DNS servers to allow users, computers, applications, and other elements of the network to easily find domain controllers and other resources on the network. DNS is a hierarchical, distributed database that maps user-friendly domain names (like syngress.com) to IP addresses. When a user enters a DNS name into a browser or other application, it is sent to a DNS server, which looks up the IP address for that domain. This IP address is sent back to the client, which uses the numeric address to locate and communicate with the computer at this address.

Figure 2.4 illustrates name resolution using DNS. In this example, a user wants to connect with the syngress.com domain. As shown in step 1 of this figure, because machines use IP addresses to locate and communicate with each other on a TCP/IP network, the client contacts the DNS server and requests the IP address of syngress.com. In step 2, the DNS server checks its database to find the IP address that maps to this particular domain name. After finding it, step 3 is performed, and the DNS server sends the information back to the client, informing it that the IP address of syngress.com is 209.164.15.58. Now that the client has this information, the client performs step 4, by connecting to syngress.com using the numeric address.

Figure 2.4 Name Resolution Using DNS

WINS Servers

The *Windows Internet Name Service* (WINS) is another method of name resolution that resolves IP addresses to NetBIOS names, and vice versa. *NetBIOS* names are used by pre-Windows 2000 servers and clients, and they allow users of those operating systems to log on to Windows Server 2003 domains. They are supported in Windows Server 2003 for backward-compatibility with these older systems. By implementing a WINS server, you allow clients to search for computers and other resources by computer name, rather than by IP address.

WINS is similar to DNS in that user-friendly names are mapped to IP addresses within a database. When clients attempt to connect to a computer or resource using its NetBIOS name, they can send a request to a WINS server to provide the IP address of that resource. The WINS server searches its database for the name-to-address mapping and returns the IP address to the requesting client. Once the client has this address, it can connect to and communicate with the computer or resource.

Web Servers

Web servers allow organizations to host their own Web sites on the Internet or a local intranet. An *intranet* is a local area Network (LAN) that uses the same technologies that are used on the Internet, so that users can access Web pages and other resources using Web

browsers and other Web-enabled applications. Implementing a Web server in an organization allows users to benefit by accessing information, downloading files, and using Web-based applications.

Web Server Protocols

Microsoft's Windows Server 2003 Web server product is *Internet Information Services (IIS) 6.0*, which is included with Windows Server 2003. IIS allows users to access information using a number of protocols that are part of the TCP/IP suite, including the following:

- **Hypertext Transfer Protocol (HTTP)** Used by the World Wide Web Publishing service in IIS. Allows users to access Web pages using a Web browser like Internet Explorer or other Web-enabled applications. By connecting to sites created on your Web server, users can view and work with Web pages written in the Hypertext Markup Language (HTML), Active Server Pages (ASP), and Extensible Markup Language (XML). This allows users to not only view static information, but also to benefit from Web-based programs.
- **File Transfer Protocol (FTP)** Used for transferring files between clients and servers. Using this service, clients can copy files to and from FTP sites using a Web browser like Internet Explorer or other FTP client software. By using such software, clients can browse through any folders they have access to on the FTP site, and they can access any files they have permissions to use.
- **Network News Transfer Protocol (NNTP)** Used for newsgroups, which are also called discussion groups. The NNTP service in IIS allows users to post news messages. Other users can browse through messages stored on the server, respond to existing messages, and post new ones using a newsreader program. For example, a group of users could have a discussion group that deals with a certain project, so that members of the team can exchange ideas and discuss problems in a forum that can be viewed by all members of the group. Another group could also be created that allows employees to post messages regarding items for sale, charitable events, or other things that you might see on a typical bulletin board. NNTP allows organizations to incorporate such message groups into the way that employees exchange information with one another.
- **Simple Mail Transfer Protocol (SMTP)** Used to provides e-mail capabilities (as described in the discussion of the mail server role later in this chapter). The SMTP service that is installed with IIS isn't a full e-mail service, but provides limited services for transferring e-mail messages. Using this service, Web developers can collect information from users of a Web site, such as having them fill out a form online. Rather than storing the results of the form locally in a file, the information can be e-mailed using this service.


Web Server Configuration

Although a Web server can facilitate a company’s ability to disseminate information, it isn’t an actual role that is configured using the Configure Your Server Wizard. It is installed as part of the application server role, which we’ll discuss later in this chapter. The Configure Your Server Wizard provides an easy, step-by-step method of configuring Web servers through the application server role; however, it isn’t the only way to install IIS. You can also install IIS through the Add or Remove Programs applet in the Windows Control Panel.

Using Add or Remove Programs to install IIS takes a few extra steps, but it allows you to perform the installation without installing other services and features available through the application server role. To use Add or Remove Programs to install IIS, follow these steps:

1. Select **Start | Control Panel | Add or Remove Programs**.
2. Click the **Add/Remove Windows Components** icon to display the **Windows Components Wizard**, which provides a listing of available components to install.
3. In the list, select **Application Server** and click the **Details** button to view the **Application Server** dialog box, shown in Figure 2.5.

Figure 2.5 Installing IIS through the Application Server Dialog Box in the Windows Components Wizard



4. The **Application Server** dialog box contains a number of subcomponents. To install IIS, select the check box for **Internet Information Services (IIS)**, and either click **OK** to install the default components or click **Details** to view even more subcomponents that can be installed within IIS.
5. When you’ve made your selections, click **OK** to return to the **Windows Components Wizard**.

6. Click **Next** to have Windows make the configuration changes you requested from your selection.
7. Once the Wizard has finished copying the necessary files and changing system settings, click **Finish** to complete the installation process and exit the Wizard.

Database Servers

Database servers are used to store and manage databases that are stored on the server and to provide data access for authorized users. This type of server keeps the data in a central location that can be regularly backed up. It also allows users and applications to centrally access the data across the network. A large number of the databases used in your organization can be kept on one server or a group of servers that are specifically configured to protect data and service client requests.

The Configure Your Server Wizard does not include a configurable role for database servers. A database server is any server that runs a network database application and maintains database files, such as Microsoft SQL Server or Oracle. SQL Server is a high-performance database management system. It is used for data storage and analysis, and it provides users with the ability to access vast amounts of data quickly over the network. Because SQL Server provides additional measures of security that would not otherwise be available (as discussed in the “Securing Database Servers” section later in this chapter) and processing occurs on the server, transactions can occur securely and rapidly.

Data stored in database management systems is generally accessed through user interfaces that are developed by an organization or third parties. For example, a company might create custom applications in Visual Basic (or some other programming language), or use ASP on the Web server to display information that is stored in a database. While the user interacts with the data through the user interface, the data is actually stored in the SQL Server or Oracle database located on a database server.

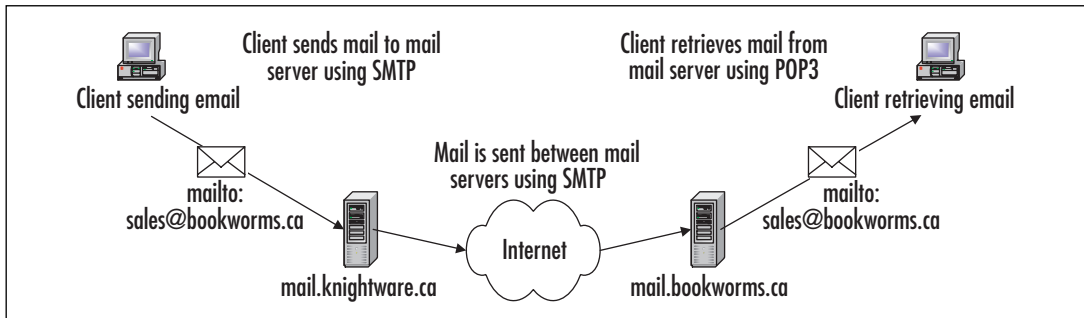
Mail Servers

Mail servers enable users to send and receive e-mail messages. Users send e-mail to other users through at least one mail server. When the message arrives, the destination mail server stores the message until it is retrieved by the user. If the mail server does not handle the e-mail account for an intended recipient, it will transfer the message to a mail server that does. In this way, mail servers will work together to ensure a message reaches its intended audience.

When a server is configured to be a mail server, two protocols are enabled: SMTP and Post Office Protocol (POP3). As shown in Figure 2.6, SMTP is used by clients and mail servers to send e-mail. POP3 is used by clients when retrieving e-mail from their mail server. Each of these protocols is part of the TCP/IP protocol suite and installed when TCP/IP is installed on a computer. However, even if TCP/IP is installed on Windows

Server 2003, the services provided by mail servers still need to be enabled by configuring the machine to take the role of a mail server.

Figure 2.6 How E-mail Is Transmitted and Retrieved



E-mail addresses determine which mail server and client the e-mail should go to. Each e-mail address uses the format of *account@domain*. The first part of the address specifies the account the e-mail is destined to reach, and the second part specifies the domain in which this account resides. In the example in Figure 2.6, a message destined for *sales@bookworms.ca* is sent from the *knightware.ca* domain. Because the mail server in *knightware.ca* recognizes that the message is being sent to a user in another domain, it uses the SMTP protocol to send it to the mail server in the *bookworms.ca* domain. When the *bookworms.ca* mail server receives this e-mail, it will see it is for the account named *sales* and put it in the mailbox for that user. The client that uses the *sales* account can then use the POP3 protocol to retrieve his or her e-mail from the mail server.

Certificate Authorities

Certificate authorities (CAs) are servers that issue and manage certificates. *Certificates* can be used for a variety of purposes, including encryption, integrity, and verifying the identity of an entity, such as a user, machine, or application. Certificates can be used to prove an entity is who (or what) they claim to be, in much the same way that your birth certificate is used to prove your identity. They are digitally signed files that contain data a wide range of information, often including a cryptographic key, information about whom or what the key is issued to, an expiration date, where the validity of the certificate can be checked, and which CA signed the certificate. Certificates are typically part of a larger security process known as a *Public Key Infrastructure (PKI)*.

PKI

PKI is a method that uses unique identifiers called *keys*, which are mathematical algorithms used for cryptography and authentication. There are two different kinds of keys used in PKI: public keys and private keys.

For data confidentiality, the public key is used to encrypt session keys and data; the private key is used for decryption. The public key is openly available to the public. 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 PKI.

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 it 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 pass it on to the CA. When the CA responds to the request, it sends it to the RA, which 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 who 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 PKI 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 this 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 used to decrypt the message.

Conversely, if a person wants to send a message and validate 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.

Certificates

Certificates use PKI by binding the value of a public key to the person or thing that holds the private key. The certificate stores information that identifies its holder and contains a copy of the key value. When communicating with another party that has a corresponding key, data exchanged between the two can be securely transmitted using encryption.

Certificates may be used for a number of different purposes. Windows 2003 Server computers acting in the role of a Web server may use certificates to authenticate users or to authenticate Web servers themselves. In doing so, the certificate provides proof of the identity of a particular user or machine. Mail servers can also benefit from certificates, because they are used to allow e-mail to be digitally signed. This provides proof of the integrity and

origin of a message. In sending secure mail, certificates are used with Secure/Multipurpose Internet Mail Extensions (S/MIME), which allows the e-mail to be sent encrypted across a network.

Certificates may also be used by different protocols to ensure secure communication, as in the case of Internet Protocol Security (IPSec) or Transport Layer Security (TLS). Encrypting communication between clients and servers with these protocols allows data to be transmitted and users to be authenticated with little (or no) chance of others intercepting and viewing the information. By using certificates for authentication, encryption/decryption of data, and secure communication, Windows 2003 Servers Certificate Services can provide enhanced security to a network.

Certificates can contain a variety of facts about a user's or machine's identity and about the certificate itself. Data included in a certificate may include the following:

- The value of a key issued by a CA
- Information about the person, machine, or other entity that was issued the certificate, which may include their name, e-mail address, or other data
- Information about who issued the certificate
- The digital signature of the issuer, which ensures the certificate is valid
- How long the certificate is valid

Because different systems must be able to understand the format of a certificate, specific standards are used in the generation of a certificate. Windows 2003 Server supports X.509, which is a standard that specifies the syntax and format of digital certificates. X.509 is a popular standard for digital certificates, published by the International Organization for Standardization (ISO). It dictates how information is organized in the certificate and what information is included. An X.509 certificate includes facts about the user to whom the certificate was issued, information about the certificate itself, and can also include information about the issuer of the certificate (who is referred to as the CA). To prevent the certificate from being used indefinitely, it also contains information about the period for which the certificate is valid.

Certificate Services

Certificate Services is used to create a CA on Windows Server 2003 servers in your organization. With Certificate Services, you can create a CA, format and modify the contents of certificates, verify information provided by those requesting certificates, issue and revoke certificates, and publish a Certificate Revocation List (CRL). The CRL is a list of certificates that are expired or invalid, and it is made available so that network users can identify whether certificates they receive are valid.

Certificate Services supports implementing a hierarchy of CAs, so that a single CA isn't responsible for providing certificates to the entire network or authenticating the entire intranet or Internet. This isn't to say that multiple CAs must be used in an organization, but

it is one possibility. Using a hierarchy of CAs is called *chaining*, where one CA certifies others. In this hierarchy, there is a single root authority and any number of subordinate CAs.

A *root authority* (or root CA) resides at the top of the hierarchy. Because the hierarchy uses a parent-child relationship, all subordinate CAs reside beneath the root authority. The root CA is the most trusted CA in the hierarchy—any clients that trust the root CA will also trust certificates issued by any CA below it. This makes securing a CA vital (as discussed in the “Securing CAs section later in this chapter).

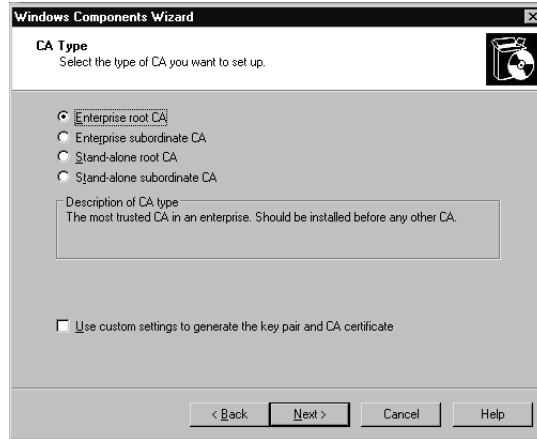
Subordinate CAs are child CAs in the hierarchy. They are certified by the root authority and bind its public key to its identity. Just as the root CA can issue and manage certificates and certify child CAs, a subordinate CA can also perform these actions and certify CAs that are subordinate to it in the hierarchy.

In addition to having different levels of CAs in an organization, there are also different types of root and subordinate CAs that can be used. *Enterprise CAs* use AD to verify information that is provided when requesting a certificate and to store certificates within AD. When the certificate is needed, it is retrieved from directory services. *Stand-alone CAs* can be used in environments that do not use AD (CAs do not require AD).

As with IIS, Certificate Services isn’t an actual role that can be set up with the Configure Your Server Wizard. Instead, you must follow these steps:

1. Select **Start | Control Panel | Add or Remove Programs**.
2. Click **Add/Remove Windows Components** to display the **Windows Components Wizard**, which provides a listing of available components to install.
3. In the list of available components, click the check box beside the **Certificate Services** item so it is checked. A warning message will appear, stating that after Certificate Services is installed, the name of the machine cannot be changed. This is because the server’s name is bound to the CA information stored in AD, and any changes to the name or domain membership would invalidate certificates issued by this CA.
4. Click **Yes** to continue with the installation. (Clicking **No** will cancel it.)
5. You are presented with the window shown in Figure 2.7, which allows you to specify the type of CA that will be set up. As mentioned earlier, you have the option of creating an enterprise root CA, an enterprise subordinate CA, a stand-alone root CA, or a stand-alone subordinate CA.

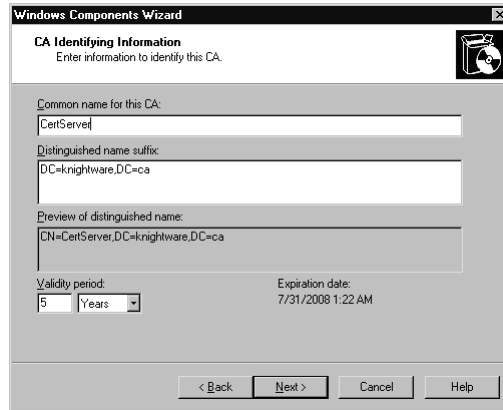
Figure 2.7 Choosing a CA Type in the Windows Components Wizard



6. For this example, we will assume that this is the first CA being created and AD is used. Select **Enterprise root CA** and click **Next**
7. You are then presented with a window shown in Figure 2.8, which allows you to provide information to identify the CA you’re creating. Enter a common name and distinguished name suffix for the CA. Distinguished names are used to provide each object in AD with a unique name. A distinguished name 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, which is the common name of the object, and includes such things as user accounts, printers, and other network elements represented in the directory.
 - OU, which is the Organizational Unit. OUs are containers in the directory, which 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, which is a domain component. This is used to identify the name of the domain or server, and the DNS suffix (for example .com, .net, .edu, .gov, and so forth).


When combined, these components of a distinguished name are used to show the location of an object. In the case of the CA being created here, the common name is CertServer, and the distinguished name suffix is the domain components. This makes the distinguished name CN=CertServer,DC=knightware,DC=ca, which you can see in the preview in Figure 2.8.

Figure 2.8 Entering CA Identifying Information in the Windows Components Wizard



8. Optionally, you can change the **Validity period** of certificates issued by the CA. As shown in Figure 2.8, the default validity period is five years. You can modify this by specifying a different number and whether the period is in **Years**, **Months**, **Weeks**, or **Days**.
9. Click **Next** when you are finished entering CA identifying information.
10. This will bring you to the **Certificate Database Settings** window, shown in Figure 2.9, where you can specify the location of the certificate database and log file. By default, the database and log are named after the common name you specified for the CA, and each is stored in the **System32** folder of the *%systemroot%* (for example, C:\Windows\System32). Click **Next** to continue.

Figure 2.9 Choosing Certificate Database Settings in the Windows Components Wizard



11. A message box will appear informing you that IIS must be stopped before installation can continue. Clicking **No** will return you to the previous window. Clicking **Yes** will stop the service and cause Windows to make the configuration changes you requested from your selection. If ASP is not enabled on the machine, a message box will interrupt the process, asking if you want to enable ASP. Clicking **Yes** will enable ASP and continue the installation.
12. After the Wizard has finished copying the necessary files and changing system settings, click **Finish** to complete the installation process.

Application Servers and Terminal Servers

Application servers and terminal servers provide the ability for users to access applications over the network. Rather than running solely on the client's machine, all or parts of these programs run on the server. This frees resources on the client machine and enables users to benefit from newer application technologies.

Application Servers

Application servers allow users to run Web applications and distributed programs from the server. Web applications are programs that use Internet technologies to provide functionality and are accessible across networks and the Internet using Web browsers like Internet Explorer. These programs are often created using ASP or XML. Applications can be created in a wider variety of programming languages (such as Perl, Visual Basic, and Visual C++). Distributed applications divide the program so that part of it runs on the client while the rest runs on one or more servers. For example, a distributed program might have a user interface that is installed on the client's machine, which allows the user to access a SQL Server database. In reality, the program might access a number of other network-aware programs, which correlate data from a number of different database systems and return it to the client. By using the application server role, the server is configured to provide greater reliability and performance to these applications.

Because Web applications require Internet technologies, when Windows Server 2003 is set up as an application server, IIS subcomponents such as ASP can be installed. As explained earlier in this chapter, IIS is a Web server that comes with Windows Server 2003 and can be used to make Web applications available to users on the network. If IIS has been installed, the application server role will appear as a configured role in the Manage Your Server tool. This is despite the fact that only some components for the application server role have been installed. To modify the installed components, you can either use the Windows Components Wizard or the Configure Your Server Wizard.

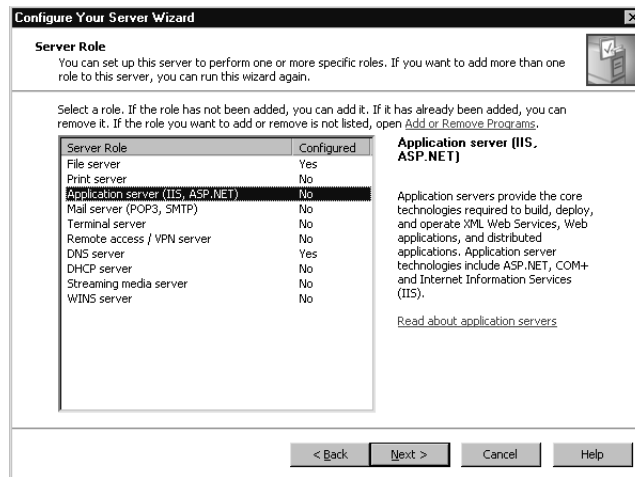
As an example of configuring a server role, in Exercise 2.1, we will set up an application server in Windows Server 2003.

EXERCISE 2.01

ADDING AN APPLICATION SERVER ROLE TO WINDOWS SERVER 2003

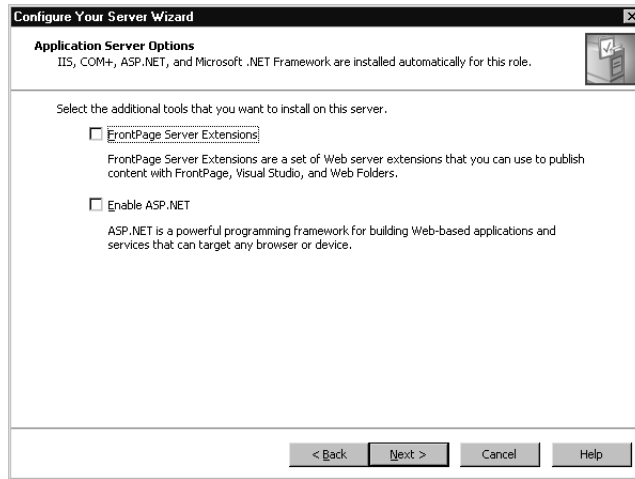
1. Select **Start | Administrative Tools | Manage Your Server**.
2. When **Manage Your Server** starts, click the **Add or remove a role** button.
3. When the **Configure Your Server Wizard** starts, read through the information on the **Preliminary Steps** window, and then click **Next**.
4. After the Wizard checks your network settings and operating system version, the **Server Role** window will appear. From the list, select **Application server (IIS, ASP.NET)**, as shown in Figure 2.10. Then click **Next** to continue.

Figure 2.10 Choose the Application Server Role




5. The **Application Server Options** window appears, as shown in Figure 2.11. Here, you can add components that are used with IIS. Note that IIS will be installed regardless of what you select on this page. Select the **FrontPage Server Extensions** check box to add Web server extensions that allow content created with FrontPage, Visual Studio, and Web Folders to be published to the IIS Web site. Select **Enable ASP.NET** to allow Web-based applications created using ASP.NET to be used on the site. After selecting the options you wish to add, click **Next** to continue.

Figure 2.11 Select Application Server Options



6. The **Summary of Selections** window, shown in Figure 2.12, provides a list of components that will be installed as part of the application server configuration. Review these settings, and then click **Next** to begin installing these components.

Figure 2.12 Review the Summary of Selections




7. After copying files, the **Windows Components Wizard** will open and continue the installation. Once it has completed, you will be returned to the **Configure Your Server Wizard**. Click **Finish** to complete the installation.

Terminal Servers

Terminal servers allow remote access to applications using thin-client technology. This makes the user's machine act as a terminal emulator (similar to the concept of a dumb terminal). The user connects to the terminal server using client software installed on their machine, logs on to the Terminal Services session, and is presented with a user interface (normally a Windows Server 2003 desktop). Keystrokes and mouse clicks generated by the user at the client are sent to the terminal server. Updated screen images are sent back from terminal server to the client system. When working in a session, the user is essentially working at the server. All processing is occurring at the server, which is being interacted with through the client software.

A benefit of Terminal Services is that users can run programs that they might otherwise be unable to use. For example, a user running an older version of Windows might need to use Office XP, but she doesn't have the minimal requirements to install it. Through Terminal Services, she can connect and be presented with a Windows Server 2003 desktop. If Office XP is installed on the terminal server, the user can open and use the application. Because all processing is actually occurring on the server, the user can run applications that are impossible to install on her local system.

There are a wide variety of clients that can use Terminal Services. Client software is available for Windows 3.11 and later, as well as Macintosh and UNIX. Internet Explorer can also be used to access a terminal server, using the Web client software.




Planning a Server Security Strategy

The only truly secure network is one that is totally inaccessible. No one would be able to misuse applications, damage equipment, delete data, or mistakenly modify information. In providing this level of security, however, the network would also become useless, because it could not provide the services and resources needed by users. Security is always a trade-off between usability and protection. When planning security, you need to find an acceptable balance between the need to secure your network and the need for users to be able to perform their jobs.

In creating a security plan, it is important to realize that the network environment will never be completely secure. If people are willing to invest enough time, effort, and money into hacking a system, they will probably find a way in. The goal is to make it difficult for intruders to obtain unauthorized access, so it isn't worth their time to try or continue attempting to gain access. It is also critical to protect servers from potential disasters and to have methods to restore systems if they become compromised.

A good security plan considers the needs of a company and tries to balance it with their capabilities and current technology. As you'll see in the sections that follow, this means identifying the minimum security requirements for an organization, choosing an operating system, and identifying the configurations necessary to meet these needs. To develop a security plan, you must identify the risks that potentially threaten a network, determine what

countermeasures are available to deal with them, figure out what you can afford financially, and implement the countermeasures that are feasible.



Choosing the Operating System

In planning a strategy for server security, you will need to determine which operating systems will be used in the organization. Different network operating systems provide diverse features that can be used as part of your security strategy. If you're setting up a new network and need to choose a server operating system, or you're unfamiliar with what operating systems are used on an existing network, you will not know what features can be used for managing and maintaining security.

Of course, there are non-Microsoft network operating systems available to use on your server, but we will consider only the following Windows server systems here:

- Windows NT Server 4
- Windows 2000 Server
- Windows 2000 Advanced Server
- Windows 2000 Datacenter
- Windows Server 2003 Standard Edition
- Windows Server 2003 Enterprise Edition
- Windows Server 2003 Datacenter Edition
- Windows Server 2003 Web Edition

One of the first considerations for the operating system you choose will be the minimum system requirements for installing the operating system. Obviously, if your existing server cannot handle a particular version of Windows, you will not be able to install it. If this is the case, you will need to upgrade the hardware, purchase a new server to support the operating system you want, or choose an operating system that does match the current server's hardware. The minimum system requirements for Windows server operating systems are shown in Table 2.1.



NOTE

All of the Windows server operating systems also require a CD-ROM or DVD drive (except Window NT Server 4, which does not use a DVD drive), VGA or higher resolution monitor, keyboard, and mouse.

Table 2.1 Minimum System Requirements for Windows Server Operating Systems

Server	Computer/Processor	Memory (RAM)	Hard Disk	CPU Support
Windows NT Server 4	486/33 MHz or higher/Pentium, or Pentium Pro processor	16MB; 32MB recommended	Intel and compatible systems: 125MB available hard disk space minimum. RISC-based systems: 160MB available hard disk space	Up to 4 CPUs (retail version); Up to 32 CPUs available from hardware vendors
Windows 2000 Server	133 MHz or higher Pentium-compatible CPU	At least 128MB; 256MB recommended; 4GB maximum	2GB with 1GB free space; additional free space required for installing over a network	Up to 4 CPUs
Windows 2000 Advanced Server	133 MHz or higher Pentium-compatible CPU	At least 128MB; 256MB recommended; 8GB maximum	2GB with 1GB free space; additional free space required for installing over a network	Up to 8 CPUs
Windows 2000 Datacenter	Pentium III Xeon processors or higher	256MB	2GB with 1GB free space; additional free space required for installing over a network	8-way capable or higher server (supports up to 32-way)
Windows Server 2003 Standard Edition	133 MHz	128MB	1.5GB	Up to 4 CPUs
Windows Server 2003 Enterprise Edition	133 MHz for x86-based computers; 733 MHz for Itanium-based computers	128MB	1.5GB for x86-based computers; 2GB for Itanium-based computers	Up to 8 CPUs
Windows Server 2003 Datacenter Edition	400 MHz for x86-based computers; 733 MHz for Itanium-based computers	512MB	1.5GB for x86-based computers; 2GB for Itanium-based computers	Minimum 8-way capable machine required; maximum 64
Windows Server 2003 Web Edition	133 MHz	128MB	1.5GB	Up to 2 CPUs

Beyond the minimum requirements, you will need to look at the features available in different versions and editions of Windows, and how they can be used to enhance network security. The progression from one version to another has offered improvements and additions to security, with Windows Server 2003 offering the most security features. By identifying which features are necessary for your organization, you can create a network that provides the necessary functionality and security.

Security Features

Windows 2000 offers a number of new security features that were not previously available in Windows NT. Many of the features we'll discuss next were implemented in Windows 2000 and have been updated in Windows Server 2003. In addition, new features have been added that make Windows Server 2003 the most secure Windows server product Microsoft has ever marketed.

Windows 2000 Server was the first version to provide encryption of data over the network and in the file system. IPsec allows encryption of data across the network. EFS uses a public key system to encrypt data on hard disks. Encryption ensures that unauthorized parties are unable to view the data if they gain access to it.

Windows 2000 was also the first version to provide built-in support for smart cards. *Smart cards* are generally the size of a credit card and have the ability to store data. When a smart card is inserted into a smart card device, it provides information that can be used for authentication and other purposes. With smart cards, the security of a network can be greatly enhanced because it is necessary to physically possess the card to log on.

A major advance that first appeared in Windows 2000 was Kerberos authentication. Kerberos version 5 is an industry-standard security protocol that uses mutual authentication to verify the identity of a user or computer, as well as the network service that is being accessed. In Windows 2000 Server and later, Kerberos is the default authentication service.

With Kerberos, each party to a transaction proves that they are who they claim to be through the use of *tickets*. A Kerberos ticket is encrypted data that is issued for authentication. Tickets are issued by a *Key Distribution Center* (KDC), which is a service that runs on every domain controller. When a user logs on, the user authenticates to AD using a password or smart card. Because the KDC is part of AD, 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 a ticket when a particular service is needed.

As mentioned earlier in this chapter, AD is a directory service that was first introduced in Windows 2000 Server. Because AD was not available when Windows NT 4 was released, it cannot be installed on a Windows NT server. Once AD is installed on Windows 2000 Server or Windows Server 2003, the server becomes a domain controller that can be used for authentication and management of user accounts and other objects in AD.

When AD is installed, a number of features and tools become available. There are three graphical tools that can be used with Windows 2000 Server or Windows Server 2003:

- **Active Directory Users and Computers** This utility allows you to administer user and computer accounts, groups, printers, OUs, contacts, and other objects stored in AD. Using this tool, you can create, delete, modify, move, organize, and set permissions on these objects.
- **Active Directory Domains and Trusts** This utility allows you to manage domains and the trust relationships between them. Using this tool, you can create, modify, and delete trust relationships; create and remove user principal name (UPN) suffixes; raise the domain mode (Windows 2000 Server only); and raise domain and forest functional levels (Windows Server 2003 only).
- **Active Directory Sites and Services** This utility allows you 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.



EXAM WARNING

Active Directory Users and Computers, Active Directory Domains and Trusts, and Active Directory Sites and Services are tools that are installed with AD. These tools are not available on servers that have not been configured as domain controllers. They are the primary tools for interacting with AD, and they allow you to configure different aspects of the directory.

A new feature in Windows Server 2003 is that AD 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 some of the attributes that are common to each of these objects. This makes it faster to manage users, because you do not need to make changes to one account at a time.

Windows Server 2003 AD also provides the ability to drag and drop objects into containers. To use this feature, select an object with your mouse, hold down your left mouse button to drag the object to another location (such as an OU), and release the button to drop the object into the container. This ability also makes it easy to add user and group objects to groups. Dragging and dropping a security principle's object (user, computer, or group) into a group adds it to the group membership.

In addition to these graphical tools, Windows Server 2003 also provides a number of command-line utilities for managing AD. Using these tools, you can perform management tasks through the textual interface of the command prompt. These tools allow administrators to manually enter commands to run operations from a command prompt or use the commands in batch files and scripts that can be scheduled to run at specific times.

Another new Windows Server 2003 feature is that domain controllers 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 drive on a server fails, you can use the backup to restore the data to a new drive and have the server up and running again. This same process can be used to restore AD to a new domain controller, so you do not need to replicate the entire directory across the network. Allowing domain controllers to be added to an existing domain through the use of backups is of great benefit when you are setting up a new domain controller across a slow WAN link from the nearest existing domain controller.

Functional Levels

When a Windows Server 2003 domain controller is created on a network, AD is installed with a basic set of features. Additional features can be enabled, depending on the operating systems running as domain controllers and the functional level that is configured for the domain or forest.



NOTE

Windows 2000 contained two modes: mixed and native. In Windows Server 2003, these are now called *functional levels*, but they remain unchanged. Just as Windows 2000 installed in mixed mode, Windows Server 2003 installs in the Windows 2000 mixed functional level. In Windows 2000, there was only one level of forest operation. Modes existed only at the domain level. With Windows Server 2003, there are domain functional levels and separate forest functional levels. In order to raise the forest functional level, the functional level of all domains in the forest must be set to the appropriate level.

Domain Functional Levels

The domain functional level determines which servers are supported in a domain and the features that are available in AD. When one or more Windows 2003 Server computers are installed on a domain, the domain functional level can be set for AD. At lower levels, older versions of Windows servers can still be used in the domain, but more advanced features for AD are sacrificed. At the highest level, only Windows 2003 Server machines can be used in the domain, and a full set of these advanced features become available. By not setting the domain functionality to an appropriate level, you may be forfeiting a number of the features you need for your network.

There are four different levels of functionality for AD:

- **Windows 2000 mixed** Allows domains to contain Windows NT Backup domain Controllers (BDCs) that can interact with the PDC emulator in a Windows Server 2003 AD domain. In this level, the basic features of AD are avail-

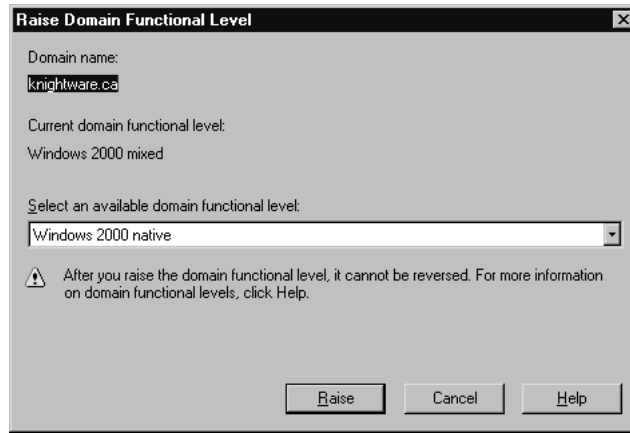
able. However, you cannot use additional group nesting, universal security groups, or security ID histories (SIDHistory) when moving accounts between domains. Because it accommodates the widest variety of domain controllers on your network, this is the default level of functionality when a Windows Server 2003 domain controller is installed.

- **Windows 2000 native** The highest mode available for Windows 2000 and the next highest level for Windows Server 2003 domain controllers. This functional level removes support for replication to Windows NT BDCs, so these older servers are unable to function as domain controllers. In this level, only Windows 2000 and Windows Server 2003 domain controllers can be used, and support for universal security groups, SIDHistory, and group nesting becomes available.
- **Windows Server 2003 interim** New in Windows Server 2003, this level is used when your domain consists of Windows NT and Windows Server 2003 domain controllers. 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 domain has never had (and will not have) Windows 2000 domain controllers, this is the level used for performing an upgrade.
- **Windows Server 2003** The highest functionality level for AD, this level is used when there are only Windows Server 2003 domain controllers in the domain. When this level is set for the domain, a number of additional features are enabled, which we'll discuss shortly.

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 allowed control of which operating systems were supported and the features that were available in AD. Windows 2000 mixed mode provides backward-compatibility with older operating systems like Windows NT 4, allowing Windows NT BDCs to still be used in a domain. Windows 2000 native mode restricted the domain to using only Windows 2000 Server machines on the network, and it provided an expanded feature set for AD. In Windows 2003 Server, these modes are now referred to as *functional levels*, and they allow Windows 2003 Server to provide backward-compatibility to domain controllers using these operating systems. In addition to these functional levels, Windows 2003 also introduces two new domain functional levels that were not available in the previous versions: Windows Server 2003 interim and Windows Server 2003.


The tool used to raise domain and forest functional levels is **Active Directory Domains and Trusts**. To raise a domain level, right-click the domain in the left console pane and click **Raise Domain Functional Level** in the context menu. The **Raise Domain Functional Level** dialog box appears, as shown in Figure 2.13. Select the functional level that you want, and then click **Raise**.

Figure 2.13 Raising the Domain Functional Level



When raising the domain functional level, it is important to remember that it is a one-way change. After raising the level, you cannot lower it. For example, if you raise the domain from Windows 2000 mixed to Windows Server 2003, you cannot return the level to Windows 2000 mixed again. This means that you cannot add Windows NT BDCs or Windows 2000 domain controllers to the domain after the upgrade. If you attempt to change the domain functional level after raising it to Windows Server 2003, a dialog box similar to the one shown in Figure 2.14 will be displayed.

Figure 2.14 Attempting to Change a Domain Functional Level After Raising the Functional Level



After all domain controllers are running Windows Server 2003 and the domain functional level has been raised to Windows Server 2003, new features are automatically available. One such feature is the domain controller renaming tool, which allows you to rename a domain controller without needing to demote it first. This can be useful when you need to restructure the network or simply wish to use a more meaningful name for a particular

domain controller. When you use this tool, AD and DNS entries for the renamed domain controller are automatically updated.



NOTE

You can also rename domains using the domain rename utility (`rendom.exe`). Using this tool, you can change the NetBIOS and DNS names of a domain, including any child, parent, domain tree, or forest root domains. By renaming domains, you can move them in the DNS hierarchy. For example, you can change the name of `dev.web.syngress.com` to `dev.syngress.com`, placing the `web.syngress.com` and `dev.syngress.com` domains on the same level of the hierarchy. You can even rename a domain so that it becomes part of a completely different domain tree. The only domain that you cannot reposition in this manner is the forest root domain.

The Windows Server 2003 domain functional level also provides a new attribute for user and computer accounts. The `lastLogonTimestamp` is added to user and computer objects, and it is replicated within the domain to all domain controllers, so that the last time these accounts were used to log on to the domain can be recorded. This way, a history of the user or computer account is created.

Another feature that becomes enabled when the domain functional level is raised is the ability to add a password to `InetOrgPerson` accounts. `InetOrgPerson` is an object class in AD that is used to create accounts that represent users in non-Microsoft directory services, and it is used in the same way as a user object. Other network operating systems, such as Novell NetWare, use their own implementations of a directory service, which are not always compatible with AD. `InetOrgPerson` is used to assist applications written for other directories or when migrating from these directory services to AD. Object classes are sets of attributes used to determine which attributes an object may have when it is created. Using the `InetOrgPerson` class, you can create a type of user account that is compatible with accounts from other directory services.

The features we've covered so far are only available in the Windows Server 2003 functional level. However, other features for the Windows Server 2003 level may also be available when lower functional levels are implemented. Windows 2000 native and Windows Server 2003 functional levels provide the ability to nest security and distribution groups in one another. Security groups are used to assign permissions and rights to groups of accounts, rather than modifying each account individually. Distribution groups are used to send bulk e-mail to large groups of users as a single entity. By nesting groups, one group can be added as a member of another group, saving the need to repeatedly add the same accounts to the membership of various groups.

Limited group nesting is available for domains running in Windows 2000 mixed mode. When this functional level is used, group nesting for distribution groups is allowed, but there is limited support for security groups. You can nest security groups only if you are

InetOrgPerson Object Class

The InetOrgPerson object class and the attributes it contains originate 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 the Lightweight Directory Access Protocol (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. The schema is part of AD and defines the classes of objects and the attributes that can be used in AD. In Windows Server 2003, an InetOrgPerson object class is included in the AD schema as a type of user class that can be used by LDAP applications that require this type of object and when migrating to AD from other directory services. This saves administrators from needing to extend the schema to create a new InetOrgPerson object class.

adding global groups to the membership of domain local security groups. Aside from this, nesting isn't permitted.

Another benefit of the Windows 2000 native or Windows Server 2003 functional level is that universal security groups can be used. (Domains that have the functional level set to Windows 2000 mixed do not allow universal security groups to be created.) Universal security groups can contain accounts and groups from any domain in the forest, and they can also be assigned permissions to resources in any domain in the forest. In this situation, the group can contain user accounts, global groups, and universal groups from any domain in the forest, and it can be assigned permissions to resources in any domain. Universal distribution groups can be used at any functional level, including Windows 2000 mixed.

In summary, some features are available but limited in the Windows 2000 mixed functional level. In other cases, however, support for a particular feature isn't available at all. Windows 2000 native or Windows Server 2003 functional levels provide the ability to convert groups. Each of these higher functional levels allows conversion between security groups and distribution groups. In addition, the Windows 2000 mixed functional level does not support SIDHistory, which allows user and computer accounts to be moved from one domain to another without affecting existing permissions. By failing to raise the functional level of a domain, you make several features unavailable to it.

Forest Functional Levels

In addition to the domain functional level, you can also set the functional level of a forest. A domain functional level is individually set for each domain. The forest functional level is set for the entire forest and thereby affects all domains within that forest. There are three different forest functional levels:

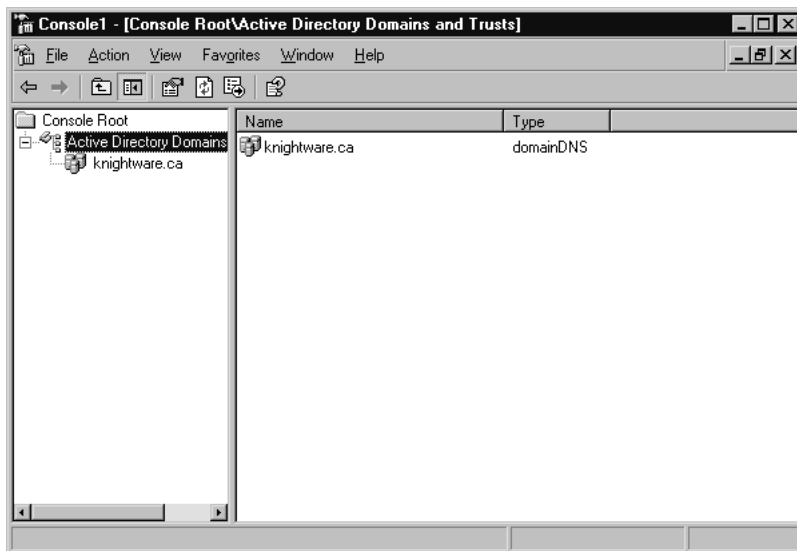
- Windows 2000
- Windows Server 2003 interim
- Windows Server 2003

By default, the functional level of a forest is set to Windows 2000. The Windows 2000 forest functional level allows Windows NT, Windows 2000, and Windows Server 2003 domain controllers on the network. However, it also provides fewer features than the higher functional levels. Elevating the functional level of a forest enables additional features. At the Windows Server 2003 interim level, domain controllers running Windows NT Server 4 and Windows Server 2003 can exist within the forest. This level is used when directly upgrading from Windows NT 4 to Windows Server 2003. When the default level is raised to Windows Server 2003, additional features in AD become available.

To raise the forest functional level, all domains in the forest must consist only of domain controllers running Windows Server 2003. In addition, the functional level of all domains must be set to Windows 2000 native or higher. After the functional level has been raised, all domains will have their functional level set at Windows Server 2003, even if it was set at Windows 2000 native prior to the forest level being elevated.

Like domain functional levels, forest functional levels are raised using **Active Directory Domains and Trusts**. As shown in Figure 2.15, this tool has an **Active Directory Domains and Trusts** node in the left pane. Right-click this node and click **Raise Forest Functional Level** in the context menu. You will see a dialog box that is similar to the one for raising the domain functional level (see Figure 2.14). Select the new functional level from the drop-down list, and then click **Raise** to complete the task.

Figure 2.15 Using Active Directory Domains and Trusts



As with domain functional levels, raising the forest functional level is a one-way change. After raising the level, you cannot lower it. Therefore, it is important that you decide which domain controllers exist on your network or may be added in the future prior to raising the level. If older operating systems are used for domain controllers in the forest, you will need to upgrade them before raising the level, and you will not be able to add these older systems after you make this change.

By raising the functional level to Windows Server 2003, new features become available to the forest. One such feature is the ability to create forest trusts. *Forest trusts* are one or two-way transitive trust relationships between two different forests. 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, meaning that both domains trust one another, users in each domain can access resources in the other domain. This expands the network, so users are able to use services and resources in both forests.



NOTE

Forest trusts are new in Windows Server 2003. They involve a great deal of complexity that does not exist in other trust relationships. It is important to note that when a forest trust exists, the Global Catalog for each forest remains separate. Much of the additional complexity stems from this fact. When a user who is logged on to a domain in one forest attempts to access resources in a domain located in the other forest, special pointers in the local forest's Global Catalog must be present. The default settings often allow for a free exchange of users in each direction the trust allows. For maximum security, these pointers should be manually configured by an administrator, so that only specific domains or resources on each side of the trust are accessible from across the trust.

To improve the performance of replication across the network, the Windows Server 2003 level allows linked value replication. To ensure that all domain controllers have a duplicate copy of AD, directory data is replicated between them. *Linked value replication* improves replication by having less information copied between domain controllers. Rather than treating the entire membership of a group as a single unit of replication, linked value replication allows individual members of groups to be replicated (instead of the entire group).

When the functional level is raised to Windows Server 2003, you can make additional modifications to the schema by disabling classes and attributes. When a particular type of object or an attribute is no longer needed in an object, the class or attributes within it can be deactivated. The ability to disable schema objects was available in Windows 2000, but Windows Server 2003 provides the ability to reactivate them again when needed. If schema objects are no longer required, you can deactivate them, and then reactivate them later if the situation changes. (Although classes and attributes can be disabled, they cannot be deleted.)

Now that we've discussed raising the domain and forest functional levels, let's look at the procedure for doing it. Exercise 2.2 will walk you through the process of raising both of these functional levels, so that all of the features discussed earlier are available for use.

EXERCISE 2.02

RAISING DOMAIN AND FOREST FUNCTIONALITY

The following steps should not be performed on a production network. This exercise assumes that all domain controllers 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. Select **Start | Administrative Tools | Active Directory Domains and Trusts**.
2. When **Active Directory Domains and Trusts** opens, expand the **Active Directory Domains and Trusts** node and select your domain.
3. Select **Action | Raise Domain Functional Level**.
4. In the **Raise Domain Functional Level** dialog box, select **Windows Server 2003** from the drop-down list, and then click the **Raise** button.
5. A warning message will appear, informing you that this action will affect the entire domain and 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. Select the **Active Directory Domains and Trusts** node.
8. Select **Action | Raise Forest Functional Level**.
9. In the **Raise Forest Functional Level** dialog box, select **Windows Server 2003** from the drop-down list, and then click the **Raise** button.
10. A warning message will appear, informing you that this action will affect the entire forest and cannot be reversed. Click **OK**.
11. After you raise the level, a message box will inform you that the action was successful. Click **OK**, and then exit the utility.

EXAM
70-293
OBJECTIVE
1.4.1

Identifying Minimum Security Requirements for Your Organization

Before you can begin implementing security measures, you need to know what needs protecting. Different organizations have different needs, so the systems and data that are essential to one company may be superfluous to another. For this reason, the security planning process involves considerable analysis. You need to determine which risks could threaten a company, what impact these threats would have on the company, the assets that the company needs to function, and what can be done to minimize or remove a potential threat.

Risk is the possibility of experiencing some form of loss. This isn't to say that a risk will become a real problem, only that it has the potential of happening. To address risks, you need to determine which events and factors in an organization are potential threats, and then devise ways to deal with them before they become actual problems. There are many different risks that can affect an organization, and the types of risks will often vary from business to business. The following are the main types of threats:

- Environmental threats, such as natural and man-made disasters
- Deliberate threats, where a threat was intentionally caused
- Accidental threats, where a threat was unintentionally caused

Environmental threats can be natural disasters, such as storms, floods, fires, earthquakes, tornadoes, and other acts of nature. The types of disasters that can occur generally vary from one geographical region to another. For example, a business in California might be more prone to earthquakes, while an organization in Canada might be at risk of severe snowstorms. When dealing with this type of disaster, it is important to analyze the entire company's risks, considering any branch offices located in different areas that may be prone to different natural disasters.

Human intervention can create problems as devastating as any natural disaster. Man-made disasters can also occur when someone creates an event that has an adverse impact on the company's environment. For example, faulty wiring can cause a fire or power outage. In the same way, a company could be impacted by equipment failures, such as the air conditioning breaking down in the server room, a critical system failing, or any number of other problems.

The deliberate threat type is one that has appeared numerous times in the news over the last number of years. These types of threats result from malicious persons or programs, and they can include potential risks such as hackers, viruses, Trojan horses, and various other attacks that can damage data and equipment or disrupt services. This type of threat can also include disgruntled employees who have authorized access to such assets and have the ability to harm the company from within.

Many times, internal risks are not malicious in nature, but accidental. Employees can accidentally delete a file, modify information with erroneous data, or make other mistakes

that cause some form of loss. Because people are fallible by nature, this type of risk is one of the most common.

Each business must identify the risks it may be in danger of confronting and determine what assets will be affected by a potential problem. Assets are property and resources that have value to the company, and they can include the following:

- **Hardware** Servers, workstations, hubs, printers, and other equipment.
- **Software** Including commercial software (which is purchased off the shelf) and in-house software (which is developed by programmers working for the company).
- **Data** Including documents, databases, and other files needed by the business.
- **Personnel** Employees who perform necessary tasks in the company (for example, the network administrator who knows how to restore damaged systems from a backup).
- **Sundry equipment** Office supplies, furniture, tools, and other assets needed for the business to function properly.
- **Facilities** The physical building and its components.

As you can see, any number of risks could result in the loss of a wide variety of assets. For example, a fire could destroy a building, including the facilities containing servers that store critical software and data. It might also injure key personnel who are necessary for the business to function. With one disaster, an entire company can be crippled.

When identifying minimum security requirements, it is important to determine the value and importance of assets, so you know which are vital to the company's ability to function. You can then prioritize risk, so that you can protect the most important assets of the company and implement security measures to prevent or minimize potential threats.



TEST DAY TIP

Questions dealing with identifying the minimum security requirements will be mixed with issues directly related to Windows Server 2003. They will test your knowledge by matching the minimum requirements shown in a scenario against the features and functionality of Windows Server 2003.

Determining the value and importance of assets can be achieved in a number of ways. Keeping an inventory of assets owned by the company will allow you to identify the equipment, software, and other property owned by the company. By referring to this list, you can see the possessions of the business, and you can update this list to reflect the current monetary value. For example, you could see that a new server has specific software and hardware installed on it, and it would cost a specific amount to replace. In the same light, an older

server may cost less to replace, but contain sensitive data that makes it valuable to the company.

To determine the importance of data and other assets, and thereby determine what is vital to secure, you can meet with department heads. Doing so will help you to identify the data and resources that are necessary for people in each department to perform their jobs. For example, a Human Resources department might specify that a particular database is used for critical information and that specific software is needed. At the same time, you could find that a folder on the server contains a Web site used by employees to advertise items for sale, upcoming events, and other material that is not work-related. By discovering this information, you make great strides in being able to protect the work-related material while expending little to no effort in preserving non-work-related files. You will also probably find that you need the assistance of the Accounting department to determine appropriate values.

In addition to interviewing different members of an organization, review the corporate policies for specifications of minimum security requirements. For example, a company may have a security policy stating that all data is to be stored in specific folders on the server, and that the IT staff is required to back up this data nightly. Such policies may not only provide insight on what is to be protected, but also what procedures must be followed to provide this protection.

Organizational policies are not the only method of acquiring information on security requirements. Companies may be required to protect specific assets by law or to adhere to certain certification standards. For example, hospitals are required to provide a reasonable level of security to protect patient records. This may include implementing firewalls to prevent hackers from accessing these files through the Internet. If such requirements are not met, an organization can be subject to legal action.

Identifying Configurations to Satisfy Security Requirements

To protect assets from risks that were identified as possible threats to a business, countermeasures must be implemented. Servers will need certain configurations to provide security, and plans must be put into practice. By applying methods to protect assets, the potential loss can be minimized or removed, and the security requirements of a business can be met.


Compare the risks faced by an organization with an operating system's features to find support that will address certain threats. Configuring the server to use these services or tools can assist in dealing with potential problems. For example, installing AD and using domain controllers on a network can heighten security and provide the ability to control user access and security across the network. In the same way, configuring a file server to use EFS so that data on the server's hard disk is encrypted can augment file security. Using security features in an operating system allows you to minimize many potential threats.

The same technique should be used when determining which roles will be configured on servers. As described earlier in this chapter, different server roles provide different ser-

vices to a network. By comparing the functionality of a server role to the needs of a company, you can identify which roles are required. For example, if you need secure communication and transmission of data on a network, configuring IPSec will be a viable solution. Similarly, configuring servers to be DNS or WINS servers will provide name resolution. Domain controllers allow you to benefit from AD. By understanding what people need in your organization, you can determine which server roles must be configured.

Although it may be tempting to configure a server with every possible role, this can cause problems. When a server is configured to play a certain role in an organization, a number of different services, tools, and technologies may be installed and enabled. Because there is a possibility these may be exploited, you should avoid this risk by never installing more roles than are needed. Always disable any unneeded services on the server.

Although roles are helpful, running a wizard to configure servers in a particular role isn't enough to create a secure environment. Additional steps should be followed to protect these servers and the data, applications, and other resources they provide. By customizing servers in this manner, you can ensure that the company will be able to benefit from Windows Server 2003 without compromising security. We'll discuss these steps in the "Customizing Server Security" section later in this chapter.



Planning Baseline Security

Security templates allow you to apply security settings to machines. These templates provide a baseline for analyzing security. Templates are .inf files that can be applied to computers manually or by using Group Policy Objects (GPOs).

Security Templates and Tools

There are numerous settings, or customizable security policies, that you can apply through security templates, including the following:

- **Account Policies** Include password policies, Kerberos policies, and account lockout policies.
- **Local Policies** Include user rights, audit policies, and other security options.
- **Event Log** Include configuration options for the Application, System, and Security event logs that can be viewed through Event Viewer.
- **Restricted Groups** Used to specify group memberships.
- **System Services** Used to configure permissions and startup options for services.
- **Registry** Used to specify permissions and for auditing Registry objects.
- **File System** Used to specify permissions and for auditing files and folders.

You can create and edit security templates using the Security Templates snap-in for the Microsoft Management Console (MMC), as explained in the "Creating Custom Security

Templates” section later in this chapter. This tool allows you to manage your own templates, but you can also use predefined templates that come with Windows Server 2003. The next sections describe the predefined templates and the tools for working with security settings.

Predefined Templates

The Windows Server 2003 predefined templates are located in the `%systemroot%/Security\Templates` directory. The following templates are available:

- **compatws.inf** Relaxes security settings on a workstation or server, so that otherwise incompatible applications have a chance of working.
- **DC security.inf** Contains the default security settings for a domain controller.
- **hisecdc.inf** Contains high-level security settings for domain controllers.
- **hisecls.inf** Contains high-level security settings for workstations.
- **rootsec.inf** Contains the default security settings for the system volume (`%systemdrive%`).
- **iesaccls.inf** Contains settings to lock down Internet Explorer.
- **securedc.inf** Contains enhanced security settings for domain controllers.
- **securews.inf** Contains enhanced security settings for workstations.
- **setup security.inf** Contains the default security settings for a default installation of Windows Server 2003.

These templates are described in more detail in the following sections.

Compatws Template

The `compatws` template is used to provide users with access to applications that do not function properly with full system security in place. The `compatws` template relaxes user permissions so that programs are more likely to run without errors. It also removes any members of the Power Users group. Many administrators solve their application problems by adding users to the Power Users group. However, members of this group also have the ability to create users, groups, shares, and printers. Overall, this template erodes system security and should be used with caution.

DC Security Template

The DC security template is created when a server is first promoted to being a domain controller. It contains a number of default settings, including settings for the file system, Registry, and system services. This template allows you to reapply these default security settings. Registry keys and system services that have been added or modified since the initial installation may be overwritten, as may permissions on new files. Therefore, considerable planning should be done before applying this template to a domain controller in your network.

Hisecdc Template

The hisecdc template is used to apply high-level security settings to a domain controller. Using this template will cause the domain controller to require encrypted authentication. Using this setting will also prevent most pre-Windows 2000 computers from being able to communicate with the server, because the domain controller will require clients to communicate using NTLM version 2 (NTLMv2). Finally, this template will cause many applications to malfunction.

Hisecws Template

The hisecws template applies settings similar to those in the hisecdc template, but it is designed for use with workstations and servers that are not configured as domain controllers. When this template is applied to a computer, all of the domain controllers that have accounts for users that can log on to the client must be running Windows NT 4.0 Server with Service Pack 4 installed, Windows 2000 Server, or Windows Server 2003. Also, any domain controllers in domains that the client is a member of must be running Windows 2000 Server or Windows Server 2003.

Clients are also unable to connect to computers using LAN Manager for authentication or from machines running operating systems earlier than Windows NT 4.0 Service Pack 4 using an account on the local machine. In addition, attempts to connect to a server running Windows NT 4 where the time on each machine has a difference of 30 minutes or more will fail. If the client connects to a computer running Windows XP, the time difference between them cannot exceed 36 hours.

The hisecws template also modifies settings to control memberships in security-sensitive groups. Once applied, all users are removed from the Power Users group, and only members of the Domain Admins group and the Administrator account are kept as members of the computer's local Administrators group.

As with the hisecdc template, applying the hisecws template will cause many applications to malfunction because of the enhanced security. This template should be very carefully tested before deployment.

Rootsec Template

The rootsec template is used to define security settings for the system volume. It is used to set permissions at the root of the system drive, so that original settings can be reapplied. This can be particularly useful if the permissions on the system drive are inadvertently modified. This template can also be modified to apply the same root permissions on other volumes. In doing so, it will overwrite inherited permissions on child objects, but will not overwrite any explicit permissions on child objects.

IesacIs Template

The iesacIs template is used to lock down security settings used by Internet Explorer (IE), which can be used to access data on the Internet or on a corporate intranet. Using this template, you can enhance security by enforcing stricter settings on Internet Explorer.

Securedc Template

The `securedc` template is used on domain controllers to enhance security while minimizing the impact on applications. This template also configures servers to refuse LAN Manager responses. Computers running operating systems such as Windows for Workgroups, Windows 95, and Windows 98 use LAN Manager to authenticate to servers. For these clients to be able to connect to a domain controller with the `securedc` template applied, the clients will need to have a patch or the Active Directory Client Extensions Pack installed on them.

Securews Template

The `securews` template provides the same settings as the `securedc` template, but it applies to workstations or servers that are not configured as domain controllers. It is designed to enhance security without impacting on applications that are running on the computer. This template also affects authentication, because it limits the use of NTLM by configuring clients accessing the machine to respond with NTLMv2 responses.

When this template is applied, the domain controllers that contain user accounts for those who will log on to the client must run Windows NT 4.0 with Service Pack 4 or higher, Windows 2000, or Windows Server 2003. Additionally, there are requirements dealing with time. If the domain contains Windows NT 4 domain controllers, the clocks between the domain controllers running this operating system must have their time synchronized within 30 minutes of one another. Computers also will not be able to connect to servers running Windows 2000 or Windows NT 4 if their clocks are off by more than 30 minutes from the server. Computers will not be able to connect to a Windows XP machine if their clocks are off by more than 20 hours.

Servers that have this template applied to it also have limitations. The server won't be able to connect to clients running LAN Manager and will need to be authenticated using NTLMv2. However, NTLMv2 can be used to authenticate to Windows 2000 or Windows Server 2003 servers if the clocks on the client and server are within 30 minutes of one another. If the server is running Windows XP, the two machines must be synchronized within 20 hours of one another.

Setup Security Template

The setup security template is created when a computer is installed, and it varies from one machine to another, depending on whether its operating system was upgraded or a clean installation. Because of this, it should never be applied to a group of computers using Group Policy or manually to other systems, unless you have carefully reviewed its settings. This template allows you to reapply a system's default security settings. Use the DC security template for domain controllers, not the setup security template.

Security Configuration and Analysis


A tool that makes significant use of security templates is the Security Configuration and Analysis tool. This tool is an MMC snap-in that allows you to analyze and configure system settings. Using it, you can perform the following tasks:

- Analyze security settings for local and group policies.
- Apply security templates to the local Windows Server 2003 computer.
- Export settings to template files, so they can be applied later either manually or by using Group Policy.

The Security Configuration and Analysis tool assists you in determining whether a computer has an adequate security configuration by comparing the current settings to those in a security template. One or more templates are applied to a database, which is used to analyze the difference between the database settings and the current computer configuration. In viewing the results, you are able to determine what changes will be made to the machine if the template is applied. You can alter the settings to ensure that the desired configuration results are obtained, and apply them to the computer individually or to a range of computers using a GPO.

When the Security Configuration and Analysis snap-in is loaded into MMC, the console tree in the left pane shows the Security Configuration and Analysis node, as shown in Figure 2.16. When you initially select this node, it will provide information in the details pane (right pane) on how to open or create a database that can be used to analyze or configure the computer. After you have opened or created a database, the left pane is populated with a log or nodes containing settings that can be configured. You can then select any of these nodes and modify settings that can be applied to the local machine or multiple machines using Group Policy.

Figure 2.16 Initial Information Provided by the Security Configuration and Analysis Tool

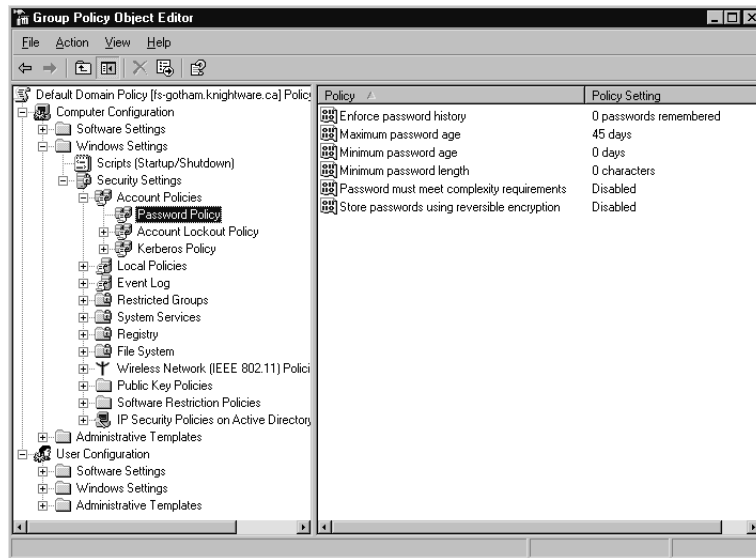


Group Policy Object Editor

The Group Policy Object Editor is another tool that allows you to view and modify security settings. Because this tool is also loaded into the MMC, it has the same basic appearance as the Security Configuration and Analysis snap-in. A tree of information appears in the left pane, and details on selected items appear in the right pane. GPOs can be applied to manage security settings at the OU, site, and domain level.

As shown in Figure 2.17, security settings are available under **Computer Configuration** and **User Configuration**. The settings under **Computer Configuration** apply to settings that affect the computer, and those under **User Configuration** apply to users. The policies that appear in this snap-in are those that have already been configured in the GPO.

Figure 2.17 Configured Policies in the Group Policy Object Editor



Using the Group Policy Object Editor, you can import policies stored in templates or export current settings to a template file that can then be used to configure other computers. These are topics we’ll discuss later in this chapter, in the “Enforcing Default Security Settings on New Computers” section.

Secedit

Secedit is a command-line tool that allows you to analyze and configure computers using templates, and to automate security configurations. Commands are entered from the textual interface of the command prompt, which means that these commands can be added to scripts and batch files to automatically configure a machine. Unlike the other tools we’ve discussed so far, Secedit cannot be used to modify or export a template.

There are several commands that can be used with Secedit to specify which actions to perform. The different parameters for Secedit include the following:

- **secedit /analyze** Used to analyze the security settings of a computer.
- **secedit /configure** Used to apply the security settings in a template to a computer.
- **secedit /export** Allows you to export the security settings in the database to a template.
- **secedit /import** Used to import a template into the database so that its settings can be used to analyze the machine or to configure its security settings.
- **secedit /validate** Used to validate the syntax of a template before importing it into the database.
- **secedit /GenerateRollback** Used to create a rollback template that can be used to restore the computer's security settings to the way they were before applying a configuration template.



EXAM WARNING

The Secedit command-line tool and Security Configuration and Analysis snap-in are the only tools that allow you to analyze security settings by having them compared to a security template. No other tools in Windows Server 2003 have this ability.

The following sections describe each of these commands and their parameters in more detail.

Analyze

The `secedit /analyze` command provides the ability to compare the security settings in a template to those of a computer. The syntax for this command is as follows:

```
secedit /analyze /db FileName.sdb [/cfg FileName] [/overwrite] [/log
  FileName] [/quiet]
```

As is the case with each of the Secedit commands, the command's parameters allow you to specify additional options. The parameters for `secedit /analyze` are shown in Table 2.2.

Table 2.2 Parameters for the `secedit /analyze` Command

Parameter	Description
<code>/db FileName.sdb</code>	Used to specify the database that is used in performing the operation.
<code>/cfg FileName</code>	Used to specify the security template that is to be imported into the database and used for the operation.
<code>/overwrite</code>	Used to specify that the database is to be emptied before the security template is imported into it. When this setting isn't used, security templates are accumulated in the database, so that multiple templates can be used in the process. Any conflicting settings existing in the database are overwritten as the next template is imported.
<code>/log FileName</code>	Specifies the log file used to record events related to the command. By default, if this parameter isn't specified, events will be logged to <code>%windir%\Security\Logs\scesrv.log</code> .
<code>/quiet</code>	Ensures that the user is not prompted for input during the process.

Configure

The `secedit /configure` command is used for configuring security settings on a computer, by applying the settings in a database to the machine. With this command, the template can be imported into a database and applied to the local machine. The syntax for this command is as follows.

```
secedit /configure /db FileName.sdb [/cfg FileName ] [/overwrite][/areas
  Area1 Area2 ...] [/log FileName] [/quiet]
```

The command's parameters are the same as those listed in Table 2.2, with the addition of `/areas Area1 Area2`. This parameter is used to specify what security settings are exported to the template. When this parameter is used, security areas can be specified. When it isn't used, all settings are exported. The following security areas can be specified:

- **SECURITYPOLICY** Includes account and audit policies, event log settings, and security options.
- **GROUP_MGMT** Includes settings for restricted groups.
- **USER_RIGHTS** Includes settings for user rights assignments.
- **REGKEYS** Sets Registry permissions.
- **FILESTORE** Sets file system permissions.
- **SERVICES** Includes system service settings.

Export

The `secedit /export` command allows you to export settings to a template. Using this command, you can take the settings from a computer, export it to a template, and then import it to another machine or GPO so that multiple computers now share the same configuration. The syntax for this command is as follows:

```
secedit /export /db FileName.sdb [/mergedpolicy] [/cfg FileName ]
        [/areas Area1 Area2 ...] [/log FileName] [/quiet]
```

The command's parameters are the same as those listed in Table 2.2, with the addition of **/areas Area1 Area2**, explained in the previous section, and **/mergedpolicy**, which is used to merge the security settings of the domain and local computer into a single template file.

Import

The `secedit /import` command is used to import a security template into a database, so it can be applied to the computer or used in analysis. The syntax for this command is as follows:

```
secedit /import /db FileName.sdb /cfg FileName [/overwrite]
        [/areas Area1 Area2 ...] [/log FileName] [/quiet]
```

The command's parameters are the same as those listed in Table 2.2, with the addition of **/areas Area1 Area2**, described earlier in the “Configure” section.

Validate

The `secedit /validate` command is used to validate the syntax of a template before importing it into the database. This command is particularly useful when you've created a new security template and want to ensure that it does not have errors before using it for configuration or analysis. The syntax for this command is as follows:

```
secedit /validate FileName
```

Unlike the other commands we've discussed, this command has only one parameter: *FileName*. The *FileName* parameter is used to specify the name of the template to be validated.

GenerateRollback

When applying a configuration template to a machine, the `secedit /GenerateRollback` command provides the option of creating a template that can be used to roll back settings on the machine. Before a security template is applied, the current settings of the computer are exported into a template file. If you wish to restore the old settings of the computer after the security template is applied, you can use the rollback template. The syntax for this command is as follows:

```
secedit /GenerateRollback /cfg FileName.inf /rbk FileName.inf [/log
    FileName] [/quiet]
```

Table 2.3 describes these parameters.

Table 2.3 Parameters for the secedit /GenerateRollback Command

Parameter	Description
/cfg <i>FileName.inf</i>	Used to specify the security template that will be used in creating the rollback template.
/rbk <i>FileName.inf</i>	Used to specify the name of the rollback template to be created.
/log <i>FileName</i>	Specifies the log file used to record events related to the command. By default, if this parameter isn't specified, events will be logged to %windir%\Security\Logs\scesrv.log.
/quiet	Ensures that the user is not prompted for input during the process.

Planning Secure Baseline Installation Parameters

Because applying a security template can have a major impact on a computer, it is important that you take preliminary steps to ensure that the template can be applied correctly and will not make unwanted changes. By reviewing information about the template and performing an analysis of changes that will be made after the template is applied, you can ensure the computer will be configured correctly.

Before applying a security template, you should review its settings. Each of the templates addresses different levels of security and/or different settings that will be applied to the computer. Although template settings can be customized, you should determine whether a particular template configures the computer the way you want. If the wrong settings are applied, you need to either manually correct them or use a rollback template that was created before you applied this template.

The only predefined templates that will return a computer to an original state are the setup security and DC security templates. As we discussed earlier, the setup security template contains settings from when the computer was installed, and it is specifically created for each computer. This template can be used on workstations, stand-alone servers, and member servers, but domain controllers should not have this template applied to them. To return a domain controller to the state it was in when it was first promoted, use the DC security template. In both cases, any changes that have been made to settings since the template was initially created are not applied.

Using Security Configuration and Analysis to Analyze a Computer


By analyzing a computer with Security Configuration and Analysis, you can determine whether a machine has adequate security settings or if additional configuration is required.

The analysis is performed by adding one or more security templates to a database, which is used for comparison against the computer's current settings. In comparing this information, you can see where possible problems exist between your current configuration and the ones stored in the template.

Analyzing a computer begins by opening the MMC with the Security Configuration and Analysis snap-in installed. Then you can analyze a computer by performing the following steps:

1. In the left pane of the console, right-click **Security Configuration and Analysis** (see Figure 2.16) and select **Open Database** from the context menu. (Note that the context menu options also appear on the **Action** menu when **Security Configuration and Analysis** is selected.)
2. The **Open database** dialog box, shown in Figure 2.18, lists all the existing databases. To open an existing database, select the database from the list and click **Open**. To create a new database instead, enter the name of the new database in the **File name** text box, and then click **Open**. If you are opening an existing database, you will then be returned to the **Security Configuration and Analysis** tool, and you can skip to step 4. If you are creating a new one, the **Import Template** dialog box appears.

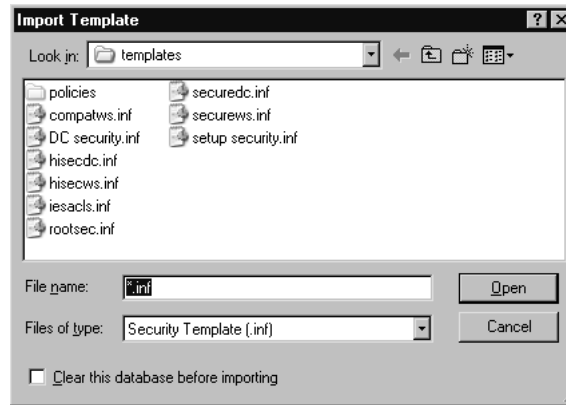
Figure 2.18 Opening an Existing Database or Creating a New One



3. As shown in Figure 2.19, the **Import Template** dialog box displays a list of the security templates stored in the `%systemroot%\Security\Templates` folder. This folder contains predefined security templates, but you can browse the hard disk for other security templates that you've created or downloaded and stored else-

where. Select a template from the list and click **Open**. The template is imported into the database, and you're returned to the **Security Configuration and Analysis** tool.


Figure 2.19 Importing a Template



You can add more templates by right-clicking the **Security Configuration and Analysis** node again and selecting **Import Template** from the context menu. When multiple templates are added to the database used for analysis, the templates are merged together so that all settings are used for comparison. These templates are added one at a time, and any conflicts between them are resolved by the order in which they are imported. For example, if you added the *compatws* template and then the *securews* template to the database, the settings in the *securews* template would take precedence because it was the last one to be imported. If another template is then added and conflicts with the current composite template in the database, this new template's settings would take precedence over the previous settings. To import a template into the database without having it appended to existing settings, check the **Clear this database before importing** check box in the **Import Template** dialog box. Any existing settings in the database will be purged, and only the settings in the template being imported will be used.


4. After you've opened or created a database and added the necessary templates, you are ready to begin taking steps to analyze the existing security settings. Select the **Security Configuration and Analysis** node, right-click it, and click **Analyze Computer Now**.
5. As shown in Figure 2.20, the **Perform Analysis** dialog box appears. Here, you can enter the name and path of a log file that will be used to record errors in the process. After clicking **OK**, another dialog box informs you that analysis of the computer is being performed.

Figure 2.20 Entering the Analysis Log File Path



- When the analysis is complete, the left pane of the **Security Configuration and Analysis** tool is populated with information about the settings that have been analyzed. As shown in Figure 2.21, the left pane shows different areas of security. When selected, these display results of the analysis for that area in the right pane. A side-by-side comparison is offered, showing database settings used for analysis and the computer's current settings. This allows you to quickly determine if changes need to be made to the current settings or if they provide the level of security desired for your organization.

Figure 2.21 Viewing the Results of a Security Analysis



When an analysis is performed, the results are organized into areas of security, and visual flags are used to indicate discrepancies. The following flags may appear beside entries in the results:

- A red X indicates that the entry does not match the corresponding setting in the database.
- A green check mark indicates that the entry in the database and the computer's setting match.
- An exclamation mark indicates that an entry in the database does not correspond to any setting on the computer. This may appear if a security setting for a group or other object is in a template added to the database, but the group or object isn't one that is used on the computer being analyzed.
- A question mark indicates that although the setting is on the computer, there is no corresponding entry in the database. This may indicate that the account you are using when performing the analysis does not have the appropriate permissions to analyze a security area or object, or that the entry was not used in any of the templates added to the database.
- No highlight indicates that the entry isn't defined in the database and isn't used on the system.

To modify settings in the database, double-click an entry. For example, double-clicking the **Maximum password age** entry brings up a corresponding dialog box, which allows you to change the number of days before a password will expire. Once you're finished making the modifications, you can save these changes to a new template file by selecting the **Security Configuration and Analysis** node and clicking **Action | Export Template**. In the **Export Template To** dialog box, shown in Figure 2.22, you can specify the name of the new template and where it should be saved. As you'll see in the next section, you can then use your new template to apply the settings to the computer and other machines on your network.

Figure 2.22 Exporting a Template



EXERCISE 2.03

ANALYZING SECURITY USING SECURITY CONFIGURATION AND ANALYSIS

1. Select **Start | Run**, type **MMC**, and click **OK**.
 2. In the blank console that appears, click **File | Add/Remove Snap-in**.
 3. When the **Add/Remove Snap-in** dialog box appears, click the **Standalone** tab, and then click the **Add** button.
 4. In the **Add Standalone Snap-in** dialog box, select **Security Configuration and Analysis** from the list and click **Add**.
 5. Click **Close** to return to the previous screen. The **Security Configuration and Analysis** entry should appear in the **Add/Remove snap-in** dialog box. Click **OK** to close the dialog box.
 6. The console tree in MMC should now contain a **Security Configuration and Analysis** node. Select this node and click **Action | Open Database**.
 7. When the **Open database** dialog box appears, type the name of a new database in the **File name** text box and click **Open**.
 8. When the **Import Template** dialog box appears, select **hisecdc** if you are working on a domain controller, or select **hisecws** if you are working on a workstation or server that isn't configured as a domain controller. Then click **Open**.
 9. When the **Security Configuration and Analysis** console appears, select the **Security Configuration and Analysis** node in the left pane and click **Action | Analyze Computer Now**.
 10. When the **Perform Analysis** dialog box appears, click **OK** to accept the default path and filename for the error log to be created.
 11. When the analysis is complete, browse through the settings and identify differences between the security settings in the database and the machine.
-



Enforcing Default Security Settings on New Computers

Security settings can be enforced on local computers or through AD. By using security templates in conjunction with the Security Configuration and Analysis snap-in, you can configure a local computer's security settings. Security templates can also be imported into the group policy of a domain, site, or OU in AD, so that the settings can be applied to multiple computers.

Using Security Configuration and Analysis to Apply Templates a Local Computer

The Security Configuration and Analysis tool allows you to configure local computers by applying the settings in a security template to the local policy. The settings will apply only to the computer on which Security Configuration and Analysis is being run. They will not affect other machines in the domain.

The initial steps for configuring a local computer are similar to the steps involved in running an analysis. In the Security Configuration and Analysis console, select the **Security Configuration and Analysis** node in the left pane and click **Action | Open Database**. As described earlier in the “Using Security Configuration and Analysis to Analyze a Computer” section, use the **Open database** dialog box (see Figure 2.18) to either open an existing database or create a new one. If you are opening an existing database, you will be returned to the **Security Configuration and Analysis** tool. If you are creating a new database, the **Import Template** dialog box (see Figure 2.19) appears. In the **Import Template** dialog box, select the security template that will be applied to the local machine and click **Open**. The template is imported into the database, and you're returned to the **Security Configuration and Analysis** tool. You can add other templates by selecting the **Security Configuration and Analysis** node again and clicking **Action | Import Template**. Check the **Clear this database before importing** check box if you want only the settings in the template being imported to be used in the database.

After you've added the templates to the database, you return to the **Security Configuration and Analysis** tool. You can apply the template by selecting the **Security Configuration and Analysis** node again and clicking **Action | Configure Computer Now**. In the dialog box that appears (see Figure 2.20), specify the filename and path of the error log file created for this process. Clicking **OK** in this dialog box will begin the configuration of the computer.

Using Group Policy Object Editor to Apply Templates

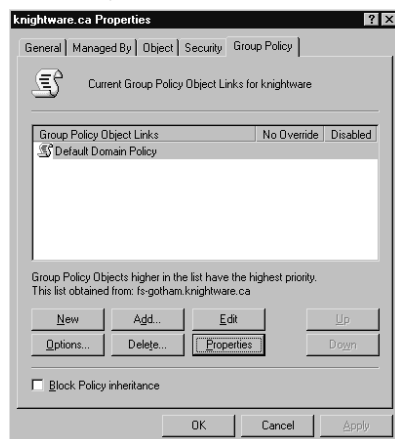
AD allows security templates to be applied at the domain, site, and OU level by using GPOs. When a security template is imported into a GPO, any computers that have the GPO applied to them will automatically receive the configured settings. The Group Policy Object Editor tool allows you to view and modify settings in a GPO.

You can view and modify the group policies of domains, sites, and OUs using tools that are installed on domain controllers. You can access the group policy configuration of a site through Active Directory Sites and Services. To access domain and OU settings, use **Active Directory Users and Computers**. By selecting a site in **Active Directory Sites and Services** and clicking **Action | Properties**, you can access the group policy configuration of that site. To see the group policy settings of a domain or OU, select it in **Active Directory Users and Computers**, and then click **Action | Properties**.

As shown in Figure 2.23, the **Group Policy** tab of a domain, site, or OU **Properties** dialog box allows you to view linked group policies. This tab includes a list of the group policies that are currently linked to this domain. Beneath the list are the following buttons for working with the GPO:

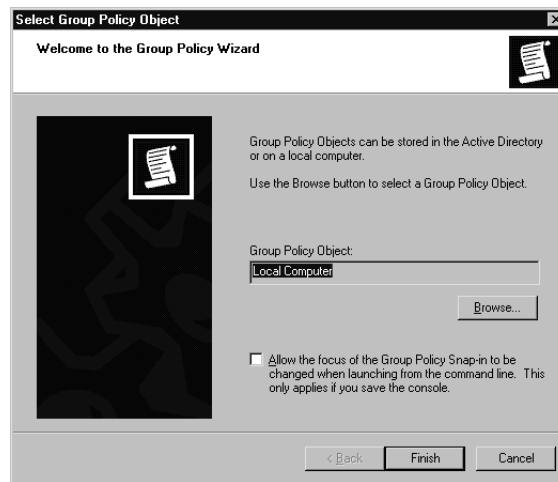
- **New** Allows you to create a new GPO.
- **Add** Allows you to link an existing group policy to the domain, site, or OU.
- **Edit** Displays the Group Policy Object Editor, which can be used to configure the GPO.
- **Options** Displays a dialog box containing two options for the GPO. The **No Override** option specifies that group policies lower in the hierarchy cannot override the settings in this policy. The **Disable** option specifies that settings in this group policy are not to be applied.
- **Delete** Removes a selected group policy from the domain, site, or OU. There are two options. The **Remove the link from the list** option removes the link so it no longer appears in the listing. The **Remove the link and delete the Group Policy Object permanently** option removes the link so it no longer appears in the listing and also deletes it so it cannot be used in the future.
- **Properties** Displays properties of the group policy. You can configure permissions associated with a selected GPO and see where else it may be linked.

Figure 2.23 Viewing Group Policy Properties of a Domain



To open the Group Policy Object Editor, click the **Edit** button on the **Group Policy** tab. You can also open this tool using the MMC, by adding the Group Policy Object Editor snap-in. After you've added this snap-in, you are prompted to choose whether you want to open the local computer policy or browse for a group policy in AD, as shown in Figure 2.24. If the default choice of opening the local computer policy is used, any modifications you make will apply only to the computer on which you are working. Remember any local policy settings you configure can be overridden by a group policy applied at the site, domain, or OU level.

Figure 2.24 Selecting a Group Policy



As shown in Figure 2.25, the Group Policy Object Editor has two panes. The left pane contains a tree view that allows you to browse through various policy settings. This tree is divided into two separate sections: **Computer Configuration** (which applies to computer accounts) and **User Configuration** (which applies to user accounts). Located beneath each of these is a **Windows Settings | Security Settings** node, which contains groups of settings that you can view and modify. When you select a node in the left pane, policy settings appear in the right pane. When you double-click one of these policy settings, you'll see a dialog box that allows you to modify the entry. Each entry has different values that you can set.

Figure 2.25 Group Policy Object Editor

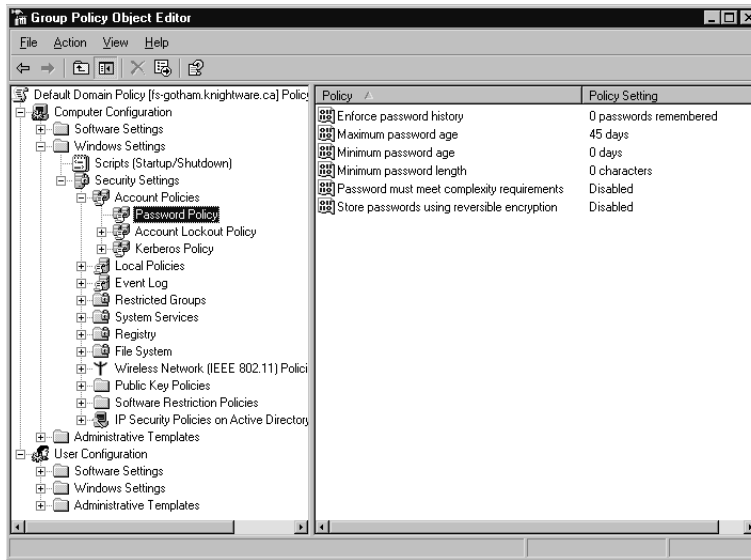
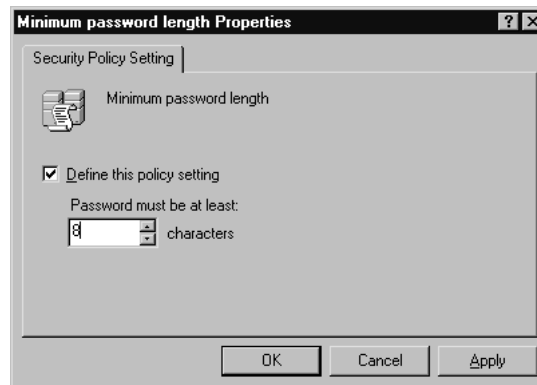


Figure 2.26 shows the **Minimum password length Properties** dialog box. Notice the **Define this policy setting** check box, which is common to all of the policies in the Group Policy Object Editor tool. If you check this option, you can then modify the value associated with that policy.

Figure 2.26 Viewing Minimum Password Length Properties



You can also import security templates into policies that are viewed through the Group Policy Object Editor. Right-click the **Security Settings** node and select **Import Policy** in the context menu. You will see a dialog box that displays the default directory for predefined templates. If necessary, browse to and select a template, and then click **Open** to import the template into the policy.

EXAM
70-293
OBJECTIVE
1

Customizing Server Security

Security templates contain predefined configurations, which are a great starting point, but usually, they do not fulfill the needs of many organizations. You may need to make some changes to match the organizational policies of your company. Similarly, configuring roles for servers requires additional steps to make the servers secure from attacks, accidents, and other possible problems. By customizing server security, you can implement security measures that will fulfill the unique needs of your organization.

Because every organization is different, the security needs of one company may vary from those of another. Before making security changes, you should consult corporate policies, as well as any requirements pertaining to organizational certifications or relevant laws. You should also use the other methods of identifying the security requirements of an organization that were discussed earlier in the chapter. Once you have an understanding of the organization's needs, you can determine what customization needs to be done to enhance security.

EXAM
70-293
OBJECTIVE
1.3
1.3.1

Securing Servers According to Server Roles

As you saw earlier in this chapter, servers can be configured in any number of different roles. You can use the Configure Your Server Wizard to configure the server for that role. Although this procedure may install and enable a number of different services, tools, and technologies, additional steps usually are required to ensure the server's security. Some tasks are unique to the server's role, but others should be applied to all servers on your network.

Security Issues Related to All Server Roles

Any server used by members of an organization might be at risk of attacks by hackers and malicious programs, as well as accidents or other disasters. You will want to consider taking a number of countermeasures to ensure that any server is well protected.

Physical Security

As the term suggests, *physical security* addresses the need to protect servers from physical threats. Such threats may affect any number of assets in an organization and can result in widespread damage. These types of threats always involve some level of tangible risk. Taking steps to prevent physical interaction with equipment and implementing methods to ensure that equipment is safe from environmental threats will help promote physical security.

A large part of physical security involves protecting systems from unauthorized physical access. Even if you've implemented strong security that prevents or limits access across a network, it will do little good if a person can sit at the server and make changes or (even worse) pick up the server and walk away with it. If people have physical access to a server, any number of events could occur. They could knock out network cables, bump the server over, spill a drink on electrical components, or unplug it. Physical security controls access to hardware and software, so that people are unable to damage or steal devices and the data

they may contain. If people do not have physical access to systems, the chances of unauthorized data access are reduced.

To prevent physical contact, all servers in an organization should be locked in a secure area. If it can be justified, a dedicated server room should restrict access. If a company's facilities are limited and there is only a single server involved, it should be kept in a locked closet to prevent anyone from touching it. In addition to the server itself, all installation CDs and backup tapes used by the server should be kept under lock and key.

Physical security also involves protecting servers and other assets from environmental disasters. Natural disasters can occur at any time, and they are largely dependent on the geographical location of an office. For example, a branch office in Tornado Alley would need to be able to withstand twisters, and a California branch office might need to withstand earthquakes and mudslides. In both areas, however, Uninterruptible Power Supplies (UPSs) should be installed to provide electricity during power outages, and systems to extinguish fires need to be in place. By considering natural risk sources within an area, you can determine which measures need to be taken to reduce or remove risks.

Physical security not only includes natural disasters, but also those caused by the workplace environment. If a server room isn't properly ventilated with temperature control, the server could overheat or experience issues with electrostatic discharge. In wireless networks, poor environmental conditions could also cause sensitive data to be accessed by other parties who pick up the signals. As data is transmitted, unauthorized parties using special equipment could intercept the packets of data sent over the wireless network. In addition, servers need to be stored in stable areas that adhere to the environmental requirements of the equipment.

Knowing When to Stop Securing Systems

Security is an ongoing process, but there comes a time when you need to decide that enough is enough. No system can be absolutely secure, and every level of security you add restricts access and functionality. For this reason, security is a trade-off, and you need to decide when you've reached an acceptable level.

A major consideration for security is cost versus benefit. At no time should the cost of securing an asset exceed the value of that asset. For example, a server may be configured as a file server and contain sensitive data, which means a higher level of security is needed than for other resources. In providing this security, you don't want to pay more for security than the equipment or data is worth. If the server cost \$7,000, and it would cost the company \$5,000 to replace the data, any security costs over this collective amount would negate any benefits from securing the server. Once it approaches the point where the company is spending an unreasonable amount of money to protect data or equipment, you've exceeded the optimal level of security. Keep in mind that you may need to work with your company's Accounting department to come up with the appropriate numbers.

Service Packs and Hotfixes

At times, software vendors may release applications or operating systems with known vulnerabilities or bugs, or these problems may be discovered after the software has been released. *Vulnerabilities* are weaknesses in programming code that can be exploited. *Bugs* are defects that may cause the software to function incorrectly. To remedy these issues, manufacturers will release service packs, patches, or bug fixes after they have brought their product to market. *Service packs* contain updates that may improve the reliability, security, and software compatibility of a program or operating system. *Patches* and *bug fixes* are used to repair errors in code or security issues. Failing to install these may cause certain features to behave improperly, make improvements or new features unavailable, or leave your system open to attacks from hackers or viruses. In most cases, the service packs, patches, or bug fixes can be acquired from the manufacturer's Web site.

Updates for Windows operating systems are made available on the Windows Update Web site, which can be accessed through an Internet browser by visiting <http://windowsupdate.microsoft.com>. The Windows Update Web site determines what software is recommended to secure your system, and then allows you to download and install it from the site.

Windows Update provides updates for only Windows operating systems, certain other Microsoft software (such as Internet Explorer), and some additional third-party software, such as drivers. To update most third-party programs installed on the computer, you will need to visit the manufacturer's Web site, download the update, and then install it.

Windows 2000, Windows XP, and Windows Server 2003 also provide an automated update and notification tool that allows critical updates to be downloaded and installed without user intervention. When enabled, this tool regularly checks Microsoft's Web site for updates, and if one or more are found, automatically downloads and installs the update. You can also just have it notify you that updates that are available. Because this tool requires connecting to Microsoft over the Internet, it can be used only if the servers or workstations have Internet access.

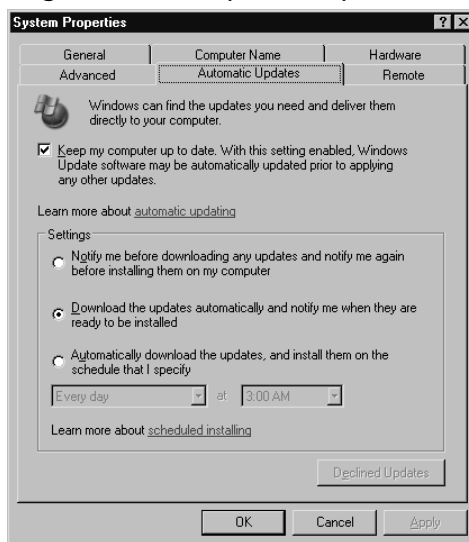
In some situations, administrators may not want Windows Server 2003 to automatically download and install software without their approval, or they may not want computers to connect to the Microsoft Web site in this manner. In these cases, the Automatic Updates service should be disabled or configured so that it is used for notification only. These settings can be accessed by selecting **Start | Control Panel | System** and clicking the **Automatic Updates** tab in the **System Properties** dialog box. As shown in Figure 2.27, the **Automatic Updates** tab provides a number of settings that allow you to configure whether updates are automatically acquired and installed on the computer, when updates occur, and whether intervention is required. These settings include the following:

- **Keep my computer up to date** Enables Automatic Updates on the machine. When this selected, the other settings in this list may be configured.
- **Notify me before downloading any updates and notify me again before installing them on my computer** Informs users that an update is available and asks them if they would like to download it. If the user chooses to have the

update downloaded, Automatic Updates will prompt the user when the download is complete, asking if the update should be installed.

- **Download the updates automatically and notify me when they are ready to be installed** Causes any updates to be downloaded from the Microsoft Web site without any notification. Once the update has completed downloading, the user is asked if the update should be installed.
- **Automatically download the updates, and install them on the schedule that I specify** Causes any updates to be downloaded from the Microsoft Web site without any notification. When this option is chosen, you can specify the time when the update can be installed without user intervention.

Figure 2.27 Choosing Automatic Updates Options



Deciding Whether to Apply an Update

Even though service packs, bug fixes, and patches are designed to fix problems with an operating system or application, you cannot be sure that they will not cause problems themselves. An example of this is Windows NT 4.0 Service Pack 6, which caused major problems after being applied. This service pack was removed from the Microsoft Web site and soon replaced with Service Pack 6a. It's usually a good idea to wait a few days or a week to see if other customers of the manufacturer experience any issues before installing an update.

Continued

To ensure an update works properly and does not cause major problems with computers on your network, it is wise to apply the update to a test machine or computer in a lab environment before applying it to computers on the production network. A test machine has the same configuration and programs installed as other network machines, but it isn't actually used for business purposes. Testing updates on such a computer is especially important when applying updates to servers, because server changes can affect large numbers of users if problems arise.

Although new versions of Windows provide a method of automatically applying updates from Microsoft's Web site, you can configure this service to either notify you before applying an update or disable the service so that Windows will not regularly check for updates. If this service is disabled, you must regularly check for updates manually by visiting Microsoft's site, downloading the updates, and applying them. You should never merely install Windows and leave it without applying critical updates. Failing to apply certain updates may leave your system vulnerable to attack or cause elements of the system to function unexpectedly.

Antivirus Software

Viruses, Trojan horses, and other malicious programs are a threat to any organization, especially if the organization is connected to the Internet. If these programs infect a network, data and systems can be damaged or destroyed. Worse, infection might cause critical information (such as passwords or files) to be transmitted to other sources. To prevent these malicious programs from causing problems, antivirus software should be installed on servers and workstations throughout the network.

When antivirus software is installed, it will scan for viruses and clean them using information stored in signature files. *Signature files* are used to identify viruses and let the software know how to remove them. Because new viruses appear every month, signature files need to be updated regularly by downloading them from the vendor's Web site.

Unnecessary Accounts and Services

Hackers and malicious programs can use insecure elements of a system to acquire greater access and cause more damage. To keep these entities from exploiting elements of your system, you should disable any services that are not needed. If a service has a weakness for which a security patch has not been developed, it could be exploited. By disabling unneeded services, you are cutting off possible avenues of attack. In doing so, you will not affect any functionality used by computers and users, and you can avoid any security issues that may be related to them.

Certain accounts in Windows Server 2003 should also be disabled or deleted. If an account is no longer being used, it should be removed to avoid a person or program using it to obtain unauthorized access. Even if an account will not be used temporarily (for example, during an employee's leave or vacation), the account should be disabled during the

user's absence. If an employee has left permanently or a computer has been removed from the network, these accounts should be deleted.

There are other accounts that you should consider disabling due to their access level. The Administrator account has full access to a system and is a well-known account. Windows Server 2003 and previous versions of Windows all have an account named Administrator that has the ability to do anything on a server. Because hackers already know the username of this account, they only need to obtain password to achieve this level of access. Although the Administrator account cannot be deleted, it can be disabled and renamed. If you create new user accounts and add them to the Administrators group, and disable the Administrator account, attackers will find it more difficult to determine which account to target.

Another account that is disabled by default, and should remain so, is the Guest account. This account is used to provide anonymous access to users who do not have their own account. Like the Administrator account, the Guest account is created when Windows Server 2003 is installed. Because there is the possibility that this account could accidentally be given improper levels of access and could be exploited to gain even greater access, it is a good idea to leave this account disabled. By giving users their own accounts, you can provide the access they need and audit their actions when necessary.

For any user, group, or computer account, it is important to grant only the minimum level of access needed. Employees can accidentally or maliciously modify data or use systems inappropriately. To prevent users from causing such problems, you should never give them more access than they require. You want users to be unable to access anything beyond the scope of their role within the organization. This will assist in keeping other data and systems on the network protected. Determining what level of security a user needs to perform his or her job usually requires some investigation. All users often have their own personal directories for storing files, but they also typically need additional access to databases, programs, and files stored on various servers. To determine how much access a user or group needs, you should begin by discussing the user's duties with management. By understanding the job a user performs, you will be able to determine which resources the user needs to access.

Strong Passwords

Passwords are a key component of the default method of authentication for Windows and other software (such as database management systems). They are used to prevent unauthorized access to computers, networks, and other technologies by forcing anyone who wants access to provide a specific piece information, which should be known only to the authorized user.

Strong passwords are more difficult to crack than simple ones. These types of passwords use a combination of keyboard characters from each of the following categories:

- Lowercase letters (*a–z*)
- Uppercase letters (*A–Z*)

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

The length of a password also affects how easy it is to crack. The more characters used, the more variations of letters, numbers, and special characters the password can contain. You can use security templates and group policies to control how long a password is valid, the length of a password, and other aspects of password management. If you specify a minimum password length of at least seven characters, it will be harder to exploit the account accessed with this password.

In addition, you should avoid using passwords that contain your username, real names, or company name, because these make passwords easier to guess. You should also avoid using passwords that contain actual words that appear in the dictionary, because hacking programs can be used to crack such passwords.

Another requirement that is important to having secure passwords is making sure that each time users change their passwords, they use passwords that are different from previous passwords. All too often, users will use the same password over and over, modifying it slightly. For example, they might have the password “pass1” one month, and then change it to “pass2” the next. In other cases, they might simply change the password each month to the name of the current month (January, February, and so on). Again, ensuring each new password is different from previous passwords will make it more difficult for unauthorized persons to determine current passwords.

To ensure domain controllers are secure, there are a number of password requirements that are enforced by default on Windows 2003 domain controllers:

- The password cannot contain any part of the user’s account name.
- It must be a minimum of six characters in length.
- It must contain characters from three of the four categories: lowercase letters, uppercase letters, numbers, and special characters.

NTFS

Windows Server 2003 supports the FAT, FAT32, and NTFS file systems. Of these, NTFS provides the highest level of security. Using NTFS, you can do the following:

- Set permissions on individual files and folders.
- Control which accounts have access to file system resources.
- Implement file encryption, which prevents unauthorized users from accessing files and folders.
- Implement disk quotas, which allows you to control how much hard disk space users may use.

Using NTFS greatly enhances the security and management of files.

Disk partitions can be formatted with NTFS when a server is initially installed. If a volume is formatted as FAT or FAT32, you can convert it to NTFS. You can convert partitions to NTFS by using the command-line tool `convert.exe`. This tool changes existing partitions into NTFS partitions, without adversely affecting any files on the hard disk.



EXAM WARNING

NTFS is an important part of security on Windows NT, Windows 2000, and Windows Server 2003 systems. Without NTFS, permissions cannot be set on individual files or folders. In Windows Server 2003, other features such as disk quotas and EFS are not available without NTFS.

Regular Backups

It is also important to perform regular data backups. When backups are performed, the data on a computer is copied to other media (such as tape), which can then be stored in another location. If a problem occurs with the source data, you can restore any files that were damaged or lost. For example, if a user accidentally deletes a file or a server's hard drive crashes, a backup can be restored and all files returned to their previous state.

Windows Server 2003 also provides Automated System Recovery and the Recovery Console for restoring systems that have failed.

Recovery Console is a text-mode command interpreter that can be used without starting Windows Server 2003. It allows you to access the hard disk and use commands to troubleshoot and manage problems that prevent the operating system from starting properly. With this tool, you can do the following:

- Enable and disable services.
- Format hard disks.
- Repair the master boot record and boot sector.
- Read and write data on FAT16, FAT32, and NTFS drives.
- Perform other tasks necessary to repairing the system.

You can start Recovery Console from the installation CD for Windows Server 2003, or you can install it on an x86-based computer. When installed on the computer, Recovery Console can be run from a multiple-boot menu that appears when the computer is first started. Either method will start the same program and allow you to enter different commands to repair the system.

Automated System Recovery (ASR) allows you to back up and restore the Registry, boot files, and other system state data, as well as other data used by the operating system. An ASR set consists of files that are needed to restore Windows Server 2003 if the system cannot be started. When you create an ASR set, the following items are backed up:

- System state data
- System services
- Disks that hold operating system components

In addition, ASR creates a floppy disk that contains system settings. Because an ASR set focuses on the files needed to restore the system, data files are not included in the backup.

You should create an ASR set each time a major hardware change or a change to the operating system is made on the computer running Windows Server 2003. For example, if you install a new hard disk or network card, or apply a security patch or service pack, you should create an ASR set. Then, if a problem occurs after upgrading the system, you can use the ASR set to restore the system to its previous state (but only after you've attempted other methods of system recovery).

ASR should not be used as the first step in recovering an operating system. In fact, Microsoft recommends that it be the last possible option for system recovery and be used only after you've attempted other methods. In many cases, you'll be able to get back into the system using Safe Mode, the Last Known Good Configuration or other options.

To create an ASR set, use the Windows Server 2003 Backup utility. On the **Welcome** tab of the Backup utility, click the **Automated System Recovery Wizard** button. This starts the **Automated System Recovery Preparation Wizard**, which takes you through the steps of backing up the system files needed to recover Windows Server 2003 and creating a floppy disk containing the information needed to restore the system.

Securing Domain Controllers

The methods described in the previous sections can improve the security of a server in any role, but they are particularly important for domain controllers. Physical security and strong passwords are needed to prevent unauthorized parties from modifying accounts or other aspects of the domain. In addition, methods to protect the server from malicious programs and tampering should be implemented. Examples include applying updates, installing antivirus software, and formatting all partitions as NTFS.

The effects of an insecure domain controller can be far-reaching. Information in AD is replicated to other domain controllers, so changes on one domain controller can affect all of them. This means that if an unauthorized entity accessed the directory and made changes, every domain controller would be updated with these changes. This includes disabled or deleted accounts, modifications to groups, and changes to other objects in the directory. Because all Windows 2000 Server domain controllers store a writable copy of AD, additional steps must be taken to secure the directory.

It is important that group membership is controlled, so that the likelihood of accidental or malicious changes being made to AD is minimized. This especially applies to the Enterprise Admins, Domain Admins, Account Operators, Server Operators, and Administrators groups.

Because anyone who has physical access to the domain controller can make changes to the domain controller and AD, it is important that these servers have heightened security. Consider using smart cards to control authentication at the server console.

Encryption should also be used to protect data and authenticate users. As mentioned, NTFS partitions allow file encryption, and Kerberos provides strong authentication security. In Windows Server 2003, Kerberos is the default authentication protocol for domain members running Windows 2000 or later.

Securing File and Print Servers

File and print servers also need additional security. In addition to setting permissions on files and folders, regularly performing backups, and using antivirus software, organizations may also need to implement greater levels of protection such as encryption. Similarly, print servers need to be protected from improper use and must be configured to prevent unauthorized users from wasting print resources.

File Servers

Because file servers are used to store data in a central location, it is important that they are kept secure. Although file servers allow the data to be accessed by other users, you need to ensure that only those who are authorized are able to use the files. For this reason, it is especially important that volumes on a file server are formatted as NTFS and appropriate permissions are set on files and folders. As an added measure of security, these disks should also use EFS.

EFS is used to encrypt data on NTFS volumes. When EFS is used, unauthorized users and malicious programs are prevented from accessing the content of files, regardless of their permissions. Although the process involved in the encryption and decryption of data can be quite complex, EFS file encryption is completely transparent to the user.

When a user specifies that a file is to be encrypted using EFS, parts of the file are individually encrypted with *file encryption keys*. These keys are stored in the file header and encrypted using a public key that corresponds to the user who encrypted the file. When the user accesses the file, the file encryption keys are decrypted using the private key that corresponds to the public key that was used to encrypt them. Because this key is held privately by the user who encrypted the file, no one else can access it. The decrypted file is stored in memory, and the original file remains stored in the file system remains encrypted.

When a folder is encrypted with EFS, you have the option of encrypting all files and subfolders inside it. If this option is used, any files that are created in or copied to folder or subfolders are automatically encrypted. If encryption is not specified at the folder level, only the files and subfolders that a user explicitly specifies will be encrypted.

Ensuring That Data Is Encrypted

EFS is an important part of keeping data secure on a file server because it prevents unauthorized parties from viewing and modifying data. When folders are encrypted with EFS, you can have all files and subfolders encrypted as well. If this option isn't used, files on the hard disk will not be encrypted. This is an important issue when users are working with applications that create temporary files.

Temporary files are used to store information while the person is working on a document, spreadsheet, or other file. These files are created by the application and may contain a duplicate working copy of a file. Although applications that use temporary files are supposed to remove them when the files are closed or the program shuts down, this isn't always the case. If an authorized user opens an encrypted file using a program that creates temporary files, an unencrypted temporary file may be created. Potentially, a hacker who could not access the data in the encrypted file might be able to open the temporary file and view the data inside it.

To ensure that temporary work files are not accessible, you should encrypt the temporary folder you specify in the application. In this way, when the application creates a temporary file, it will automatically be encrypted, eliminating a potential target for hackers.

Although EFS is an important part of securing a file server, this does not mean that every file on the network is a candidate for being encrypted with EFS. As mentioned, only files on NTFS volumes can be encrypted with EFS. If a volume is formatted as NTFS, files that have the System attribute or are located in *%systemroot%* (for example, C:\Windows) cannot be encrypted. Also, if the file or folder you want to encrypt is compressed, you cannot use encryption. The opposite is also true: if a file or folder is encrypted with EFS, it cannot be compressed.



NOTE

In Windows 2000, there was no visual indication of which files and folders were encrypted. This made management of EFS difficult because you needed to examine file and folder properties when attempting to ascertain which ones were encrypted. In Windows Server 2003, Microsoft developers included an option that colors encrypted files and folders green, so that they can be easily spotted in Windows Explorer and other applications.

Another important limitation of EFS is that it encrypts data only on NTFS volumes. When a file is accessed remotely on a file server, Windows Server 2003 decrypts it and sends it across the network in unencrypted form. For data to be encrypted during transmission, other technologies like IPSec must be used.

IPSec ensures that data is sent securely over the network by encrypting packets and authenticating the identity of the sender and receiver. When using IPSec, a policy is applied to both the sender's and receiver's computer, so the systems agree on how data will be encrypted. Other computers that intercept traffic between the machines will be unable to decipher the information contained in the packets.

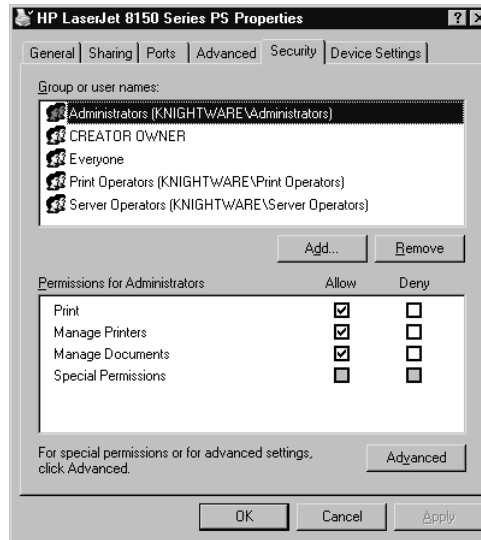
Print Servers

Files that are being printed may also require protection. IPSec can be implemented to protect the transmission of data being sent to printers. After all, if a document can be captured while being sent to a printer, a hacker can view its information just as if it were being accessed directly from a server.

Physical security issues can be very important for printers. Anyone with access to a printer can remove printed documents from it. This is especially critical for printers that are routinely used to print sensitive documents or financial instruments like checks. A sensitive document may reside on a highly secure file server, but once it is printed, anyone standing by the printer could simply pick it up and walk away. To prevent this from happening, such printers should be located in secure areas that are not accessible to the public and other unauthorized users.

Just as files can have permissions assigned to them, so can printers. Printer permissions are used to control who can print and manage network printing. As shown in Figure 2.28, they are set on the **Security** tab of a printer's properties. Using printer permissions, you can allow or deny the following permissions for users:

- **Print** Allows users to print documents.
- **Manage Printers** Allows users to perform administrative tasks on a printer, including starting, pausing, and stopping the printer; changing spooler settings; sharing the printer; modifying permissions; and changing property settings.
- **Manage Documents** Allows users to perform administrative tasks relating to documents being printed. It allows users to start, pause, resume, reorder, and cancel documents.

Figure 2.28 Setting Permissions for a Printer


Although different permissions exist for printing, only the Print permission gives the ability to print a document. For example, when only the Manage Documents permission is given, the user has the ability to manage other people's documents but cannot send documents to the printer for printing. Because those who manage printers may need to print test pages to determine if the printer is working properly, the Manage Printers permission can be set only if the Print permission is given.

Because the Print permission is assigned to the Everyone group, all users have access to print to a printer once it is shared on the network. For most printers, it's usually a good idea to remove this permission and add the specific groups within your organization that should have access to the printer.

Securing DHCP, DNS, and WINS Servers

DHCP, DNS, and WINS servers often provide the ability to connect to the network and find other computers. DHCP is used to provide IP address and configuration information to clients. DNS and WINS servers are used to resolve names to IP addresses (and vice versa). If you do not secure these servers, malicious persons and programs may be able to prohibit users from connecting to the network, redirect traffic to other locations, and impact the ability to use network resources.

DHCP servers do not require authentication when providing a lease. Any client that contacts the DHCP server can obtain a lease and connect to the network. In addition to receiving an IP address as part of the lease, clients may also be automatically configured with WINS or DNS server information. To avoid this, it is important that you restrict physical and wireless access to your network. This helps to prevent unauthorized persons from successfully connecting to your network and obtaining a valid DHCP lease. In addition,

auditing should be enabled on the DHCP server so that you can review requests for leased addresses. By reviewing the logs, you may be able to identify possible problems.

Just as DHCP is an unauthenticated protocol, so is the NetBIOS naming protocol used by WINS. WINS was designed to work with NetBIOS over TCP/IP (NetBT), which does not require any authentication. Because a user does not need to provide credentials to use WINS, it should be regarded as available to unauthorized persons or programs. These entities could request a massive number of names to be registered or resolved by the WINS server, so that the server becomes bogged down and unable to process other requests. This type of attack is called a denial of service (DoS) and is designed to overload systems and prevent access for legitimate users.

Rogue servers can also be a problem on the network. When a client requests a DHCP lease, it does so by broadcast. If an unauthorized person puts a DHCP server on the network, the incorrect IP address and configuration information could be provided to clients. This isn't the case if the rogue DHCP server is running Windows 2000 or Windows Server 2003, because these must be authorized in AD. If the server determines that it is not authorized, the DHCP service will not start. However, pre-Windows 2000 and non-Windows DHCP servers require no authorization and can be effectively used as rogue DHCP servers in a Windows Server 2003 environment. Handing out bogus DHCP leases that do not expire can be a very effective DoS technique. Because of this, it is important to monitor network traffic for DHCP server traffic that does not come from your network's authorized DHCP servers.

Restricting access to DHCP tools and limiting membership in groups that can modify DHCP settings are other important steps in securing a DHCP server. To administer DHCP servers remotely using the DHCP console or Netsh utility, you need to be a member of the Administrators group or the DHCP Administrators group. By restricting membership in these groups, you limit the number of people who can authorize a DHCP server to service client requests.



TEST DAY TIP

Many people get DHCP, DNS, and WINS confused with one another. Remember that DHCP is used to assign IP addresses to clients, while DNS and WINS are used for name resolution. To avoid confusing DNS and WINS, remember that DNS is the Domain Name System. Remembering its name will help you associate that DNS is used to resolve DNS names to IP addresses and vice versa. Through the process of elimination, this will make it easier to remember that WINS is used to resolve NetBIOS names to IP addresses and vice versa.

Securing Web Servers

Because IIS provides a variety of services that allow users to access information from the Web server service, it provides potential avenues of attack for unauthorized users, malicious

programs, and other sources. For this reason, it is not installed by default. If you do not need a Web server on your network, IIS should remain uninstalled. If it has been installed on servers that do not need it, make sure to uninstall it.

Once IIS is installed on Windows Server 2003, it is locked down to prevent any unneeded services from being exploited. By default, IIS will provide only static content to users. If dynamic content is used on the server, you will need to enable the necessary features. For example, if your site is going to use ASP, ASP.NET, Common Gateway Interface (CGI), Internet Server Application Programming Interface (ISAPI) or Web Distributed Authoring and Versioning (WebDAV), each of these will need to be enabled before they can be used. As with Windows Server 2003 itself, any components that are not needed should be disabled.

Another default setting of IIS is that it will not compile, execute, or serve files with dynamic extensions. For example, if you have Web pages written as ASPs with the extension .asp, IIS won't provide users with this content. These are not allowed by default because of Microsoft's new security initiatives. Dynamic content can contain malicious code or have weaknesses that can be exploited. If files that provide dynamic content need to be used on the Web server, you must add the file extensions to the Web service extensions list. Any file types that are not needed should not be added.

An important part of protecting Web servers is using firewalls. Firewalls prevent direct access between a network and clients by having traffic pass through the firewall, which determines if the traffic should be blocked or allowed. In other words, it acts as a buffer between the Web server and clients using it or between the internal network and other networks like the Internet. Rules can be set up on the firewall controlling what kinds of traffic may pass and who can perform certain actions. For example, the firewall might prevent AVI files from being transmitted from the Internet for general users but not administrators. Or, it might prohibit executable downloads to prevent virus-infected files or Trojan horses from being installed on clients.

Securing Database Servers

When securing databases, you should take advantage of security features offered by the database software. Microsoft SQL Server, for example, provides two methods of authenticating clients to access data: Windows Authentication Mode and Mixed Mode. When Windows Authentication Mode is used, the SQL Server administrator has the ability to grant logon access to Windows user accounts and groups. If Mixed Mode is used, users can be authenticated through either Windows authentication or separate accounts created within SQL Server.

Regardless of the authentication mode used, like many database applications, SQL Server allows you to control access to data at a granular level. Permissions can be set to determine the operations that a user can perform on the data contained in the database. In many database applications, you can set permissions at the server, database, or table level. While one account might have the ability to create tables and delete data in all databases, another may only be able to view data in a single database. These permissions are different

from those that can be set through AD and NTFS, and they apply only within the database program.

Database servers may also need to be secured through other roles that are used to access the database. For example, IIS is set up through the application role, and Web pages on the server can be used to access data stored in a database. Similarly, applications that are developed and made accessible from a terminal server may be used to view and manipulate database information.

To control access to the database server, you can use settings configured through a *data source name* (DSN). A DSN is commonly used by compiled and Web-based programs to gain access to data that is stored in data management systems and data files. A DSN contains information on the database name, the server it resides on, and the directory in which it's stored (if a data file is used). It also holds the username, password, and driver to use when making the connection. Programs use information in the DSN to connect to the data source, make queries, and manipulate data. To create or modify a DSN, use the Data Sources (ODBC) applet (select **Start | Administrative Tools | Data Sources (ODBC)**).

Because a DSN provides the username and password to use when connecting to the data source, a number of security-related issues arise from its use. Any passwords that are used should follow the recommendations for strong passwords that were discussed earlier in this chapter. In cases where a DSN is being used to connect to a SQL Server database, you also have the option of using Windows authentication or SQL Server authentication. If SQL Server authentication is used, you can enter the username and password of an account created in SQL Server. However, you should avoid entering the name of any accounts with access higher than the user will need. For example, entering the system administrator account (**sa**) would provide a DSN with full access to SQL Server and could maliciously or accidentally cause problems. To avoid possible damage to data or access violations, you should provide the username and password of a SQL Server account that has restricted access.

Securing Mail Servers

When Windows Server 2003 is configured with the mail server role, it should be set up to require secure authentication from e-mail clients. As mentioned earlier, clients retrieve their e-mail from mail servers using the POP3 protocol. Client software and the mail server's POP3 service can be configured to accept only passwords that are encrypted in order to prevent them from being intercepted by unauthorized parties.

In Windows Server 2003, the Microsoft POP3 Service uses Secure Password Authentication (SPA) to ensure that authentication between the mail server and clients is encrypted. SPA is integrated with AD, which is used to authenticate users as they log on to retrieve their e-mail. In cases where domain controllers are not used, SPA can authenticate to local accounts on the mail server. When the POP3 service is configured to accept only authentication using SPA, clients must also be configured to use encrypted authentication.

If they are not, clients will attempt to authenticate using cleartext (which is plaintext, or unencrypted data) and will be rejected by the mail server.

To prevent mail servers from filling up with undeleted or unchecked e-mail, disk quotas should also be implemented. E-mail can include attachments, which are files that are sent with messages. Over time, mail left on the server can fill up hard disk space and affect server performance. By using disk quotas, users can be limited to a specific amount of hard disk space. Disk quotas can be used only on NTFS partitions. When NTFS is used, permissions can also be set on the directories that store e-mail, preventing unauthorized parties from accessing it on the server.

Securing CAs

In addition to the basic server hardening techniques mentioned earlier in this chapter, a CA needs additional levels of security applied to it. Recall that a root CA resides at the top of the hierarchy, with subordinate CAs existing below it. Because the root CA is the most trusted one in a hierarchy, any CAs below it automatically trust it. These subordinate CAs use the root CA's public key and bind it to its own identity. In doing so, the subordinate can also issue certificates to users and computers.

Because of the trust between root and subordinate CAs, if the root CA is compromised, subordinate CAs continue trusting it. This compromises all certificates issued by the CAs in the hierarchy. As a security measure, you should disable the root CA's ability to issue certificates online and allow only child CAs to perform this function. An offline root CA is more difficult to compromise, since physical access to it is required.

Additional benefits can be derived from the use of enterprise CAs. When a user requests a certificate from an enterprise CA, that CA is able to validate the information provided by the user through AD. This can provide an extra measure of security. Stand-alone CAs require manual inspection and approval of requests by a CA administrator. Manual processes are typically much more error-prone than automated ones.

When certificates are found to be invalid, they should immediately be revoked. After a certificate is revoked, the CRL should be immediately updated and published. The CRL is used to inform the world of certificates that are no longer valid. If the certificate is invalid, the software used to check it often allows the user to decide whether or not to trust the certificate holder.



NOTE

Although you can publish a CRL immediately, that does not necessarily mean that all hosts will begin to use the new list. CRLs are cached on local hosts and will not be refreshed until the update period is reached. As a result, the old list that allows invalid certificates will continue to be used until a host checks back in for an update. As an administrator, you can determine how frequently the CRL is updated at the host level. You'll need to balance your security needs against the network traffic requirements of a CRL update and choose an appropriate interval for your organization.

Securing Application and Terminal Servers

Application and terminal servers are also configurable server roles that need additional steps to ensure that they are secure. Users are able to access applications across the network and execute them on servers using each of these roles. Because of the importance of many network-accessed applications, and the damage that can be done if they are exploited, it is essential that these roles are protected.

Application Servers

Application servers provide access to a wide variety of data on the network, and they need to be hardened using the methods discussed earlier. The tasks users perform on a network often rely on their ability to use specific software and to be assured that all data is secure. To achieve these goals, hard disks storing these applications and the files they generate should be formatted with NTFS.

There is also a need for in-house applications, which are developed by programmers working for the company, to use the latest development tools. Older application development tools may have vulnerabilities that can be exploited. For example, a program developed using an older version of Visual Basic could be decompiled, allowing a hacker to view the code used to create the program. Code generated for in-house applications often contains sensitive information such as server names and authentication information. In addition, older development software is often not able to take advantage of the latest advances in security.

Application servers need to use the general security methods discussed for other servers. They may need to connect to other servers to acquire information or provide services. For example, an application hosted on an application server may use a database server to acquire and process data before returning it to end users. Because the two work in conjunction, the database server must also be secured. Even though the application server might be exceptionally secure, if there are security issues on the database server, the data might be compromised, which can potentially affect the ability of the application to do its job.

Servers configured in the application server role also have IIS 6.0 installed by default. IIS lets the application server provide Web-based applications to users of the network. Because the application server may have a Web server installed on it, steps need to be taken to ensure the Web server is also secure, as discussed earlier in this chapter.

Terminal Servers

Because terminal servers provide access to applications and data, they also need to be configured to ensure that users and hosts do not achieve unauthorized access. By setting permissions on connections, you can control who can access a server and perform specific tasks. This is in addition to the permissions that can be set on files accessed by users in a terminal server session. By limiting access in these ways, you can control who is able to use files and applications and what actions they are able to perform.

Terminal servers can also be configured to use specific levels of encryption. When a communications link is established between a client and the terminal server, the data transmitted between them can be encrypted to prevent others from being able to view and use it. The following encryption levels can be set:

- **High** This is the default level. It uses 128-bit encryption, which may not be supported by all clients. If clients do not support this level of encryption, they will be unable to connect to the terminal server.
- **Low** This level provides only one-way encryption. Clients send data to the server using 56-bit encryption, but any data sent from the server to the client is unencrypted.
- **FIPS compliant** This level encrypts data using Federal Information Processing Standard (FIPS) encryption algorithms and is mandated for use by portions of the U.S. government.
- **Client compliant** This level encrypts data using the strongest possible key strength supported by the client. Because the level of encryption depends on the client, it may be a good idea to use it if legacy clients or a mix of clients are used on the network. However, if you have strong security requirements, this level does not allow you to specify the encryption level clients will use, so it should not be used.

EXAM
70-293
OBJECTIVE
1.3.2

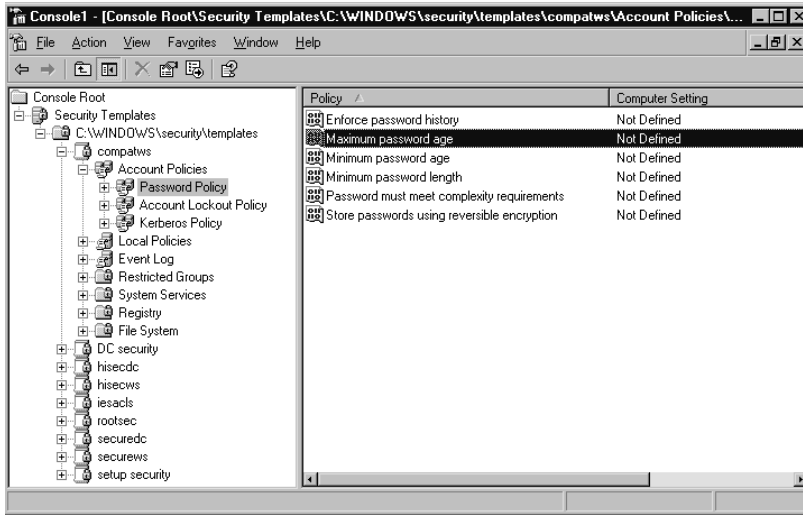
Creating Custom Security Templates

Earlier in this chapter, we discussed how you can use predefined security templates to modify security settings. Although these templates contain settings that can be used for a number of purposes, they may not have the settings you specifically need for your organization. In such cases, you may want to create custom security templates.

You can create custom security templates in a number of ways. As described earlier, modifying the results of an analysis using Security Configuration and Analysis, and then exporting the changes to a new template file, is one way to create a custom security template. In addition, you can create custom security templates using the Security Templates snap-in. The Security Templates snap-in allows you to modify existing templates and create new ones from scratch.


As shown in Figure 2.29, Security Templates consists of two panes. The left pane contains the **Security Templates** node. When expanded, this node reveals the default template location (`%systemroot%\Security\Templates`) and the child nodes that contain policy templates. Each policy node contains groups of settings that, when selected, appear in the right pane of the utility.

Figure 2.29 The Security Templates console



To define the settings for a particular policy in the group, right-click the policy in the right pane and selecting **Properties**. Each dialog box contains different options that are relevant only for that setting. For example, the Properties dialog box for the **Maximum password age** policy is shown in Figure 2.30. The **Define this policy setting in the template** check box enables the configurable settings in this dialog box. In the case of the **Maximum password age** policy, the settings in the Properties dialog box allow you to control the number of days until the password expires. Clicking **OK** applies any changes you have made to the policy.


Figure 2.30 Setting Maximum Password Age Properties



In addition to modifying existing templates, you can create new templates with Security Templates. Right-click the folder in which you want to store the new template

and click **New Template**. In the dialog box that appears, shown in Figure 2.31, enter a name for the template in the **Template name** text box. Optionally, you can enter a description in the **Description** text box. After you click **OK**, the new template will appear in the list.

Figure 2.31 Adding a New Security Template



When a new template is created, all of the settings in it are undefined. In other words, there are no restrictions set within the template, because all settings are bypassed. To configure security settings for the template, you must go through each node and make the necessary changes to the policy settings, as discussed earlier. Exercise 2.04 guides you through the process of creating a new template and modifying a setting in it.

EXERCISE 2.04

CREATING A NEW TEMPLATE USING SECURITY TEMPLATES

1. Select **Start | Run**, type **MMC**, and click **OK**.
2. When MMC opens, click **File | Add/Remove Snap-in**.
3. In the **Add/Remove Snap-in** dialog box, click the **Standalone** tab to select it (if necessary).
4. Click the **Add** button. When the **Add Standalone Snap-in** dialog box appears, select **Security Templates** from the list and click **Add**.
5. Click **Close** to return to the previous window. A **Security Templates** entry should appear in the **Add/Remove snap-in** dialog box. Click **OK** to close the dialog box.
6. The console tree in the MMC should now contain a **Security Templates** node in the left pane. Expand this node to display the **%systemroot%\Security\Templates** node. (Note that **%systemroot%** will be replaced with your actual Windows directory location, such as **C:\Windows**.)

Expand this node to display all of the security templates that are stored in this directory.

7. Right-click the `%systemroot%\Security\Templates` node and select **New Template** from the context menu.
 8. In the dialog box that appears, type **TestTemplate** in the **Template name** text box, and then click **OK** to continue. The new template should now appear in the left pane of **Security Templates**.
 9. Expand the **TestTemplate** node to view the child nodes within it.
 10. Expand **Account Policies**.
 11. Click the **Password Policy** node to display the policy settings it contains.
 12. In the right pane, double-click the **Maximum password age** policy to display the **Properties** dialog box.
 13. Select the check box next to **Define this policy setting in the template**.
 14. Change the value in the **Password will expire in** box to **90**.
 15. Click **OK**. The **Suggested Value Changes** dialog box will appear. Because the Minimum Password Age value hasn't been set, this policy will automatically be adjusted to 30 days.
 16. Click **OK** to accept the change, and then exit the **Properties** dialog box.
 17. Select **File | Save As**.
 18. When the **Save As** dialog box appears, enter a name for this template in the **File name** text box, and then click **Save** to save this new template.
-

Deploying Security Configurations

As mentioned earlier in this chapter, security configurations can be deployed either manually on the local computer or to multiple systems using AD. You learned how to use Security Configuration and Analysis and Secedit tools to apply a security template to a single computer. Also, when you use GPOs to deploy security templates, you must use Active Directory Users and Computers (for GPOs at the domain and OU levels) or Active Directory Sites and Services (for GPOs at the site level).

Computers have local security policies, which reside on the machine and affect only that particular computer. A user who logs on to the computer is subject to the policy settings that have been configured. The security policy can control a wide range of settings,

including whether the user's actions are audited, the resources the user is allowed to use, and whether the user can even access the computer.

GPOs can be applied to any Windows 2000 or later computer that has joined a domain. They can also be applied to user accounts. Security settings configured in GPOs override those made at the local computer level. Because policies can be set at the site, domain, and OU levels in AD, a computer or user may be subject to a wide combination of security settings. Policy settings are cumulative and applied in the following order:

1. Site-level GPOs that affect the computer account
2. Domain-level GPOs that affect the computer account
3. OU- and sub-OU level GPOs that affect the computer account
4. Site-level GPOs that affect the user account
5. Domain-level GPOs that affect the user account
6. OU- and sub-OU level GPOs that affect the user account

By default, all settings applied will be in effect for the user and computer. However, it is also possible that some settings may conflict between GPOs. For example, a site-level policy that applies to the computer may specify a different setting (such as a user right) than an OU setting (the same right, but configured differently) that is applied later. By default, the last setting applied is the effective setting. This means that the OU-level setting would be in effect. Administrators can modify this behavior.

When security settings are applied using GPOs, they do not immediately affect the computer, as local computer policies do. Local computer policies are stored on the computer and take effect immediately. GPO settings are stored in AD and need to be downloaded to the machine. The Group Policy settings are refreshed on computers at regular intervals. Workstations and member servers have group policy settings refreshed every 90 minutes, with a random 30-minute offset (so that all clients do not refresh at the same time and overload the domain controllers). Domain controllers are refreshed every 5 minutes because of their additional security needs. In addition, security settings in GPOs are refreshed every 16 hours, regardless of whether changes have been made to the policy.

If you do not want to wait for an automatic refresh of group policy settings to take place, you can use the **gpupdate** command to force a refresh. This command replaces the **secedit /refresh** command that was used in Windows 2000. This command has the following syntax:

```
gpupdate [/target:{computer | user}] [/force] [/wait:Value]
        [/logoff] [/boot]
```

The **gpupdate** parameters are defined in Table 2.4.

Table 2.4 Parameters for the gpupdate Command

Parameter	Description
/target:{ <i>computer</i> <i>user</i> }	Used to specify that just the computer or the user settings should be processed. By default, both are processed.
/force	Used to reapply all settings. By default, only changed settings are applied.
/wait: <i>Value</i>	Used to specify when the command prompt should become available during the processing of group policy settings. When the timeout is reached, processing continues in the background, but the command prompt is made available. Status messages will not be displayed in the console if control of it has been returned by the application. By default, it will wait 600 seconds for policy processing to finish. If 0 is used, the program won't wait. If -1 is used, it will wait indefinitely.
/logoff	Specifies that the computer should log off the current user if client-side extensions are used in the Group Policy settings that are refreshed only at logon. An example of such an extension would be those dealing with user-targeted software installation and folder redirection. Some policies, like these, cannot be applied with a background refresh.
/boot	Specifies that the computer should restart if client-side extensions are used in Group Policy settings that are only applied at bootup. An example is a software installation policy that is applied to the computer. Some policies, like this one, cannot be applied with a background refresh.
/sync	Specifies that the next foreground policy application is to be done synchronously. This type of policy is applied when the computer boots up and when the user logs on.

Summary of Exam Objectives

When Windows Server 2003 is installed on a machine, additional configuration is needed to ensure it provides the necessary functionality and is secure. Windows Server 2003 can be configured to perform up to 11 different roles. Each role provides additional tools, services, and features that can be used to enhance your network and (in a number of cases) improve security.

Part of creating a secure environment involves choosing the right operating system. In this chapter, we compared the minimum requirements for various versions of Windows. We also saw that Windows Server 2003 offers a number of new features that were not available in previous versions, while still providing backward-compatibility to older systems. By using features like Kerberos authentication, functional levels, smart card support, and the ability to create domain controllers from backups, you can create secure environments and perform tasks more easily.

You can configure server roles using security templates, applying specific settings to a machine to make it more secure. These templates can be applied to member servers, workstations, and domain controllers, by using Local Security Policy or GPOs in AD. Because not all templates will contain the settings you want for your domain, you can modify them using the Security Configuration and Analysis tool or Security Template tool. Then you can use the Security Configuration and Analysis tool or the Secedit command-line utility to configure the server with the settings stored in the template.

In addition to customizing templates, you can perform other steps to provide security to your systems. These include implementing physical security, using antivirus programs, using NTFS on hard disks, using strong passwords, and other initiatives related to the role a server plays on the network. By using these methods, you can help protect the systems in a domain and forest from various threats.

Exam Objectives Fast Track

Understanding Server Roles

- ☑ There are 11 different server roles available for Windows Server 2003. These include domain controller, file server, print server, mail server, application server, terminal server, remote access/VPN server, streaming media server, DHCP server, DNS server, and WINS server.
- ☑ Manage Your Server is a tool in Windows Server 2003 that allows you to view information about installed server roles, view additional information, and invoke other tools used for administering a server.
- ☑ The Configure Your Server Wizard steps you through the process of installing or removing server roles on Windows Server 2003 servers. The domain controller role is the only one that can be added but not removed with the Wizard.

Planning a Server Security Strategy

- ☑ Windows Server 2003 is available in Standard, Enterprise, Datacenter, and Web Editions. This version provides a number of features that were not available in previous versions. Of the different editions, the Web Edition is the only one that cannot be used as a domain controller.
- ☑ Not all editions of Windows Server 2003 can be installed on every computer. Just as there are different minimum requirements for the various versions of Windows, there are also different minimum requirements for the different editions of Windows Server 2003.
- ☑ Windows Server 2003 provides four different levels of domain functionality: Windows 2000 mixed, Windows 2000 native, Windows Server 2003 interim, and Windows Server 2003. It also supports three levels of forest functionality: Windows 2000, Windows Server 2003 interim, and Windows Server 2003.

Planning Baseline Security

- ☑ Security templates contain settings that can be applied using Local Security Policy or Group Policy.
- ☑ Security Configuration and Analysis is an MMC snap-in that allows you to analyze security settings by comparing them to entries in a database. It also allows you to apply template settings.
- ☑ The Secedit command-line tool is similar to Security Configuration and Analysis. It also allows you to analyze and apply security settings using templates.

Customizing Server Security

- ☑ Automatic Updates can be configured to automatically download and install critical updates for the Windows operating system.
- ☑ NTFS is important to the security and availability of files on a hard disk. Using NTFS, you can set permissions on files and folders, implement EFS to encrypt files, use DFS to allow users to access files from a central location, and have disk quotas control how much disk space users can use.
- ☑ Security templates can be applied to computers individually using Local Security Policy or to many computers at once using GPOs. To import security templates into a GPO, you can use the Group Policy Object Editor. To link a configured GPO to the domain or OU level, use Active Directory Users and Computers. To link a configured GPO to the site level, use Active Directory Sites and Services.

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 set up Windows Server 2003 in the role of a Web server, but when I use the Configure Your Server Wizard, there isn't a Web server role offered in the list. What should I do?
- A:** The Configure Your Server Wizard doesn't offer a Web server role, but it can be set up through the application server role. You can also set up Internet Information Services (IIS) 6.0 through Add or Remove Programs in Control Panel. Adding IIS installs all of the basic features needed to implement a Web server.
- Q:** My network consists of servers running Windows 2000 Advanced Server. It was my understanding that multiple objects could be modified in Active Directory, but I find that I'm unable to do so. Why is this?
- A:** Different versions of Windows offer different features. Windows Server 2003 allows you to select multiple objects and change some of their common attributes at the same time. This ability wasn't available in previous versions.
- Q:** Why do Windows Server 2003 domain controllers use NetBIOS names in addition to DNS names?
- A:** NetBIOS names are used to provide backward-compatibility. They are used by pre-Windows 2000 computers and allow users of those operating systems to log on to Windows Server 2003 domains.
- Q:** I want to apply security settings to computers after regular business hours so I don't disrupt work being performed during the day. What tool should I use?
- A:** Secedit is a command-line tool that allows you to configure machines using security templates. Because it is a command-line tool, it can be invoked through batch files and scripts, which you can schedule to run after regular business hours.
- Q:** I want to create a custom security template. Which programs could I use to create this file?

- A:** Security Configuration and Analysis, Security Templates, Secedit, and Group Policy Object Editor are tools that come with Windows Server 2003 that can be used to create template files. You can use Security Configuration and Analysis to view an existing template and customize it to your needs. The Security Templates snap-in can be used to create new templates and modify existing ones. The Group Policy Object Editor allows you to review a GPO's current settings and export them to a template file. Finally, Secedit can export settings to a template file from the command line.
- Q:** I have created a custom security template and applied it to the Local Security Policy of workstations in a Windows Server 2003 domain. When users log on to the domain, the settings I changed in the Local Security Policy don't take effect. Why is this?
- A:** Settings in a GPO take precedence over those in the Local Security Policy. Any setting obtained from a GPO will override those on 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 Server Roles

1. Your network consists of two machines running Windows Server 2003 Standard Edition, one machine running Windows Server 2003 Datacenter Edition, one machine running Windows Server 2003 Web Edition, and two machines running Windows Server 2003 Enterprise Edition. You want two of these machines to be domain controllers on the network. Which machines will you promote to domain controllers and how will you configure them in this role?
 - A. Configure the two machines running Windows Server 2003 Enterprise Edition to be domain controllers using the `secedit /configure` tool.
 - B. Promote the Windows Server 2003 Datacenter Edition and Windows Server 2003 Web Edition using the `DCPROMO` tool.
 - C. Configure a machine running Windows Server 2003 Standard Edition and a machine running Windows Server 2003 Enterprise Edition to be domain controllers using the Configure Your Server Wizard.
 - D. Configure machines running Windows Server 2003 Standard Edition and Windows Server 2003 Web Edition using the Manage Your Server tool.

2. Your network is upgrading from Windows NT 4 to Windows Server 2003 and will consist of two domains in a single forest. One domain is a child of the other domain and dedicated to the Sales departments in the organization. During the upgrade, all workstations will be upgraded to Windows XP and Windows 2000 Professional. When the last BDC is removed from the network, what role will the PDC emulator play on the network?
 - A. The PDC emulator will be used to modify object classes and attributes.
 - B. The PDC emulator will receive preferred replication of password changes performed by other domain controllers in the domain.
 - C. The PDC emulator in the child domain will be used to synchronize the time on all domain controllers in the forest.
 - D. The PDC emulator will be used to add new domains and remove unneeded ones from the forest.

3. The only protocol used by your network is TCP/IP, despite the fact that workstations in the organization do not have access to the Internet. A user has been accessing files on server on your network and now wants to connect to a Web server that is used as part of the company's intranet. The user enters the URL of the Web site into Internet Explorer. Which of the following servers will be used to provide information needed to connect to the Web server?
 - A. DHCP server
 - B. DNS server
 - C. WINS server
 - D. File server

4. You want to set up a discussion group that can be accessed over the corporate intranet, so that users can view and post messages in a forum that can be viewed by other employees. Which of the following services would you use to implement this functionality?
 - A. HTTP
 - B. FTP
 - C. NNTP
 - D. SMTP

Planning a Server Security Strategy

5. You are planning to use a server on your network as a Windows Server 2003 domain controller. The server has 128MB of RAM, 2GB of hard disk space, and four processors. Which of the following editions of Windows Server 2003 can you install on this server? (Select all that apply.)
 - A. Windows Server 2003 Standard Edition
 - B. Windows Server 2003 Enterprise Edition
 - C. Windows Server 2003 Datacenter Edition
 - D. Windows Server 2003 Web Edition

6. You are concerned about insecure methods of authentication being used on a network. You are currently upgrading your network to Windows Server 2003, but some servers are still running Windows NT 4 and Windows 2000 Server. Even after the upgrade, some Windows 2000 Server computers will exist in the domain. You want to implement Kerberos authentication within the domain. Which of the following operating systems will be able to use it? (Select all that apply.)
 - A. Windows NT 4
 - B. Windows 2000 Server
 - C. Windows Server 2003
 - D. None of the above

7. Your network consists of two Windows Server 2003 domain controllers, a Windows 2000 server that is used as a Web server, and a Windows NT 4 server that runs an older version of SQL Server. Your company does not have the budget to immediately replace these servers, but you want to raise the domain functional level of your domain to the highest possible level. What functional level will you raise this domain to?
 - A. Windows 2000 mixed
 - B. Windows 2000 native
 - C. Windows Server 2003 interim
 - D. Windows Server 2003

Planning Baseline Security

8. You have just promoted a Windows Server 2003 computer to be a domain controller. After the promotion, you accidentally apply the wrong security template to it. It now

has security settings than that are too high. You can automatically change the security settings back to their previous configuration using which of the following security templates?

- A. Setup security
 - B. Rootsec
 - C. IesacIs
 - D. DC security
9. You want to apply an existing security template to the local computer policy of a Windows Server 2003 computer. Which of the following tools would allow you to do this from the command line?
- A. Security Configuration and Analysis
 - B. `secedit /configure`
 - C. `secedit /import`
 - D. `gpupdate`
10. You have performed an analysis of a Windows Server 2003 domain controller using Security Configuration and Analysis. Once the analysis is complete, a red X appears beside the Enforce Password History policy. What does this mean?
- A. The policy does not match a corresponding setting for the associated entry in the database.
 - B. The entry in the database and the policy's setting match.
 - C. An entry exists in the database that does not correspond to any setting on the computer.
 - D. A setting exists on the computer that does not correspond to any entry in the database.
11. You have created a security template and now want to apply its settings to a GPO that can be linked to containers in Active Directory. Which containers can you link a GPO to in Active Directory? (Select all that apply.)
- A. Domains
 - B. Trusts
 - C. Sites
 - D. Local computer policy

Customizing Server Security

12. You have installed a new file server on the network and formatted it to use NTFS. After formatting is complete, you use EFS to encrypt a folder containing files belonging to users. If a user accesses a file belonging to him in this folder, and then copies it across the network for another user to access, which of the following will occur?
- A. The file on the hard disk and the data sent over the network will remain encrypted.
 - B. The file on the hard disk and the data sent over the network will be decrypted and remain that way.
 - C. The file on the hard disk will be decrypted, so EFS can send it encrypted over the network.
 - D. The file on the hard disk will remain encrypted, but data sent over the network will be unencrypted.
13. You have created a custom security template that you now want to import into a GPO that is linked to the domain level. Which of the following tools will you use to invoke the Group Policy Object Editor to view and modify the GPO at this level?
- A. Active Directory Users and Computers
 - B. Active Directory Sites and Services
 - C. gpupdate
 - D. Securedc
14. Your network consists of servers running Windows 2003 Server and workstations running Windows 2000 Professional. You have applied several custom security templates to GPOs linked to the OU, domain, and site levels in Active Directory. In addition to this, there are security settings that have also been applied at the local computer level of all machines that are on the network. Because several policies now affect the computer accounts within the domain, site, and OU, which of the following will occur when the user logs on to the domain?
- A. The policy setting at the local computer level will be overwritten by the OU-level GPO, which will be overwritten by the domain-level GPO, which will finally be overwritten by the site-level GPO. For this reason, major security settings must be made at the site-level GPO; all others will be overwritten.
 - B. Security settings in the GPOs will not be applied to machines running Windows 2000 that have joined the domain.
 - C. The security settings at the local computer level will override those of the GPOs.

- D. The policy settings will be cumulative and applied in the order of policies at the site level, domain level, and finally OU level.
15. You apply custom security templates to the local computer policy on a member server and to a GPO linked to an OU in Active Directory. All servers on the network are running Windows Server 2003. After performing these actions, you find that the local computer policy has taken effect, but the group policy has not taken effect on member servers within the domain. Which of the following is the reason for this, and how can you fix it?
- A. Group policy settings take effect immediately. The problem must be that the security policy was not applied properly.
 - B. Group policy settings are refreshed on member servers every 90 minutes. To force the server to refresh the group policy, use the `secedit /refresh` command.
 - C. Group policy settings are refreshed on servers every 5 minutes. To force the server to refresh the group policy, use the `gpupdate` command.
 - D. Group policy settings are refreshed on servers every 90 minutes. To force the server to refresh the group policy, use the `gpupdate` 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. B | 10. A |
| 3. B | 11. A, C |
| 4. C | 12. D |
| 5. A, B | 13. A |
| 6. B, C | 14. D |
| 7. D | 15. D |
| 8. D | |

MCSE 70-293

Planning, Implementing, and Maintaining the TCP/IP Infrastructure

Exam Objectives in this chapter:

- 2 Planning, Implementing, and Maintaining a Network Infrastructure
 - 2.1.2 Plan an IP routing solution.
 - 2.2.2 Identify network protocols to be used.
 - 2.1 Planning Network Traffic Management
 - 2.1.1 Plan a TCP/IP network infrastructure strategy.
 - 2.1.3 Create an IP subnet scheme.
 - 2.6 Troubleshoot TCP/IP addressing.
 - 2.6.1 Diagnose and resolve issues related to client computer configuration.
 - 2.6.2 Diagnose and resolve issues related to DHCP server address assignment.
 - 2.2 Plan and modify a network topology.
 - 2.2.1 Plan the physical placement of network resources.
 - 2.4 Plan network traffic monitoring. Tools might include Network Monitor and System Monitor.

Introduction

The Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite is the foundation upon which the Internet runs, and it is the protocol suite of choice for most large, enterprise-level networks today. TCP/IP is the default network and transport protocol stack for a Windows Server 2003 network, and it is important for all network administrators to be intimately familiar with the TCP/IP protocols, IP addressing, and how to plan an IP infrastructure.

This chapter deals with the TCP/IP infrastructure. You will learn about the network protocols supported by Windows Server 2003 and how to identify the protocols to be used in your network environment. We discuss the advantages of the TCP/IP protocol suite, and we also address the multiprotocol environment that is increasingly common in today's business organizations. We will review TCP/IP basics, and then get into what's new in TCP/IP for Windows Server 2003. Specifically, we'll discuss Internet Group Management Protocol version 3 (IGMPv3), IP version 6 (IPv6) support, the alternate configuration feature, and automatic determination of interface metrics.

You'll find out how to plan an IP addressing strategy, including how to analyze your addressing requirements and how to create an effective subnetting scheme. Then we will address methods for troubleshooting IP addressing problems, both those related to client configuration and those related to Dynamic Host Configuration Protocol (DHCP) server issues. You'll learn about transitioning to the next generation of IP, IPv6, and we'll introduce IPv6 utilities such as Netsh, IPsec, PING, and Tracert. We'll discuss 6to4 tunneling, the IPv6 Helper service, and connecting to the 6bone.

Next, we'll discuss the planning of the network topology. This includes analyzing hardware requirements and planning for the placement of physical resources. You'll learn how to plan network traffic management, as well as how to monitor network traffic and devices using Network Monitor and System Monitor. We'll show you how to determine bandwidth requirements and how to optimize your network's performance.

Understanding Windows Server 2003 Network Protocols

EXAM
70-293
OBJECTIVE
2
2.1.2

In order for computers to communicate with other computers and hardware resources, they must use a common messaging structure, much like a language used to speak to other network devices. There are many message structures that are standardized and designed to provide reliable, continuous, high-speed data transfer and remain independent of the device or computer's hardware and operating system. When planning a network, you need to understand how the computers will share and access information and resources on the network so that you can decide which network protocol is best suited for the task.

The networking architecture of Windows Server 2003 uses the Network Driver Interface Specification (NDIS). NDIS provides a kind of wrapper in the I/O Manager layer of Windows that allows the hardware driver to be independent of the protocols used to

communicate on your network. Additionally, this allows for multiple network adapters with virtually any device driver, without having any effect on the transport protocols used. Let's take a look at some of the details involved with networking.

EXAM
70-293
OBJECTIVE
2.2.2

Identifying Protocols to Be Used

Network protocols are composed of software components designed specifically for communication with other networked machines. A variety of protocols are used for different functions. In order to be able to select protocols for particular tasks, it is important to understand how protocols facilitate network communication.

The first concept to understand is the standard model for network communications, known as the Open Systems Interconnection (OSI) reference model. The International Organization for Standardization (ISO) developed the OSI reference model. One of this organization's responsibilities is to provide a standard by which computers can communicate worldwide, and the OSI reference model was designed to accomplish this goal.

The OSI model is based on a concept of a stack of protocols that work together to provide the means for transmitting data. The OSI model is composed of seven layers of protocols responsible for different tasks related to data transmission. Each layer of the OSI model describes a function:

- **Application layer** Defines how applications work together over the network.
- **Presentation layer** Provides a common data format for the data transmitted.
- **Session layer** Coordinates the establishment of the connection and maintains the open connection.
- **Transport layer** Provides the mechanism for ensuring error-free delivery of data.
- **Network layer** Provides the addressing for messages for all networks.
- **Data Link layer** Defines the methods for the software drivers to access the hardware that is the physical medium, such as the network jack and the cable that plugs into it.
- **Physical layer** Puts the data on the physical medium that is carrying the data.

The data that you transmit is broken up in manageable chunks called *packets*, which will be transmitted as a single unit via the OSI layers. As the packet is passed down through each layer, information is added to aid the delivery of the packet to the corresponding layers on the destination machine. The protocols that work together to provide the packaging, delivery and receipt of the data at each of these layers are known collectively as a *protocol stack*.

The size of your network and the *topology*—how it is physically laid out—have a bearing on which protocol stack will be suitable for your needs. If you have a large network, the volume of traffic may need to be managed. One of the most common solutions

to traffic problems is to break up the network into smaller, more manageable networks known as *segments*. In order for you to combine these smaller networks together, the data must be able to travel from one network to the other along one or more physical paths. This transmission of data across network segments is called *routing*. In this situation, you would need to select a protocol stack that provides routable protocols. The two most common routable protocols are TCP/IP and Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX).

Some protocol stacks are composed of nonroutable transport protocols such as NetBIOS Extended User Interface (NetBEUI). These protocol stacks are simple to configure and implement, but they are suitable only for small networks that are not segmented. Nonroutable protocols limit your capabilities to expand your network. Your network might be required to interact with another existing network. In that case, the ability of protocols in the stack to route traffic becomes an issue.

There are other factors to consider when deciding which protocol to use. One consideration is reliability. Is it necessary to ensure the transmission of data? Some protocols just broadcast data and don't ensure that the data is received by the targeted machine. You should also consider the number of machines you have on your network. If you have a lot of devices, you may need to consider a more scalable, enterprise-capable protocol suite that includes transport protocols that provide guaranteed delivery and data integrity.

Another issue is the security of the data. You may have applications that transmit sensitive data over the network. You would not want anyone to be able to monitor that traffic and view the data. This would lead to selecting a protocol stack that provides protocols that allow you to encrypt the data and potentially validate the source and authenticity of the data. An example of this would be Internet Protocol Security (IPSec).

You may need to provide network services to the clients. These network services should be easy for the clients to access. In addition, the network services may provide a variety of functions, including data access, license services, printing services, and remote-access services over modem dial-up. Network services may also provide special communication features related to the medium, such as Infrared Data Association (IrDA), Asynchronous Transfer Mode (ATM), and other features that are used to remotely access networks using existing network connections called *virtual private networks* (VPNs). In order to leverage these services, the protocol stack should provide protocols that support the required features. For instance, VPNs use Point-to-Point Tunneling (PPTP) or Layer Two Tunneling Protocol (L2TP) to establish connections, both of which require TCP/IP to connect over the Internet. Once connected, PPTP and L2TP provide a channel to encapsulate and transmit other protocols such as TCP/IP, IPX/SPX, and NetBEUI as if you were on the same physical segment as other machines on the physical local area network (LAN).

Additionally, accessing the Internet is rapidly becoming a necessity for most businesses today. In order to access the Internet, you must have TCP/IP. You may provide other means to transmit data on your network, but Internet resources support only the TCP/IP protocol stack.

On the Windows Server 2003 platform, there are two basic choices of protocol stacks: TCP/IP and IPX/SPX. There are other protocol stacks; however, TCP/IP and IPX/SPX are the most prevalent, and they support such a robust suite of protocols that it is not usually necessary to use any others.

Considerations for Selecting Appropriate Network Protocols

There are many factors to consider when you decide which protocol or protocols to implement on your network:

- **Security** Do you transmit sensitive data between machines?
- **Reliability** Will you use applications that ensure delivery and integrity of the data you transmit?
- **Ease of implementation and maintenance** What is the total cost of ownership (TCO) to implement and maintain the selected protocols? Consider the cost of equipment, training, management, implementation, and future growth of the organization.
- **Traffic** Routable network protocols allow for better management of traffic and isolate broadcast traffic, which, in turn, reduces the amount of unnecessary data that must be handled by the network hardware.
- **Number of devices** How many machines will communicate on your network segments? Too many machines on one segment could overload your network and cause slower and more unreliable data transfer. How scalable is the protocol? Does it work well on a small network and does it allow for growth?
- **Physical topology** Are you implementing or integrating with a LAN, metropolitan area network (MAN), or wide area network (WAN)? What protocols are you currently using? Where are the machines physically located in relation to other machines on your network?
- **Function** Will you need to provide access to the Internet for users or systems on your network? Do you want to prevent access to the Internet? Will there be an intranet? Will you need to provide a stream of data to multiple destinations?
- **Existing protocols** Do you have a requirement to access resources on your network using existing protocols?

Advantages of the TCP/IP Protocol Suite

The TCP/IP protocol suite has many advantages over other protocol suites like IPX/SPX and NetBEUI. These advantages are due to TCP/IP's robust, stable, extensive feature set, combined with its scalability. The suite of protocols and services that are part of the TCP/IP

standards provide valuable and prevalent services, so it is no wonder that it is the default protocol for virtually every client and network operating system in use today!

One of TCP/IP's key advantages is that it is an open, industry-standard set of protocols, which implies that there is not one single organization that controls the standards. Novell provides IPX/SPX, which makes it a vendor-specific protocol. This implies that it was developed specifically to support NetWare's architecture and may not be as robust, or even supported, on other platforms.

TCP/IP contains applications that aid in connecting different operating systems, which ensures that you will be able to communicate in prescribed methods with any system that is using TCP/IP. The architecture of TCP/IP provides scalability—the means for sizing the network so that you can expand or shrink the network as your needs change.

Another advantage of TCP/IP is that it is routable. Routable protocols can reduce network traffic by isolating logical and physical networks. Isolating the networks allows you to better manage network traffic, direct transmissions, and restrict the distribution of broadcast traffic. You can also leverage the information about the routes from one point to another to troubleshoot and isolate problems with connections. You can use dynamic routing features to reroute traffic to prevent interruption of communications. The U.S. Department of Defense Advanced Research Projects Agency (DARPA) intended for TCP/IP WANs to provide a means for ensuring reliable communications in the event that portions of the network become unavailable. Routable protocols provide forms of addressing, such as an IP address, that define mechanisms for determining how to transmit data across network segments.

Nonroutable protocols do not use addressing, so there must be a means for determining the destination for data transmitted. An example is NetBIOS naming, which provides the architecture for defining the destination resource and the methods for sending the data to one or multiple destinations. NetBIOS naming provides for two types of names: a unique name and a group name. The unique name defines the station on a network, enabling you to connect to and communicate with the server. Group names are used to provide a means to send messages to multiple machines at once, but only those machines that are part of that group will listen to the messages.

TCP/IP was designed to be platform-independent. This allows you to connect and integrate different operating system platforms and hardware, and they will be able to communicate effectively, regardless of the platform. In addition, the suite of protocols includes methods for accessing data and resources on the various platforms, such as Line Printer Daemon (LPD) for printing, File Transfer Protocol (FTP) for file exchange, and Hypertext Transfer Protocol (HTTP) for sharing platform-independent documents, images, and other media.

The Internet and the World Wide Web are accessible only via TCP/IP. There is a vast amount of educational, research, and entertainment resources available to the world on the Internet. It is also possible to use the Internet to interconnect networks around the world. You can log on a network in London from an office in Dallas and function as if you were

in the London office. You can imagine how scalable the protocol is if it can service both a small office and the whole world.

Microsoft adds the ability to develop applications that leverage these advantages by providing the Windows Sockets API (WinSock). WinSock makes it possible for developers to design scalable TCP/IP client/server applications that can interoperate with any machines that use TCP/IP on any operating system platform. Since virtually every modern operating system uses TCP/IP, this makes WinSock a very viable framework.



EXAM WARNING

There are many different protocols that make up the Microsoft TCP/IP protocol suite. Don't forget that it is a *suite* of protocols, which includes many different features that all leverage the TCP/IP protocol stack. The Microsoft TCP/IP protocol suite provides applications and protocol functions that are designed specifically for Windows enterprise networking and the Windows operating system platform. The TCP/IP protocol stack is the industry-standard, platform-independent set of protocols that work at various layers to communicate over networks.

The Multiprotocol Network Environment

Microsoft Windows Server 2003, like its predecessors, uses a layered network architecture. Since it is layered, it makes it possible to extend the functionality of networking Windows Server 2003 with third-party software components. The layered structure also provides the Windows Server 2003 platform with the ability to allow different protocols to communicate using the same structure and methods, so users can access data in the same fashion, regardless of what networking protocol is used.

For instance, it is possible for a Novell NetWare server using the IPX/SPX network protocol stack to be accessible to a Windows Server 2003 machine using IPX/SPX. You can also use Windows Explorer on the Windows Server 2003 computer to access files on the NetWare file server without requiring any special features. If you need to run Novell Directory Services (NDS) utilities on Windows, you must also install the NetWare Client Software, which uses the IPX/SPX protocol to access those services. This type of multiprotocol configuration makes integration of other systems possible using third-party software.

Windows Server 2003 products use the TCP/IP protocol stack by default. The following network protocols are supported on Windows Server 2003:

- **TCP/IP version 4** The default protocol for Windows Server 2003.
- **TCP/IP version 6** The next generation of TCP/IP.
- **IPX/SPX** Used by many networks running Novell NetWare.
- **AppleTalk** Provides the basis for Services for Macintosh and AppleTalk routing and seed routing support.

The Windows Server 2003 architecture that supports multiple protocols also allows multiple network adapters. Each adapter can use any combination of protocols or networking components, known as *binding*. It is also possible for you to change the order in which protocols are bound to the adapter. You can choose to move the most commonly used protocols on the client up to the top of the binding order to provide faster performance. For example, your LAN may access NetWare services using IPX/SPX and Windows networking services using TCP/IP. In this example, let's assume that the NetWare services are not used very often, so your primary network communications use TCP/IP. You can change the protocol binding order as follows:

1. Select **Start | Control Panel | Network Connections**.
2. From the menu bar, select **Advanced | Advanced Settings**.
3. On the **Adapters and Bindings** tab, move the primary connection (if there is more than one) to the top of the list.
4. Select a connection, and the bindings will be displayed for that adapter.
5. Under **File and Printer Sharing for Microsoft Networks**, select **Internet Protocol (TCP/IP)** and move it up to the top of the list using the arrows in the dialog box.
6. Under **Client for Microsoft Networks**, select **Internet Protocol (TCP/IP)** and move it up to the top of the list using the arrows in the dialog box.
7. On the **Provider Order** tab, move **Microsoft Windows Network** to the top of the **Network Providers** list.
8. Click **OK**.



TEST DAY TIP

Understand how to add different adapters, protocols, services, and clients. You should be able to differentiate between a client as a service and a protocol, as well as how to change the bindings and their order for each.

Using multiple protocols on your network might provide a degree of flexibility, but it can also make your job more difficult. If you use a protocol that generates a lot of unnecessary broadcast traffic, it could harm your overall network performance. From the client's perspective, network problems could be more challenging, because each protocol bound to the adapter will be attempted in the event that some network connections are unavailable.

When configuring protocols on your computer, it is always desirable to make the fewest possible changes on the client in order to simplify the administration of the network. On a TCP/IP network with more than 25 hosts, it is a good idea to implement a DHCP server. A DHCP server will allow you to define certain settings related to host name resolu-

Windows Server 2003 Networking Features

Windows 2003 Server products offer several networking features and protocols. There have been several changes to the protocols and services available. While some of the support for older protocols is no longer available, new support is available for the latest technologies. The following are some of the changes made to networking features:

- Wireless networking functionality supports the 802.11 standard and reduces the hassles associated with configuring wireless networking.
- 802.1x authentication is enabled by default. This enables tighter wireless security than is possible with Wired Equivalent Privacy (WEP) encrypted connections. 802.1x authentication also supports Extensible Authentication Protocol (EAP), which allows third-party security enhancements like smart cards and certificates.
- The NetBEUI and Data Link Control (DLC) protocols are no longer available on any Windows Server 2003 products.
- The 64-bit versions of Windows Server 2003 products do not include IPX/SPX or any of the IPX-related services, nor the routing protocol, Open Shortest Path First (OSPF).
- The infrared (IR) networking feature is supported only on Windows Server 2003 Standard Edition.
- Gateway Services for NetWare (GSNW) is not included in Windows Server 2003 products.
- Windows Server 2003 products cannot act as IPX routers.
- It is not possible to uninstall the TCP/IP protocol on Windows Server 2003. Since you cannot uninstall and reinstall, there is a new netsh command that is used to reset the TCP/IP stack. To use this command, from the command line, type **netsh ip interface reset**.

tion and topology, and automatically provide the proper address for the hosts that are configured to use DHCP. By default, all Windows XP and Windows Server 2003 machines are configured to use DHCP.

Occasionally, you might need to manually configure the IP address of your machine. The following are some servers or services that may require a static (manually configured) IP address:

- A DHCP server
- Windows Routing and Remote Access Services (RRAS)

- Domain Name System (DNS) and Windows Internet Name Service (WINS) servers
- Any other service that provides IP functionality on your servers

If you do configure the address manually, pay close attention to the information you provide in the dialog box. Errors in the configuration will hinder network communication for that machine, and in some cases, cause problems that could prevent other machines from functioning properly.

EXERCISE 3.01

CONFIGURING THE TCP/IP PROTOCOL MANUALLY

In the following exercise, you will learn how to configure an IP address manually on a Windows Server 2003 computer.

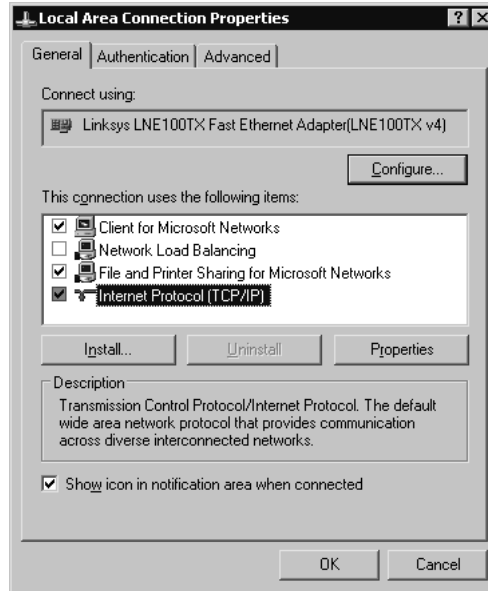
1. Open the **Local Area Connection Status** dialog box by clicking **Start | Control Panel | Network Connections** and double-clicking the appropriate local area connection. You will see the dialog box shown in Figure 3.1.

Figure 3.1 Local Area Connection Status




2. Click **Properties** to open the **Local Area Connection Properties** dialog box, shown in Figure 3.2.

Figure 3.2 Local Area Connection Properties



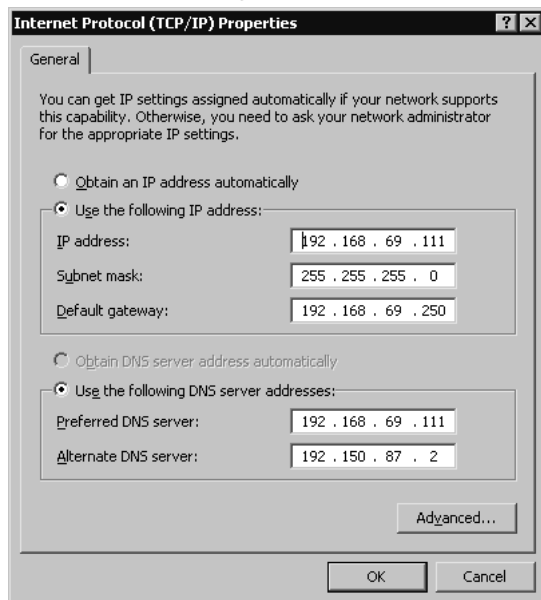
3. Click **Internet Protocol (TCP/IP)**, and then click **Properties** to open the **Internet Protocol (TCP/IP) Properties** dialog box, shown in Figure 3.3.

Figure 3.3 Internet Protocol (TCP/IP) Properties



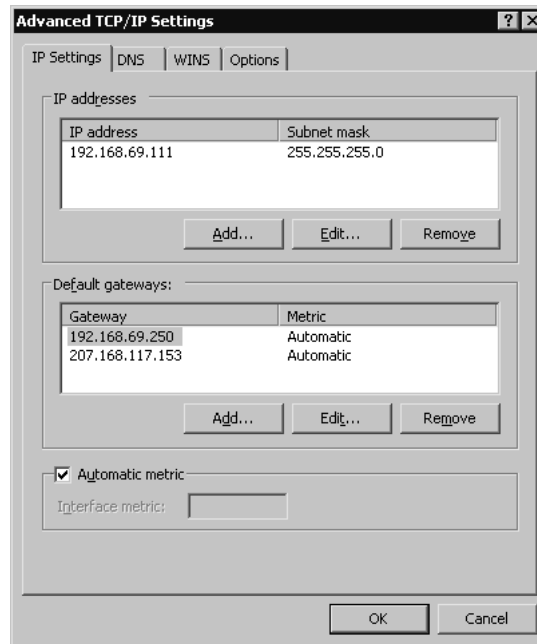
4. Click the **Use the following IP address** radio button and provide the **IP address**, **Subnet mask**, and **Default gateway**, as shown in Figure 3.4.

Figure 3.4 Internet Protocol (TCP/IP) Properties after Manual Configuration



5. Click the **Use the following DNS server addresses** radio button in the **Internet Protocol (TCP/IP) Properties** dialog box and provide at least one DNS server IP address (see Figure 3.4).
6. Click **Advanced** to open the **Advanced TCP/IP Settings** dialog box, as shown in Figure 3.5.
7. Notice the new **Automatic metric** option. Note that it is the default for all **Default gateways**.
8. Click **OK**.

Figure 3.5 Advanced TCP/IP Settings



New & Noteworthy...

Internet Control Message Protocol (ICMP) Router Discovery

ICMP is a maintenance protocol that is part of the IP layer in the Microsoft TCP/IP stack. Its functions include providing diagnostics, leveraging the use of the PING utility, and managing flow control of data to prevent traffic from saturating network links or routers. It also provides the facility that builds and maintains the routing tables, as well as determines the size of the packets that will be sent to a destination.

RRAS on Windows Server 2003 supports a new feature called *ICMP router discovery*. ICMP router discovery uses ICMP messages to “discover” the routers on the current subnet and select one to act as the default gateway. This allows DHCP clients to find a default gateway when one is not specified by the DHCP server.

This feature is disabled by default on Windows Server 2003 and Windows XP machines. In order to enable a DHCP client to perform router discovery, the client must receive a “perform router discovery” option from a DHCP server. This will enable the host to broadcast the request to all available routers. You must also set the option to **Enable router discovery announcements** on the **General** tab of the Windows Server 2003 **RRAS Properties** dialog box in order for the router to send the router advertisements.

To view your current IP configuration, you can run **ipconfig** from the command line. For more detailed information, use **ipconfig /all**. If you want to release your DHCP-assigned IP addresses from all adapters, use **ipconfig /release**. You can obtain a new lease with **ipconfig /renew**. For the **release** and **renew** commands, you can also specify the name of a specific adapter.

Reviewing TCP/IP Basics

TCP/IP on Windows Server 2003 provides a scalable, robust client/server platform that is built on industry-standard, routable, and full-featured protocols. Virtually every network operating system supports the TCP/IP protocol stack, and this allows Windows Server 2003 to integrate dissimilar systems on the network. The various protocols that make up the TCP/IP stack work together to provide network communications. These network communications provide the architecture that the Windows Server 2003 TCP/IP suite uses to leverage services such as name resolution, file transfers, and Internet access.

Every implementation of TCP/IP must follow the guidelines that are governed and managed by several agencies such as the Internet Architecture Board (IAB) and the Internet Engineering Task Force (IETF). The IAB is also responsible for managing several other groups, such as the Internet Society (ISOC), Internet Assigned Numbers Authority (IANA), and Internet Corporation for Assigned Names and Numbers (ICANN). These agencies work together to maintain an open standard using a process known as Request for Comments (RFCs) and provide the maintenance, distribution, and administrative handling of the RFCs. For information on RFCs, access the IETF Web site at www.ietf.org.

Virtually all network protocols can be mapped to the ISO's OSI reference model. The OSI model is intended to provide a general direction for developers for designing network drivers and protocols. The design intends for different components involved with network communication to be managed in a series of layers, with each layer built on top of another, having a specific set of functionality, and communicating with the adjacent layers. The layers allow for a hardware manufacturer to design a network card without regard to the operating system or applications that will be using the network card to communicate. A developer can design a client/server network application without concern for the protocols used to communicate with other machines.



EXAM WARNING

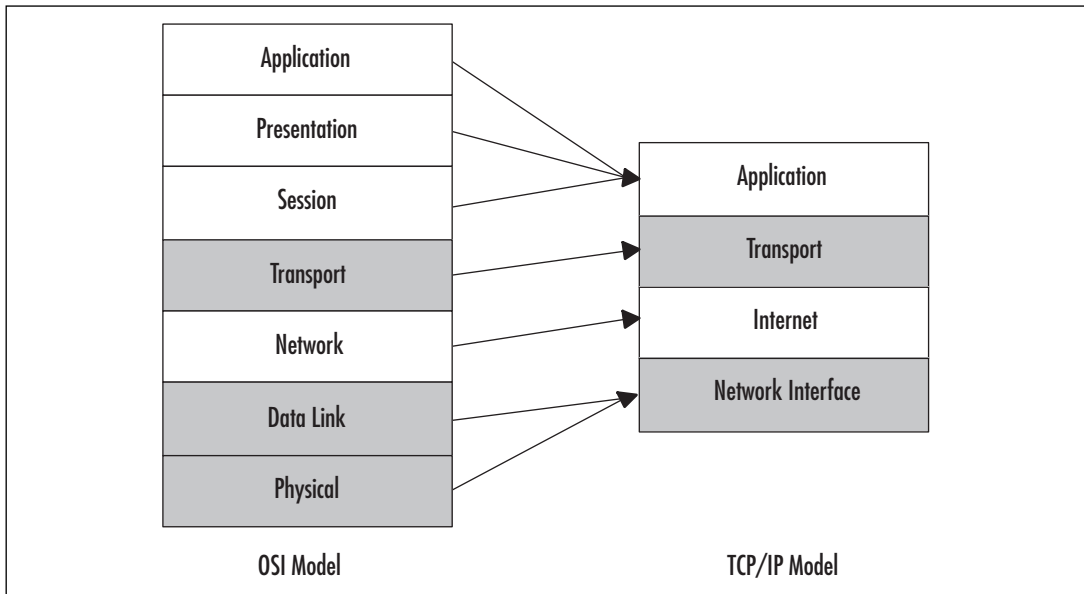
ISO is the organization that defines the standards for the OSI model. The IAB is responsible for facilitating the rules and the processes for the standards that define the Internet. The standards for the Internet are maintained by the IETF and are called RFCs.

TCP/IP uses a slightly less complex networking model that was developed by DARPA. Since the model is less complex than the OSI model, it is easier to implement and has

better performance characteristics. The DARPA model and the TCP/IP suite of protocols were designed by DARPA before the development of the OSI model. The Windows implementation of the TCP/IP protocol stack relates to the seven layers of the OSI model.

The layers in the TCP/IP model span several layers of the OSI model. As shown in Figure 3.6, the Application, Presentation, and Session layers of the OSI model are incorporated into the Application layer of the TCP/IP model. Some of the components of the TCP/IP protocol suite that operate in this layer are FTP, Telnet, HTTP, and DNS. The Application layer provides the access to the network for many applications, such as Microsoft Internet Explorer. At this layer, presentation issues such as compression and encryption are handled, and sessions are established (if applicable). Then the sending computer passes the data down to next layer, the Transport layer.

Figure 3.6 OSI Model versus TCP/IP model



The Transport layer coordinates the applications' communication sessions with other interconnected machines. The key protocols that operate at this layer are TCP and User Datagram Protocol (UDP). TCP differs from UDP in two key ways. The first distinguishing difference is that TCP is connection-oriented and UDP is connectionless. TCP expects acknowledgment from the other host for each packet of data transmitted. This is ideal for large data transfers over very large networks. FTP uses TCP ports 20 and 21 to transfer data. Because UDP is connectionless, it doesn't guarantee the delivery of the data; it just makes its best effort to deliver the packets intact. This type of data transfer is ideal for lightweight, small data transfers on a well-connected network. Trivial File Transfer Protocol (TFTP) uses UDP port 69 to initiate a connection, and the server will then dynamically

select a port number to return data from. Then the two machines continue to communicate using the new port numbers. Since TFTP uses UDP, it is well-suited for small files, such as short text files, and is faster than using FTP over TCP, since there is less overhead.



TEST DAY TIP

UDP and TCP both use the IP protocol. TCP is directed to the destination and ensures the delivery of packets by receiving acknowledgments of data delivery. UDP attempts a best-effort delivery of the datagram and does not guarantee delivery. UDP has less overhead, so it is much faster, but it is not as reliable as TCP. Both TCP and UDP use ports to differentiate between communications to and from different applications.

On a sending computer, the Transport layer passes the data down to the Internet layer. The Internet layer, which maps to the OSI model's Network layer, is responsible for addressing and routing communications over the network. IP operates here, and it is responsible for determining whether the address of the destination computer is on the same subnet as the address of the sending computer. In order to physically locate another host on the network, Address Resolution Protocol (ARP) is used for IP address-to-Media Access Control (MAC) address resolution. Other protocols that operate at this layer are ICMP, IGMP, and IPSec. The Internet layer continues the communication process by passing data to the Network Interface layer.



EXAM WARNING

ICMP provides diagnostics and error reporting. The PING utility uses ICMP to send and receive a standard packet to determine if the data delivery was timely and successful. ARP determines the physical address, or MAC address, of the destination host. IP determines whether the address is local or remote. If the address is local, it will direct ARP either to refer to its local cache or broadcast on the local subnet to resolve the MAC address. If it is determined by IP that the address is not local, ARP will resolve the MAC address of the default gateway to allow the traffic to be routed to the appropriate network.

If you are using Internet Connection Firewall (ICF) or any other firewall software, you may prevent PING from functioning if you have defined any settings or filters that block ICMP traffic. By default, ICMP traffic is disabled when you enable ICF.

The last layer (when data is being sent) is the Network Interface layer. This layer maps to the Physical and Data Link layers of the OSI model. It is responsible for the software driver-to-hardware translation and complying with the hardware communication standards such as Ethernet, ATM, and Token Ring. The Network Interface layer is isolated from the

hardware on Windows 2003 Servers by NDIS, which is a boundary layer implemented in the Microsoft networking model. This allows the protocols to function independently of the network hardware. The MAC address is part of this layer.

The Microsoft networking model corresponds to different services in the Windows architecture, to provide similar ways to access data independently of the mechanism. For instance, using Windows Explorer to access files using IPX/SPX does not seem any different to the user than accessing files using TCP/IP. Network-aware applications and network service providers operate as User mode services at the Application layer and the top of the Presentation layer.

The Presentation layer transitions data back and forth from User mode to Kernel mode. The Executive services provide Session support and transition data to the I/O Manager. The Server and Redirector (Workstation) services operate at the Session layer and are separated from the transport protocols by the Transport Driver Interface (TDI) boundary layer, which traverses the Session and Transport layers.

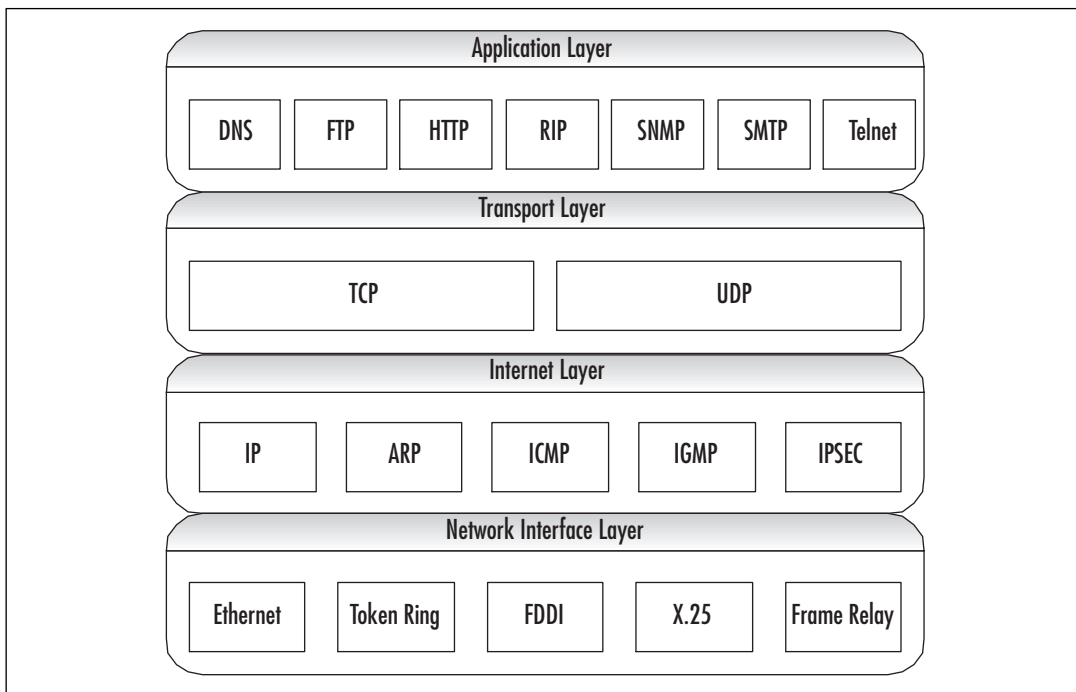
The transport protocols, such as TCP/IP and IPX/SPX, transition from the Transport layer, over the Network layer and down to the Network Interface or Data Link layer. The Data Link layer is where NDIS accesses the network adapter drivers before passing the data to the Physical layer, which allows the different protocols to be bound to different network adapters, using different physical connections.

Figure 3.7 illustrates the TCP/IP protocol suite in the TCP/IP model.

NOTE



Although we started “at the top” in describing the layers of the TCP/IP (DoD) model, it is important to remember that when they are numbered, they are referred to in reverse order (as in the OSI model). The Network Interface layer is layer 1, the Internetwork layer is layer 2, and so forth.

Figure 3.7 TCP/IP Protocol Suite and the TCP/IP Network Model

Each layer in the protocol stack provides a translation or some form of communication with the next layer. As data is passed down through the stack, each layer adds its necessary headers and protocol-specific data, and encapsulates the data from the previous layer. In some instances, the layer will establish a session with the destination host at the same layer. Once the data reaches the destination, each layer in the protocol stack will validate the header that was added by its corresponding layer, and then strip the protocol-specific information from the packet and pass it up to the next layer until it reaches the destination application.

What's New in TCP/IP for Windows Server 2003

There are many enhancements to the networking and communications components of Windows Server 2003. The TCP/IP protocol suite has been enhanced with some of the latest technologies, as well as improvements on existing functionality. For more information about other networking and communication feature enhancements, see the white paper titled "Microsoft Windows Server 2003- Technical Overview of Networking and Communication" (www.microsoft.com/windowsserver2003/techinfo/overview/net-comm.msp).

IGMPv3

Typical communications over an IP-based network are directed unicast communications. Unicast is basically a single, direct request sent from one host to another, and only the two hosts interact over the established route. For example, when you click a hyperlink in a Web browser, you are requesting HTTP data from the host defined in the link, which, in turn, delivers the data to your browser. This is useful in the Web-browsing environments we have grown accustomed to, where there is a demand for a personal, user-controlled experience.

Unicast is not useful for delivering streams of audio or video to large audiences, since a single stream of audio/video data is very costly for only one user. This is where multicast communications are effective. Multicast provides a single stream for multiple hosts. The hosts select the data by requesting the local routers to forward those packets of data from the host providing the multicast data to the subnet of the listening host. When the host decides to stop listening to the multicast traffic, IGMP is responsible for notifying the router that the host is no longer participating.



TEST DAY TIP

It is not necessary to know the differences between different versions of IGMP. It is important to be familiar with the purpose of IGMP, what its functions are, and where it fits in the OSI model.

A set of listening hosts is called a *multicast group*. IGMP is responsible for providing the functionality necessary for hosts to join and leave those groups that receive IP multicast traffic. Each of the versions of IGMP—versions 1, 2, and 3—is automatically supported by Windows Server 2003. IGMPv3 adds functionality to distribute multiple multicast sources regionally and allow the host to select the multicast source that is located closest to the host.

An example of this would be a situation in which you send a video stream broadcasting a speech from the president of your company and have several machines scattered across the United States providing the feed. Then IGMPv3 allows the hosts to provide an *include* list or an *exclude* list of those servers. The multicast routers would be responsible for forwarding the multicast traffic from the include list of servers and for preventing the forwarding of traffic from the excluded sources. As you can see, this feature can be very useful to help reduce network bandwidth utilization.

IPv6

The next generation of TCP/IP is here! Previously, it was possible to experiment with IPv6, but under the covers, the protocol stack was still dependent on IPv4 calls for WinSock functions. With the release of Windows Server 2003, the IPv6 protocol stack is designed for production use.

IPv4 has a limited number of host addresses available (2^{32} , or about 4 billion hosts). That might sound like a lot, but over the past 30 years, the pool of available addresses has been exhausted due to the popularity and growth of the Internet. With IPv6, the host address is 128 bits instead of 32, which means that we will have 2^{128} (about 340,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000) host addresses available. That means we could have about 2^{96} (about 75 trillion trillion, or 75,000,000,000,000,000,000,000,000,000,000) addresses of our very own. That should last for at least a couple of years. We will discuss transitioning to IPv6 and its features in more detail in the “Transitioning to IPv6” section later in this chapter.

Alternate Configuration

Automatic alternate configuration is an enhancement to TCP/IP that allows for a valid static IP address configuration on a DHCP-configured machine. Without an alternate configuration defined, a computer that is unable to obtain an IP address lease from a DHCP server will automatically receive an Automatic Private IP Addressing (APIPA) address from the 169.254.0.0/16 pool.

Using APIPA to Your Advantage

APIPA can be a valuable aid in assisting you with network configuration. With no effort at all, you can provide IP addressing for a TCP/IP network of Windows Server 2003 and Windows 98/2000/XP computers. APIPA is service that uses a reserved class B IP address pool (169.254.0.0/16 or a subnet mask of 255.255.0.0) to automatically provide valid IP addresses to DHCP clients in the event the computer cannot obtain a DHCP lease. This scheme is intended for smaller networks where there is no DHCP server deployed, but think of the potential use this has, not only as a way to assist your LAN users, but also to help you troubleshoot network problems and configure new servers.

One way you can help LAN users is to provide an intranet Web server that has been assigned an APIPA address. That way, if a client is unable to obtain an IP address, the user will be able to connect to this Web server. The Web server’s default home page should contain a series of simple troubleshooting procedures that the client could use, such as the following:

- Did you receive an error message on startup? Provide a list of common errors and probable solutions.
- Wait 5 minutes to see if the next DHCP request is acknowledged.
- Contact technical support at extension 5555.

Continued

Additionally, you could provide users with some basic information about what is happening or maintain a server status page to let them know that you are aware of the problem and what actions they should take. It might also be beneficial to the Information Technology (IT) staff to maintain documentation on the Web server to aid in configuring new servers, maintaining static address pools, or initiating service requests to add new equipment to the network.

Automatic Determination of Interface Metric

As noted in Exercise 3.01, “Configuring the TCP/IP Protocol Manually” and shown earlier in Figure 3.5, the automatic metric feature is enabled by default. The purpose of the automatic metric feature is to determine the speed of the interface for each default gateway and to assign the *metric*, which is the cost of using a particular route.

The metric is weighted by the number of hops to the destination. The number of hops to any host on the local subnet is one. Every router that must be used to reach the destination is another hop. When it is determined that there are multiple routes to the same destination, the metric is evaluated to determine which is the lowest metric and this the fastest route to the destination.

EXERCISE 3.02

DETERMINING THE METRIC FOR THE DEFAULT GATEWAY

In the following exercise, you will learn how to use the **route print** command to determine the metric for the default gateway on your network.

1. Open a command prompt window.
2. Type **route print**. You will see a route table, as shown in Figure 3.8.

Figure 3.8 Results of the route print Command

```

D:\>route print

IPv4 Route Table
=====
Interface List
=====
0x1 ..... MS TCP Loopback interface
0x10003 ...00 04 5a 44 c4 9f ..... Linksys LNE100TX Fast Ethernet Adapter(LNE10
0TX v4)
=====
Active Routes:
=====
Network Destination        Netmask          Gateway          Interface        Metric
-----
0.0.0.0                    0.0.0.0          192.168.69.250   192.168.69.111   20
127.0.0.0                  255.0.0.0        127.0.0.1       127.0.0.1        1
192.168.69.0               255.255.255.0    192.168.69.111  192.168.69.111   20
192.168.69.111            255.255.255.255  127.0.0.1       127.0.0.1        20
192.168.69.255            255.255.255.255  192.168.69.111  192.168.69.111   20
224.0.0.0                  240.0.0.0        192.168.69.111  192.168.69.111   20
255.255.255.255           255.255.255.255  192.168.69.111  192.168.69.111   1
Default Gateway:          192.168.69.250
=====
Persistent Routes:
None
D:\>_

```

3. Examine the route table.
4. Notice the **Network Destination** list. The destinations are described in Table 3.1.

The metric for the loopback adapter and the limited broadcast is always 1. The other addresses have a metric based on the cost of using that route for that network adapter. With multiple network adapters, a multihomed computer, the route table would indicate a different metric for each default route, but only one would be used. Table 3.2 shows a configuration with identical network adapters: one adapter on the 192.168.69.0/24 network and the other on the 192.168.70.0/24 network.

Table 3.1 Description of Routes in the Route Table

Description	Network Destination	Netmask	Gateway	Interface	Metric
Default route	0.0.0.0	0.0.0.0	192.168.69.111	192.168.69.111	20
Loopback network	127.0.0.1	255.0.0.0	127.0.0.1	127.0.0.1	1
Local network	192.168.69.0	255.255.255.0	192.168.69.111	192.168.69.111	20
Local IP address	192.168.69.111	255.255.255.255	127.0.0.1	127.0.0.1	20
Subnet broadcast	192.168.69.255	255.255.255.255	192.168.69.111	192.168.69.111	20
Multicast address	224.0.0.0	240.0.0.0	192.168.69.111	192.168.69.111	20
Limited broadcast	255.255.255.255	255.255.255.255	192.168.69.111	192.168.69.111	1

Table 3.2 Description of Routes with a Multihomed Computer

Description	Network Destination	Netmask	Gateway	Interface	Metric
Default route	0.0.0.0	0.0.0.0	192.168.69.111	192.168.69.111	20
Default route	0.0.0.0	0.0.0.0	192.168.70.100	192.168.70.100	30
Loopback network	127.0.0.1	255.0.0.0	127.0.0.1	127.0.0.1	1
Local network	192.168.69.0	255.255.255.0	192.168.69.111	192.168.69.111	20
Local IP address	192.168.69.111	255.255.255.255	127.0.0.1	127.0.0.1	20
Local network	192.168.70.0	255.255.255.0	192.168.70.100	192.168.70.100	30
Local IP address	192.168.70.111	255.255.255.255	127.0.0.1	127.0.0.1	30
Subnet broadcast	192.168.69.255	255.255.255.255	192.168.69.111	192.168.69.111	20
Multicast address	224.0.0.0	240.0.0.0	192.168.69.111	192.168.69.111	20
Multicast address	224.0.0.0	240.0.0.0	192.168.70.100	192.168.70.100	20
Limited broadcast	255.255.255.255	255.255.255.255	192.168.69.111	192.168.69.111	1
Limited broadcast	255.255.255.255	255.255.255.255	192.168.70.100	192.168.70.100	1

Note that the metric for the default route for the second network, on the adapter for the 192.168.70.100 interface, is higher than the metric for the default route on the 192.168.69.111 interface. This indicates that the 192.168.69.111 network adapter is first in the binding order. Since the metric for the default gateway for the second adapter is higher than the first network adapter, the second gateway is never used and is not necessary.

You can use the **route** command to add routes and change metrics. The command is **route add -p Destination Mask Gateway IF Metric**, where:

- **Destination** is the network destination address.
- **Mask** is the appropriate subnet mask defined for the destination network.
- **Gateway** is the address of the router interface used to interface with the network.
- **IF** is the interface you want to associate this route to.
- **Metric** is the metric for this gateway.

The **-p** parameter specifies that you want to make this route persistent, so that it will be there if you reset the adapter or restart the machine. If you do not specify **-p**, the route is temporary and will not be saved.

If you want to delete a route, use the **route delete Destination** command to remove the destination route from the route table.

You can disable the automatic metric feature by accessing the properties for the desired connection, as follows:

1. Select **Internet Protocol (TCP/IP)** and click **Properties**.
2. In the **Internet Protocol (TCP/IP) Properties** dialog box, click the **Advanced** button.
3. Uncheck **Automatic metric**.
4. Provide an **Interface metric**. The minimum value is 1.
5. Click **OK**.
6. Run the **route print** command. What changed? You will notice that all of the metric values are now 1.

You can change the values manually, which can allow you to redirect traffic over a slower interface that would normally have a higher metric.



TEST DAY TIP

You should be familiar with the route table, know how to use the route print command, and understand how to use the information in this table to troubleshoot TCP/IP connectivity problems. More details are provided in the “Creating a Subnetting Scheme” and “Troubleshooting IP Addressing” sections later in this chapter.

EXAM
70-293
OBJECTIVE
2
2.1
2.1.2

Planning an IP Addressing Strategy

Before you can implement an IP network infrastructure, there are many details that you must consider. Here, we will take a look at how to plan your network by identifying the appropriate addressing requirements and limitations that will shape the network.

Understanding subnetting is a requirement to implement your addressing scheme. You will need to identify hardware requirements, decide what class of address you will need, and determine if access to the Internet is necessary for all or just some of your hosts.

Subnetting will allow you to create logical segments on your network that will overlay the physical topology. By using a well-planned subnetting scheme, you can handle your current needs and plan for expansion for future needs. You can also make use of these segments to isolate and distribute heavy traffic, without having a major impact on other segments of your network.

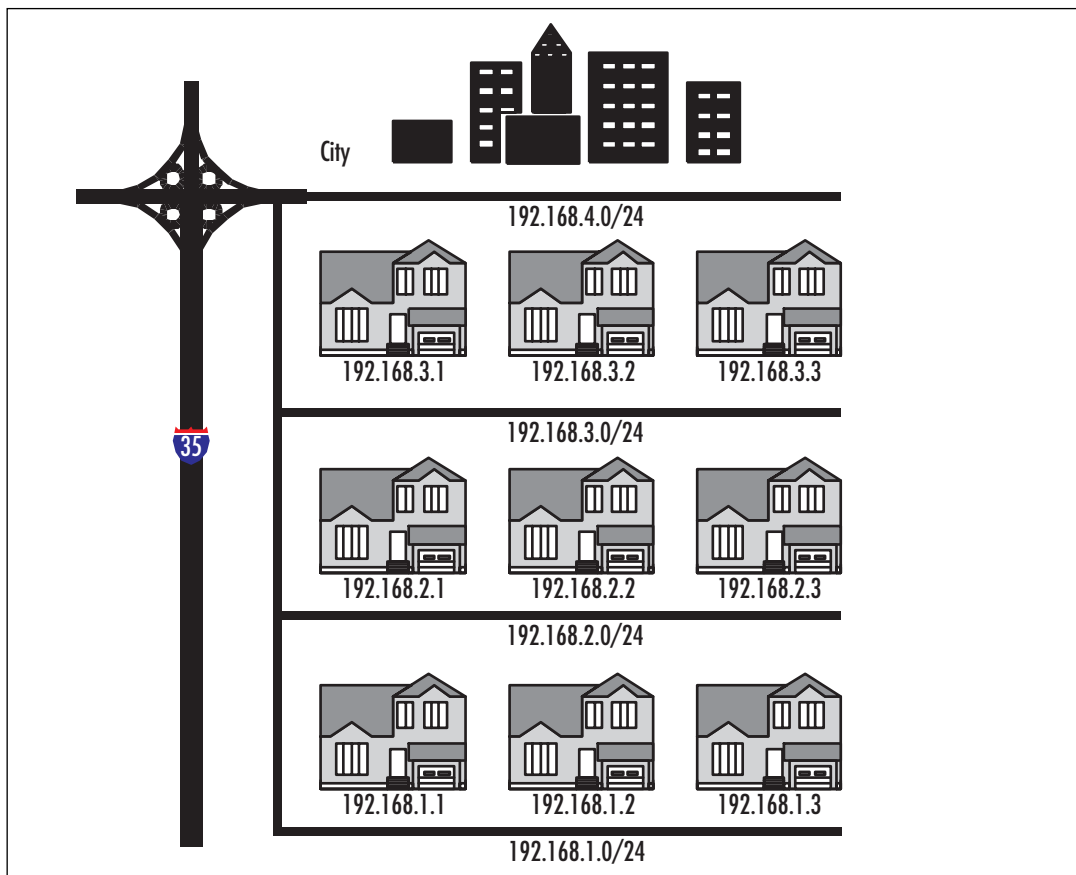
EXAM
70-293
OBJECTIVE
2.1.1

Analyzing Addressing Requirements

Every device on a TCP/IP-based network that has a network interface is referred to as a *host*. Each host must have a unique IP address. The most common analogy used to describe an IP addressing scheme is that of a street (subnetwork) with many houses (hosts). Each house (host) must have a unique address (IP address) on its street. Visualize a situation in which you are a city planner. In this analogy, the city is the entire corporate network, each street is a subnetwork, and each house is a host.

Our city, illustrated in Figure 3.9, needs streets for all of the houses for the current residents. Additionally, we might require more houses to be built for new residents. We must design the streets in such a way that will allow for traffic flow to be regulated and to minimize congestion. Also, we do not want to have so many streets that we can build only a few houses on each street before we run out of room in our city. We know that it might not be an effective use of our resources if we build a major thoroughfare and a lot of apartments in an area of the city that will have only a few residents. Some parts of our city need access to the “super highway”—the Internet or WAN—so the residents can get to other cities. We can use this example to get a concept of how to design and plan for a TCP/IP network.

Figure 3.9 IP City



Since the host IP address must be unique, the simple rule to calculate the number of hosts for our network is *one IP address per host*, plus one IP address for each additional network adapter in a host machine. We have a concept of one network in the corporate sense, but when determining address requirements, there are a few more details we must consider.

You can define IP addresses using one of the three classes available for standard IP communications: Classes A, B, and C. Before we decide which class to use, we need to determine the type of network we are implementing and how many hosts there are per segment.



EXAM WARNING

You should know the IP address classes and their ranges, the default mask for each class, and the number of hosts each class can support.

EXAM
70-293
OBJECTIVE
2.1.3

Creating a Subnetting Scheme

IP addresses are 32-bit values, often referred to as *dotted quads*. Each bit is a binary value of either 0 or 1. Since there are 8 bits, there are 2^8 combinations of 0 and 1, which equals 256 combinations, allowing for a range of 0 to 255. An address is broken down into octets consisting of four 8-bit sections. An address is usually represented by a decimal number such as 141.59.115.7, which is equal to the binary number of 10001101.00111011.01110011.00000111. Computers process only binary information, but we convert it to decimal because that is easier for us human beings to work with.

Classful Addressing

As mentioned, host addresses can belong to one of three classes of IP address, and each has a range of addresses. The range is defined by the value of the first octet. Table 3.3 shows the classes and their ranges, as well as the binary representations of the ranges. Classes D and E are also classes of IP addresses, but Class D is restricted to multicasting and Class E addresses are reserved for future use. 127.0.0.0 is reserved for connectivity testing. 127.0.0.1 is a special address that represents the local loopback adapter that resolves as *localhost*. We can ping the local host to troubleshoot the protocol stack. We will discuss this in more detail in the “Troubleshooting IP Addressing” section later in this chapter. Each class also has a default subnet mask.

Table 3.3 IP Address Classes and Their Ranges

Class	Range of Values	Default Mask	Networks	Hosts	Binary
A	0 to 126	255.0.0.0	126	16,777,214	00000001 to 01111110
B	128 to 191	255.255.0.0	16,384	65,534	10000000 to 10111111
C	192 to 223	255.255.255.0	2,097,152	254	11000000 to 11011111
D	224 to 239	Not applicable			Not applicable




TEST DAY TIP

In Table 3.3, notice that the first two bits of the first octet in each class also define the top of the range of network IDs for that class. If you take the first two bits of Class A, 01, and add the remaining six digits as ones you get 01111111, or 127. Remember that 127 is reserved, so 126 is the highest value for the network ID of a Class A network. Class B is 10 (10111111 = 191), and Class C is 11 (11011111 = 223).

The default mask for each class defines the number of networks and the number of hosts for each network. An IP address contains information about the network on which the host resides, and the address of the host. The network ID is the reference to the logical subnet, and it refers to the octets that are predefined as the network ID and implemented with the default mask. The remaining octets are for the hosts. Figure 3.10 illustrates the network and host IDs.

Figure 3.10 Network ID and Host ID



The first address in each network refers to “this network” (itself), such as 24.0.0.0/8 or 204.79.26.0/24. The last address in each network or subnetwork is the broadcast address for that segment, such as 179.54.255.255 or 204.79.26.255. We can derive the formula for determining the number of hosts per network as $2^n - 2$, where n is the number of bits available for host IDs. In Figure 3.10, we are using a subnet mask of 255.255.0.0, so the last two octets, or 16 bits, are available. If we plug that into the formula, we get $2^{16} - 2 = 65,534$ hosts per network.

Class A addresses are used for networks that have a large number of hosts. Based on the default mask, we have the first octet for networks and the last three for hosts. So, we have 126 networks and $2^4 - 2$ hosts, or 16,777,214. Likewise, with class B, the default mask is 255.255.0.0, so the first two octets are for the network IDs, for a total of 16,384, and the last two are for the hosts. So, class B networks have $2^{16} - 2$ hosts, or 65,534. Class C networks have more networks but are smaller, with $2^8 - 2$ hosts, or 254.

We could implement our network now very simply. Determine the number of hosts and the number of networks, and pick the class that fits. If you do not wish to assign a public IP address to all your machines, there is another solution. There are three banks of IP addresses that are called *private IP address* ranges. They are listed in Table 3.4. Typically, a network will need only one or two public addresses for the Internet interfaces, and everything internal to the company can use the private IP addresses internally.

Table 3.4 Private IP Addresses

Network ID	Subnet Mask	Range
10.0.0.0	255.0.0.0	10.0.0.1 to 10.255.255.254
172.16.0.0	255.240.0.0	172.16.0.1 to 172.31.255.254
192.168.0.0	255.255.0.0	192.168.0.1 to 192.168.255.254

Understanding ANDing and Binary Numbering

Once we define our subnetworks, the machines will need to communicate with other machines on the network. The determination of the host as a local or remote destination is derived by applying the subnet mask of the source host to the IP address of the destination. This process involves applying a Boolean logic method called ANDing. By ANDing the binary representation of an address and a subnet mask, the IP layer can determine if the address is on the same logical network or a different one.

In Table 3.5, we have a source and a destination host address. First, the subnet mask is applied to the source address using Boolean AND logic. To perform the AND operation, start from the left and compare each bit in the binary numbers representing the IP address and the subnet mask. If both are 1 (1 AND 1), then the result is 1; otherwise, the result is 0. After the comparison is performed with each address, if the resulting binary values are equal, then the addresses are on the same network; if they are not equal, then they are on different logical networks.

Table 3.5 Applying the Subnet Mask to IP Addresses

Source IP Address 172.16.5.16 Subnet Mask 255.255.254.0	Destination IP Address 172.16.2.251 Subnet Mask 255.255.254.0
10101100.00010000.00000101.00010000	10101100.00010000.00000010.11111011
11111111.11111111.11111110.00000000	11111111.11111111.11111110.00000000
10101100.00010000.00000100.00000000	10101100.00010000.00000010.00000000

We can use the default subnet masks to define our network, or we can use a custom subnet mask. The ability to define the subnet mask allows us to take the default network definition and “borrow” bits from the available hosts on that network in order to create smaller logical networks, or *subnets*.

EXERCISE 3.03**FUN WITH BINARY NUMBERS**

In this exercise, you will use the scientific mode of Windows Calculator to convert binary numbers to decimal numbers and vice versa.

1. Select **Start | Run** and type **calc** to launch Windows Calculator.
2. Select **View | Scientific**.
3. Make sure the **Dec** radio button is selected.
4. Using the keypad, enter the number **175**.
5. Click the **Bin** radio button. You should see 10101111.
6. In the edit box, type **11000111**.
7. Click the **Dec** radio button. You should see 199.
8. Type **75** in the edit box, and then click the **Bin** radio button.
9. Notice the binary number is 1001011. Count the number of bits. There are only 7 bits in the result. Calculator will strip leading zeros from binary values, so it is important to always “pad” the binary numbers to 8 bits when using them for IP address functions. The correct representation for 75 as an IP address octet is 01001011.
10. Use Windows Calculator to convert the binary representation of the following IP addresses to decimal IP addresses.

Binary	Decimal
11001010.01000101.01001111.00110101	
10001001.00001101.10101010.11111001	
11000111.01011111.01000000.10000001	
11000011.11011101.11101111.00000101	
00000111.11100010.00100000.11111101	
10000001.00100101.00001111.10110001	
10000011.01000100.00100000.00010110	

11. Use Windows Calculator to convert the following decimal IP addresses to a binary representation of IP addresses.

Decimal	Binary
192.178.44.121	
204.18.1.179	
10.2.2.76	

Continued

Decimal	Binary
141.22.94.107	
55.87.191.11	
187.34.59.199	
99.107.253.224	

You might begin to notice patterns with binary numbers. The values of each placeholder are similar to the decimal format, except that decimal is base 10. The first digit in a decimal number is 10^0 , or 1; the second is 10^1 , or 10; the third is 10^3 , or 100; and so on. 111 in decimal is equal to $100 + 10 + 1$.

In binary representation, each placeholder is base 2, so the first digit in a binary number is 2^0 , or 1; the second is 2^1 , or 2; the third is 2^2 , or 4; and so on. Thus, 111 in binary is equal to $4 + 2 + 1$, or 7. Table 3.6 shows a quick summary of one octet in binary.

Table 3.6 Binary Notation

128	64	32	16	8	4	2	1
2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0



TEST DAY TIP

Binary math got you down? Never fear, the standard Windows Calculator will be available for you to use during the exam. It is a good idea to be proficient with the use of the Calculator program in scientific mode, so that you don't have any doubts during the exam. Be very careful to count your digits in binary results. There are no leading zeros, so 1111111 is actually 011111111, 111111 is actually 001111111, and so on. Despite the convenience of using the Calculator program, you should still understand how to convert binary to decimal manually.

Subnetting Networks

Subnetting networks is necessary to efficiently manage network resources and control traffic on your network. When your network has grown beyond the capacity of your current infrastructure, you must change your configurations to support those changes. It is relatively simple to identify limitations that are obvious, such as the number of networks and hosts.

You can determine the number of networks by counting the number of physical locations that will need a router to connect them to other locations, such as another building or another floor in the same building. You can estimate the number of hosts needed per net-

work by counting all the IP-based resources in each physical location, including printers, desktops, servers, and other routers. Once you have that information, you can decide which class of network to use and how to break down that network into logical subnets that will be used to implement each physical or logical location. To summarize, there are three steps to subnetting:

1. Identify the number of hosts.
2. Identify the number of networks.
3. Use an assigned IP network ID or choose a private IP address, and then determine how to subnet your network.

As an example, suppose that we have 55 employees in one location, with 12 IP-based network printers, 6 servers, and 1 Internet refrigerator that orders the groceries in the break area when the stock is depleted. Our IP address block assignment provided by our Internet Service Provider (ISP) is 204.74.9.0/24. All the employees are currently located in one large, central area on the same floor. Since we have no physical boundaries to overcome, we use the default subnet mask. This would provide us with one network and 8 bits in the host portion of our address. The 8 bits give us 256 hosts, but the first host is 0, which refers to our network, and the last host is the broadcast address for the network, 255. Remember the formula is $2^n - 2$, where n is the number of bits available for host IDs. So, $2^8 - 2 = 256 - 2 = 254$ hosts per network. Since we have 74 hosts and one router, that is a total of 75 host IDs. We have plenty of room for growth, and the scheme is simple.

The first address on our network starts at 204.74.9.1 (remember 0 is “this network”) and continues to 204.74.9.254 with a subnet mask of 255.255.255.0. Table 3.7 shows an example of the network portion of our address, 204.74.9, and the host portion in the last octet, from 1 to 254.

Table 3.7 Breakdown of the Mask for IP Addresses Using a Standard Subnet Mask

Source IP Address 204.74.9.21 Subnet Mask 255.255.255.0	Destination IP Address 204.74.9.209 Subnet Mask 255.255.255.0
--- Network ID --- . --- Host ---	--- Network ID --- . --- Host ---
11001100.01001010.00001001.00010000	11001100.01001010.00001001.11010001
11111111.11111111.11111111.00000000	11111111.11111111.11111111.00000000
11001100.01001010.00001001.00000000	11001100.01001010.00001001.00000000

Notice how the results of the subnet mask are equal? Of course, this is a simple example, and we can see just by the address and subnet mask that they are both on the same network.

Now we move into the new building where everyone gets his or her very own office. The office has three stories, so we need to break up our simple network into three segments to route between floors. We must use the same IP address block provided. One

option is to borrow bits from the host IDs and create more subnetworks. The number of subnets is determined by the value of the bits that we borrow from the host IDs. In the example in Table 3.8, we have used 2 bits, shown in bold. The last octet of the subnet mask is now 11000000, or 192. The number of hosts per network is $2^6 - 2$, or 62. That should be more than sufficient, so our limitation is the number of networks.

Table 3.8 Breakdown of the Mask for IP Addresses Using a Custom Subnet Mask

Source IP Address 204.74.9.21 Subnet Mask 255.255.255.192	Destination IP Address 204.74.9.209 Subnet Mask 255.255.255.192
——— Network ID —— . —Host —	——— Network ID —— . —Host —
11001100.01001010.00001001. 00010000	11001100.01001010.00001001. 11010001
11111111.11111111.11111111. 11000000	11111111.11111111.11111111. 11000000
11001100.01001010.00001001.00000000	11001100.01001010.00001001.11000000

To determine the number of networks we have, we take the bits 11 and use the formula $2^n - 2^2$, which is 4, so we can have up to four networks. We can create a list of the networks, convert them to decimal, and get the hosts for each network, as shown in Table 3.9. Remember that the first and last hosts for each network are not assignable.

Table 3.9 Determining the Address Blocks

Subnet	Range	Hosts
1	00000001 to 00111110	204.74.9.1 to 204.74.9.62 (204.74.9.0/26)
2	01000001 to 01111110	204.74.9.65 to 204.74.9.126 (204.74.9.64/26)
3	10000001 to 10111110	204.74.9.129 to 204.74.9.190 (204.74.9.128/26)
4	11000001 to 11111110	204.74.9.193 to 204.74.9.254 (204.74.9.192/26)

Each network has 62 hosts, and there are 4 networks, so we still have 248 hosts to grow into.

We could expand this example by adding satellite offices. Without redesigning the entire subnet, we could use one of the networks that was not used in the example and subnet it further. This is called *variable-length subnetting*. One of the networks would be broken down into two smaller networks with 30 hosts by borrowing another bit. The networks would have the notation 204.74.9.0/27 and 204.74.9.33/27. The hosts for 204.74.9.0/27 are 204.74.9.1 to 204.74.9.30, and the hosts for 204.74.9.33/27 are 204.74.9.34 to 204.74.9.63.



TEST DAY TIP

If you want to use the first and last networks in this scenario, you must use Classless Inter-Domain Routing (CIDR) notation and use routing services that support CIDR. In traditional subnetting, the first network ID is all zeros, so it is “this network,” and the last network ID is all ones, which signifies the broadcast for that network.

Classless Inter-Domain Routing (CIDR)

You should see now that there are limits to the size of the network you can implement using classful IP address assignment. It has become necessary to provide more options to create larger segments to reduce the size of routing tables and overcome the depleted public IP address pool. The solution is known as Classless Inter-Domain Routing (CIDR). CIDR uses a binary format to provide the definition of network addresses.

Use the matrix in Table 3.10 to quickly identify routing and subnet information based on your requirements for the number of hosts and networks. The column of binary masks should help you calculate the networks for each subnet, and the table shows how the classful addressing scheme relates to the CIDR notation.

Table 3.10 Quick Matrix for Determining Routing and Subnet Information

Required Networks	CIDR	Binary Mask	Hosts per Subnet (2^n-2)	Subnet Mask
256 Class B	/8	11111111.00000000.00000000.00000000	16,777,212	255.0.0.0
128 Class B	/9	11111111.10000000.00000000.00000000	8,388,606	255.128.0.0
64 Class B	/10	11111111.11000000.00000000.00000000	4,194,302	255.192.0.0
32 Class B	/11	11111111.11100000.00000000.00000000	2,097,150	255.224.0.0
16 Class B	/12	11111111.11110000.00000000.00000000	1,048,574	255.240.0.0
8 Class B	/13	11111111.11111000.00000000.00000000	524,286	255.248.0.0
4 Class B	/14	11111111.11111100.00000000.00000000	262,142	255.252.0.0
2 Class B	/15	11111111.11111110.00000000.00000000	131,070	255.254.0.0
1 Class B	/16	11111111.11111111.00000000.00000000	65,534	255.255.0.0
256 Class C	/16	11111111.11111111.00000000.00000000	65,534	255.255.0.0
128 Class C	/17	11111111.11111111.10000000.00000000	32,766	255.255.128.0

Continued

Table 3.10 Quick Matrix for Determining Routing and Subnet Information

Required Networks	CIDR	Binary Mask	Hosts per Subnet (2^n-2)	Subnet Mask
64 Class C	/18	11111111.11111111.11000000.00000000	16,382	255.255.192.0
32 Class C	/19	11111111.11111111.11100000.00000000	8190	255.255.224.0
16 Class C	/20	11111111.11111111.11110000.00000000	4094	255.255.240.0
8 Class C	/21	11111111.11111111.11111000.00000000	2046	255.255.248.0
4 Class C	/22	11111111.11111111.11111100.00000000	1022	255.255.252.0
2 Class C	/23	11111111.11111111.11111110.00000000	510	255.255.254.0
1 Class C	/24	11111111.11111111.11111111.00000000	254	255.255.255.0
1/2 Class C	/25	11111111.11111111.11111111.10000000	126	255.255.255.128
1/4 Class C	/26	11111111.11111111.11111111.11000000	62	255.255.255.192
1/8 Class C	/27	11111111.11111111.11111111.11100000	30	255.255.255.224

EXAM
70-293
OBJECTIVE
2.6

Troubleshooting IP Addressing

The flexibility of TCP/IP also contributes to the complexity of troubleshooting addresses and connections. There are several tools that can help isolate and identify issues with addressing, but it is also imperative that you understand IP addressing rules and subnetting. The **ipconfig**, **ping**, and **tracert** commands are the most useful tools for identifying addressing problems with client configurations and connections to other hosts on the Internet.

EXAM
70-293
OBJECTIVE
2.6.1

Client Configuration Issues

Some of the issues that occur with manual configuration of IP addresses include duplicate addresses, invalid subnet masks, invalid default gateways, and invalid or missing host name resolution settings (such as DNS and WINS). To help identify the problem, start by typing **ipconfig /all** at a command prompt. Verify the information that is output by the command is correct, and then continue by using **ping** to help isolate the problem.

1. Ping the loopback address (127.0.0.1) to verify that the TCP/IP protocol stack is configured correctly on the local computer.
2. Ping the external IP address of the local computer to ensure the host is on the network and using a valid IP address; that is, there are no address conflicts.
3. Ping the IP address of the default gateway to verify that the default gateway is accessible and your local network configuration contains the correct subnet mask.
4. Ping the IP address of a remote host to verify that you can transmit data over the default gateway.

If you are not able to get traffic through to a site, but you are making it through the default gateway, you should use **tracert** to identify the break in the route to the destination.

EXAM
70-293
OBJECTIVE
2.6.2

DHCP Issues

DHCP is an easy way to manage IP addressing schemes for larger networks. DHCP makes it possible to boot a machine and access the network without configuring any protocol information. This eliminates many of the manual configuration issues, such as using the wrong subnet mask, duplicate IP addresses, and limited or no host name resolution. Some of the items to consider when you implement and use DHCP are lease time, number of hosts in a scope, network traffic, scope options, and topology.

When a machine acquires an IP address from a DHCP server, it acquires a *lease*. The request for the lease is a message called a DHCPREQUEST, which is broadcast by the DHCP client looking for DHCP OFFERS of a lease from a DHCP server. The *lease duration* for a DHCP address is specified in the scope set on the server and defaults to eight days. At 50 percent of the lease duration, the DHCP client sends a directed request to the DHCP server that issued the lease and requests a renewal of the lease. If no DHCPACK (acknowledgment) is received from the server, the DHCP client waits until 87.5 percent of the lease time, and then makes a final request to renew the IP address. If no DHCPACK is received at this point, the client waits until the lease is expired and starts the process over. If a DHCP client is unable to receive an IP address lease, it will use an alternate configuration, if one is specified. If there is no alternate configuration, the client will use APIPA to start the TCP/IP services and assign itself an address from the APIPA pool (169.254.0.0/16).

To determine the appropriate lease time for your network, you should consider the following:

- **Number of hosts** If the number of hosts is close to the number of total IP addresses in your DHCP server's scope, the lease should be shorter—about three days. If there are a great deal more IP addresses than hosts, a longer lease can be assigned.
- **Mobile users** If you have a small number of mobile users and the client machines do not frequently move from one network to the other, a longer lease duration is recommended. Conversely, if you have more mobile users, a shorter lease will be preferred, so that the IP addresses will be released sooner and returned to the available pool of addresses.
- **Unlimited** It is possible to set the lease duration to unlimited, but it presents a challenge if you wish to change the DHCP settings, since this setting requires the client to initiate the DHCPREQUEST.

Because they are broadcast, the DHCPREQUEST messages do not cross router boundaries, unless the router is capable of forwarding DHCP broadcast messages, in com-

pliance with RFC 2131. You can also configure a DHCP relay to forward the requests to a DHCP server.

Using DHCP can reduce IP address conflicts by preventing the need for static IP address. It also can eliminate invalid subnet masks, since they are also assigned by the DHCP server. Another advantage is the use of scope properties. By assigning scope properties, you can define default gateways, DNS servers, WINS servers, and the type of name resolution that is preferred. By managing name resolution settings, you can help eliminate broadcast traffic.

Transitioning to IPv6

IPv6, defined in RFC 2460, is now production-ready to use on most operating system platforms. At this point, it is still early in the transition from IPv4. The change to IPv6 will take some time, but with each day, it becomes more necessary due to the growing shortage of IPv4 addresses. Although the larger address space is the most immediate need, IPv6 offers other advantages over IPv4, including the following:

- Better security (built in support for IPSec)
- Support for both stateful and stateless address configuration
- An efficient hierarchical routing infrastructure
- A new header format that provides lower overhead
- Neighbor Discovery (ND) for managing nodes on the same link, replacing ARP, ICMPv4 router discovery, and ICMPv4 redirect messages
- Virtually unlimited extension headers (in comparison to IPv4's limit of 40 bytes)
- Quality of service (QoS) related header fields

The utilities and concepts associated with IPv6 are similar to IPv4, but not identical. In the following sections, we'll take a look at how to install IPv6 and start to familiarize ourselves with the new utilities used to manage it.

IPv6 on Windows Server 2003 provides a new header format that is streamlined to minimize overhead and provide more efficient processing while crossing intermediate routers. All the option fields and any other fields in the header that are not required for routing are placed after the IPv6 header. The IPv6 header also added more QoS support by adding Flow Label fields that provide special handling for a series of packets that travel between a source and destination.

ND is a set of process and messages that are used in an IPv6 environment to identify relationships between neighboring nodes. This allows hosts to discover routers on the same segment, addresses, and address prefixes. With ND, hosts can also resolve neighboring nodes and determine when the MAC address of a neighbor changes (similar to ARP in IPv4). ND also provides the process for address autoconfiguration, also referred to as *stateless address configuration*. In the absence of a stateful address configuration server, such as a

DHCP version 6 (DHCPv6) protocol server, ND provides a complex process that allows each interface to use router advertisement messages to define an IPv6 address, and then subsequently ensure the uniqueness of the selected address. Currently, the standards for DHCPv6 and IPv6 stateful addressing are still under development, so neither feature is supported on Windows XP/Server 2003 products at this time.

The new routing structure provides a hierarchical addressing and routing structure that includes a global addressing scheme. Global addresses are the equivalent of public IPv4 addresses and are accessible over the Internet. The global addressing scheme defines new ways to summarize global addresses to facilitate smaller routing tables on the Internet backbone, thus improving the efficiency and performance on the Internet.



NOTE

For detailed information and links to white papers about IPv6 in Windows Server 2003, see Microsoft's IPv6 Web site at www.microsoft.com/windowsserver2003/technologies/ipv6/default.mspx.

IPv6 Utilities

The traditional IPv4 utilities are still very useful for IPv4, but new utilities and features have been added to accommodate IPv6 functionality. To gain access to the new tools or functionality, you need to install the TCP/IP version 6 protocol.

EXERCISE 3.04

INSTALLING TCP/IP VERSION 6

In the following exercise, you will learn how to install IPv6 on your Windows Server 2003 computer.




NOTE

You can also install or uninstall IPv6 from the command line, using the **netsh interface ipv6** context (discussed later in the "Netsh Commands" section).

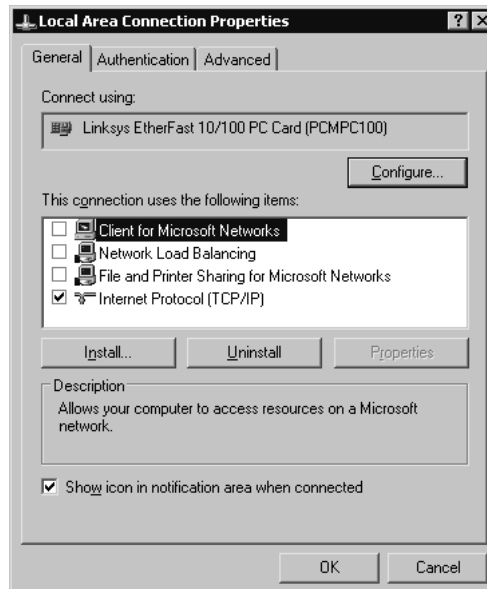
1. Open **Network Connections** and double-click the **Local Area Network** icon. You will see the **Local Area Connection Status** dialog box, as shown in Figure 3.11.

Figure 3.11 Local Area Connection Status




2. Click **Properties**.
3. In the **Local Area Network Connection Properties** dialog box, shown in Figure 3.12, click **Install**.

Figure 3.12 Local Area Connection Properties




4. In the **Select Network Component Type** dialog box, select **Protocol**, as shown in Figure 3.13, and click **Add**.

Figure 3.13 Select Network Component Type



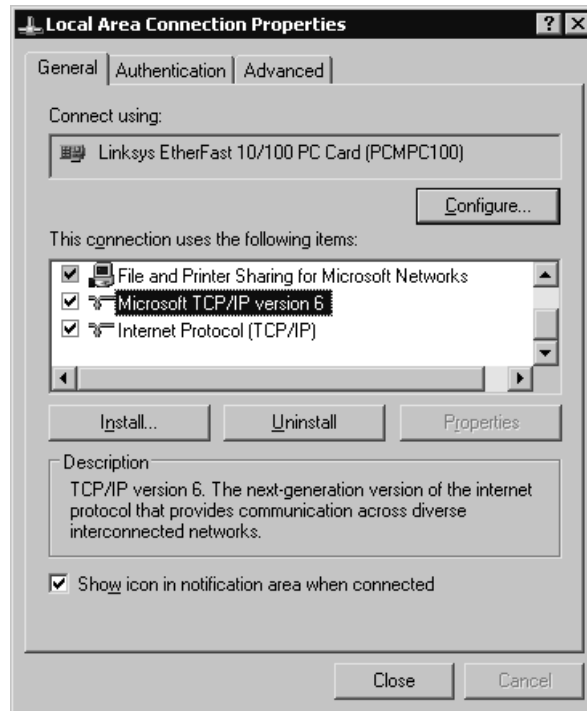
5. In the **Select Network Protocol** dialog box, select **Microsoft TCP/IP version 6**, as shown in Figure 3.14, and click **OK**.

Figure 3.14 Select Network Protocol



6. You should return to the **Local Area Connection Properties** dialog box and see that Microsoft TCP/IP version 6 is installed, as shown in Figure 3.15.

Figure 3.15 Local Area Connection Properties with TCP/IP Version 6 Installed

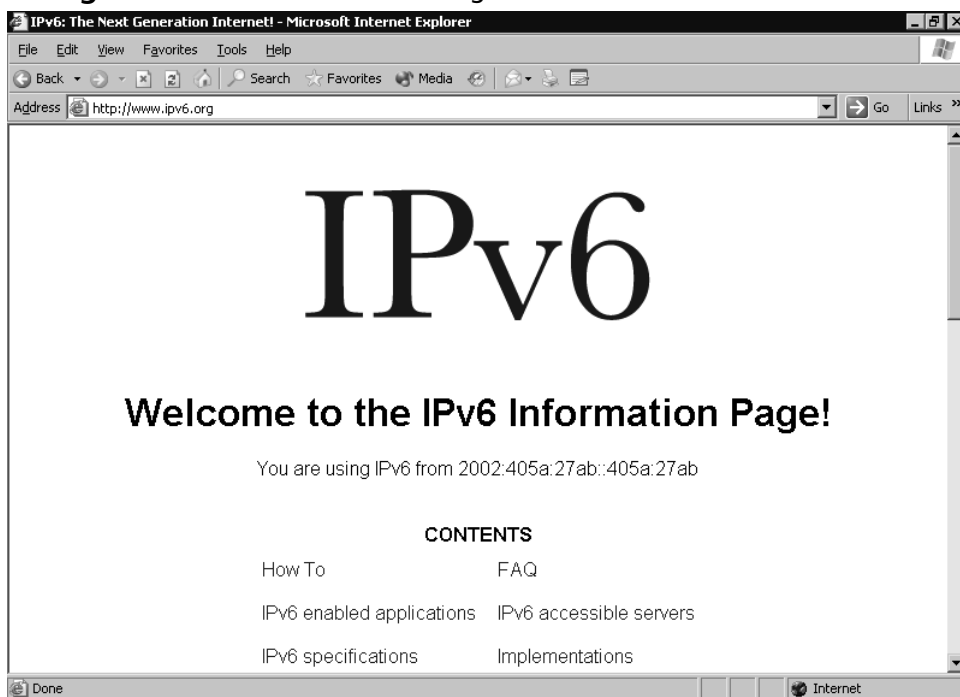


7. Click **Close**.
8. Test the TCP/IP version 6 installation by opening **Internet Explorer** and navigating to **www.ipv6.org**. You should see a line under the line "Welcome to the IPv6 Information Page!" that states, "You are using IPv6 from <your IPv6 address>," as shown in Figure 3.16. If you are behind a firewall or using 6to4 tunneling, you may not see the message that indicates you have an IPv6 address. If you are able to access the site described in step 9, then you are successfully using IPv6.

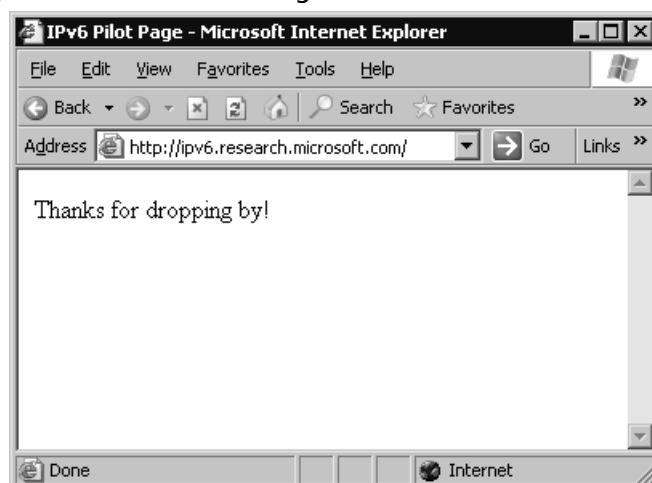


NOTE

You might need to reboot after installing IPv6.

Figure 3.16 Test the IPv6 Configuration

9. You can also navigate to an IPv6-only site from Microsoft Research. In **Internet Explorer**, navigate to <http://ipv6.research.microsoft.com>, as shown in Figure 3.17.

Figure 3.17 IPv6 Pilot Page at Microsoft Research

**NOTE**

You will not be able to browse IPv6-only Web sites with Microsoft Internet Explorer if you use a proxy server (unless the proxy server is IPv6-enabled).

Another way to test whether your IPv6 installation was successful is to run the **ipconfig** command. If IPv6 is installed, your IP address will be shown in IPv6 format, as shown in Figure 3.18.

Figure 3.18 ipconfig Results after Installing IPv6

```

C:\WINDOWS\system32\cmd.exe
Ok.
netsh interface ipv6>exit

C:\Documents and Settings\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter WAN:

    Connection-specific DNS Suffix  . : 
    IP Address . . . . . : 192.168.1.97
    Subnet Mask . . . . . : 255.255.255.0
    IP Address . . . . . : fe80::20c:29ff:fee:4176%4
    Default Gateway . . . . . : 192.168.1.7

Tunnel adapter Automatic Tunneling Pseudo-Interface:

    Connection-specific DNS Suffix  . : 
    IP Address . . . . . : fe80::5efe:192.168.1.97%2
    Default Gateway . . . . . : 

C:\Documents and Settings\Administrator>

```

Now that TCP/IP version 6 is installed, additional utilities are available with the IPv6 functionality. Other than the utilities to manage, monitor, and troubleshoot IPv6, only Telnet, FTP, and Internet Explorer actually use the IPv6 protocol stack.

netsh Commands

netsh is an interactive command-line utility that allows you to manage local or remote network configurations of active machines. netsh also supports scripting, so you can create batch configurations that run against the local machine or a specified host on the network. You can also use the Netsh utility to generate a configuration script to use as a backup configuration or as an aid to configure new machines in an identical fashion.

netsh works with the existing components installed with the operating system by using helper dynamic link libraries (DLLs). Each helper DLL contains the information necessary to execute the commands for the component to which it applies. The set of commands and features supported by the DLLs is called a *context*, and each context is unique to the networking component.

The IPv6 interface has its own context with commands to manage and display information pertaining to the routes, interfaces, addresses, and caches specific to IPv6. There are currently no graphical user interface (GUI) applications to configure IPv6, so netsh is nec-

essary for configuring IPv6 and its associated components. The component called 6to4 has a subcontext within the IPv6 context, for configuring and managing 6to4 routers and hosts. For more information about Netsh, see the Windows Help and Support Center topic titled “Netsh Overview.”

To put the netsh command into IPv6 context, type **netsh** at the command prompt, then at the **netsh>** prompt, type **interface ipv6**. Then you can use the IPv6 context commands, which include the following:

- **6to4** Changes to 6to4 context.
- **Add** Adds a configuration entry.
- **Delete** Deletes a configuration entry.
- **Dump** Shows a configuration script.
- **Install** Installs IPv6.
- **Isatap** Changes to isatap subcontext within IPv6 context.
- **Renew** Restarts IPv6 interfaces.
- **Reset** Resets IPv6 configuration.
- **Set** Sets configuration information.
- **Show** Displays information.
- **Uninstall** Uninstalls IPv6.

Ipsec6.exe

Ipsec6.exe is used to configure and implement IPSec security policies (SPs) and security associations (SAs) for IPv6. Using this utility, you can save and load security policies and security associations to a file that can be edited in a text editor. This can be a real timesaver when you implement IPSec for IPv6 on multiple machines. The command to save a configuration is **ipsec6 s *FilenameWithNoExtension***. The filename specified from the command line will be appended with the extension automatically. The extension .spd is added to security policy files, and the extension .sad is added to security association files. If you are executing this command for the first time, and there are no current policies and no current security associations, the files created can act as templates to help you get started.

Other ipsec6 commands are available to work with security policies and security associations:

- To load the configuration from these files, type **ipsec6 l *FilenameWithNoExtension***. The security policies will be loaded from *Filename.spd* and the security associations from *Filename.sad*.

- To delete security policies and security associations, type **ipsec6 d** [{sp | sa}] **[Index]** from a command line. Use the **sp** parameter with the *Index* of the policy you wish to delete, or the **sa** parameter to delete all of the security associations.
- To determine what the current security policies are, type **ipsec6 sp** **[Interface]** from the command line, where *Interface* is optional and applies to the security policies for the specified network interface.
- To view the current security associations, type **ipsec6 sa** from the command line. Note that the output from the commands to view the security policies and security associations is not formatted well for a command line, so you might prefer to save the configuration and view the files in Notepad.



TEST DAY TIP

According to Microsoft Help and Support Center documentation, the current version of IPSec for IPv6 is not recommended for use in a production environment, so you should not be concerned about anything more than being familiar with it for the exam.

IPv6 PING and Tracert Parameters

Use the following steps to use IPv6 PING to verify connectivity:

1. From a command prompt, type **netsh interface ipv6 show interface**.
2. Find the **Idx** value for **Local Area Connection**.
3. Type **netsh interface ipv6 show interface Idx**, where *Idx* is the number from the previous step. The Local Area Connection index number is usually **4**.
4. Right-click in the command window and select **Mark**. Then highlight the address. Once it is highlighted, right-click in the command prompt window. When you release the mouse button, the address will be copied to the Clipboard. Take note of your **Zone ID** for **Link**, which should match the **Idx** number in step 3.
5. Exit the **netsh** command. At a regular command prompt, type **ping**, and then right-click in the command prompt window and select **Paste**.
6. Without adding any spaces, add %<ZoneID>, where *ZoneID* is the number noted in step 4, so the command looks like this:


```
Ping fe80::204:5aff:fe08:fb4b%4
```
7. Press **Enter**. You should see four successful replies.

8. Continue by pinging another address on the same local network.
9. To test external hosts, ping the global address of another node.
10. To test name resolution with DNS or a hosts file, ping a node with **ping -6 Name**, where *Name* is the site name. The **-6** parameter tells PING to use IPv6 only.

You can use Tracert to trace the path taken by IPv6 data packets from this host to the destination host. From a command prompt, type **tracert IPv6Address%ZoneID**, where *IPv6* is a valid IPv6 address and *ZoneID* is the destination address. Alternatively, type **tracert -d -6 Hostname**, where *Hostname* is the name of the remote machine.



NOTE

Windows XP Professional includes three utilities not included with Windows Server 2003: `ipv6.exe`, `ping6.exe`, and `tracert6.exe`.

6to4 Tunneling

6to4 tunneling is used to encapsulate IPv6 data packets in IPv4 headers before they are transmitted to the destination host. 6to4 tunneling uses a 6to4 host and 6to4 routers to deliver the IPv6 data. It is an Internet standard, defined in RFC 3056, and is used for interoperability between IPv4 and IPv6 networks. 6to4 hosts and routers are defined as follows:

- **6to4 host** Any IPv6 host that is configured with at least one 6to4 address. 6to4 can be configured with the **netsh interface ipv6 6to4** commands. As you might have noticed when you ran the **show interface** command, by default, your IPv6-enabled host will have a 6to4 pseudo-interface, as well as an automatic tunneling pseudo-interface.
- **6to4 router** Uses IPv4 and IPv6 to forward 6to4 traffic to the destination 6to4 hosts. It is also possible to implement a 6to4 relay router to forward 6to4 router traffic on the IPv6 Internet.

With 6to4 tunneling, it is not necessary for IPv6 hosts (such as the computer on which you installed IPv6 in Exercise 3.4) to get an IPv6 global address prefix from their ISPs. The host can create a 6to4 address automatically.

IPv6 Helper Service

The IPv6 Helper service is responsible for automatically configuring itself with the appropriate 6to4 addresses, but it uses a specific 6to4 router on the Internet. You can test functionality with the **ping -6** command.

The 6bone

The 6bone is a dedicated IPv6 network that exists on the Internet. It began as a virtual network using IPv6 over IPv4 encapsulation. It contains links to many sites and includes a great deal of IPv6 data, testing plans, news, current events, and implementation instructions. It will be a valuable resource for managing IPv6 on your network. For more information about the 6bone, see www.6bone.net. For instructions on how to connect to the 6bone, see www.opus1.com/ipv6/whatisthe6bone.html.

Teredo (IPv6 with NAT)

Teredo is the name for IPv4 network address translator (NAT) traversal for IPv6. It provides an IPv6/IPv4 translation over NAT and address assignment. Teredo also provides the mechanism for host-to-host automatic tunneling for unicast IPv6 connectivity when IPv6/IPv4 hosts are located behind one or more NAT servers.

Currently, to provide IPv6 connectivity over the Internet, you must have a 6to4 router with a public IPv4 address, which is not always feasible. Teredo provides a mechanism for IPv6 traffic to traverse NAT and access the Internet using IPv6. Basically, IPv6 packets are sent as IPv4-based UDP messages, and this allows the IPv6 packets to pass through the IPv4 NAT server. For more information about Teredo, see the Teredo Overview document located at www.microsoft.com/windowsxp/pro/techinfo/administration/p2p/overview.asp.

EXAM 70-293
OBJECTIVE
2
2.1
2.1.2
2.2

Planning the Network Topology

The next phase in planning your TCP/IP infrastructure is planning the IP routing solution to manage the traffic on your network. This will depend on the physical location of your equipment and users, as well as on how you want to distribute the addresses. When you implement your strategy, you will also need to determine how the hosts on your network will resolve host names and implement the necessary services to provide that functionality. You will need to identify where the services such as DHCP, WINS, DNS, and so on must exist in your network to function properly and reduce the network bandwidth utilization.

Analyzing Hardware Requirements

Before you implement your network topology, you should identify the hardware needs. For each physical location, you will need to provide some sort of routing. You might need to implement a WAN solution using a T1 line, which also requires special hardware. You will need DHCP servers at each location or a DHCP relay agent. You will need to provide some form of name resolution, most likely DNS and possibly WINS. Depending on traffic and if you have a large number of users, you may decide to install switches to help manage network traffic.

For a DHCP server, the two major factors that affect performance are the amount of physical random access memory (RAM) and the speed of the disk input/output (I/O). You should always provide the largest amount of RAM possible and the fastest disk I/O for the best performance on a DHCP server. The same rules apply for WINS and DNS servers, although DNS is more dependent on network bandwidth. In any case, frequent zone updates require more RAM for better performance.

If you are using Active Directory (AD) DNS, there are other considerations related to AD, such as:

- Increased network utilization due to dynamic DNS updates related to DHCP integration and WINS reverse lookups
- Increased RAM requirements due the increased data volume

EXAM
70-293
OBJECTIVE
2.2.1

Planning the Placement of Physical Resources

The quantity of data and the type of network traffic will affect the location of IP resource servers in your enterprise. If the WAN link is slow, you might want to place DNS caching servers at each location to reduce WAN traffic related to DNS resolution. You might also consider providing a DNS server at each location to provide redundancy. In addition, by creating an AD integrated primary zone, you will allow clients to update their resource records locally. Defining which DNS servers can act as forwarders and perform iterative queries will help manage the Internet traffic.

You should also provide a DHCP server at each location. When you have multiple DHCP servers on your network, use the 80/20 rule to balance the load on the subnet: 80 percent of the scope will be on the primary server, with 20 percent on the other server. The DHCP server must have an interface on each network for which it has a scope defined, or you must locate a DHCP relay server on the same subnet as the DHCP clients.

If you implement WINS, you will need to examine the quantity of data replicated between WINS servers and the cost of WINS reverse lookups from DNS servers. You should minimize the number of WINS servers you implement in order to minimize the impact of WINS replication traffic on your network.


Use the Help and Support Center on Windows Server 2003 to see examples of performance statistics in a high traffic environment to help you gauge your enterprise needs.

EXAM
70-293
OBJECTIVE
2
2.1
2.1.1

Planning Network Traffic Management

After you decide where to place your physical equipment, users will begin accessing the services supplied by DHCP, DNS, and WINS. Other traffic comes from accessing the Internet, file sharing, and the many other network resources that will be used. You can estimate the amount of traffic at peak times by using some of the utilities provided with the operating system. The tools can be used to create baselines, identify the peak network usage areas, and identify the traffic sources.

You will also need to monitor network traffic and analyze the usage. You might be able to identify illicit network access from external sites, find Trojan horse viruses that generate broadcast storms, or just discover who is actually hogging all that Internet bandwidth. You can also determine whether your server-to-server traffic is managed well, or if it is necessary to modify the physical location of equipment.



Monitoring Network Traffic and Network Devices

Every network administrator should be familiar with two key utilities:

- **Network Monitor** Allows you to capture data, identify the source, and analyze the content and format of the message.
- **System Monitor** Allows you to monitor other resources and determine the performance of those resources.

Using Network Monitor

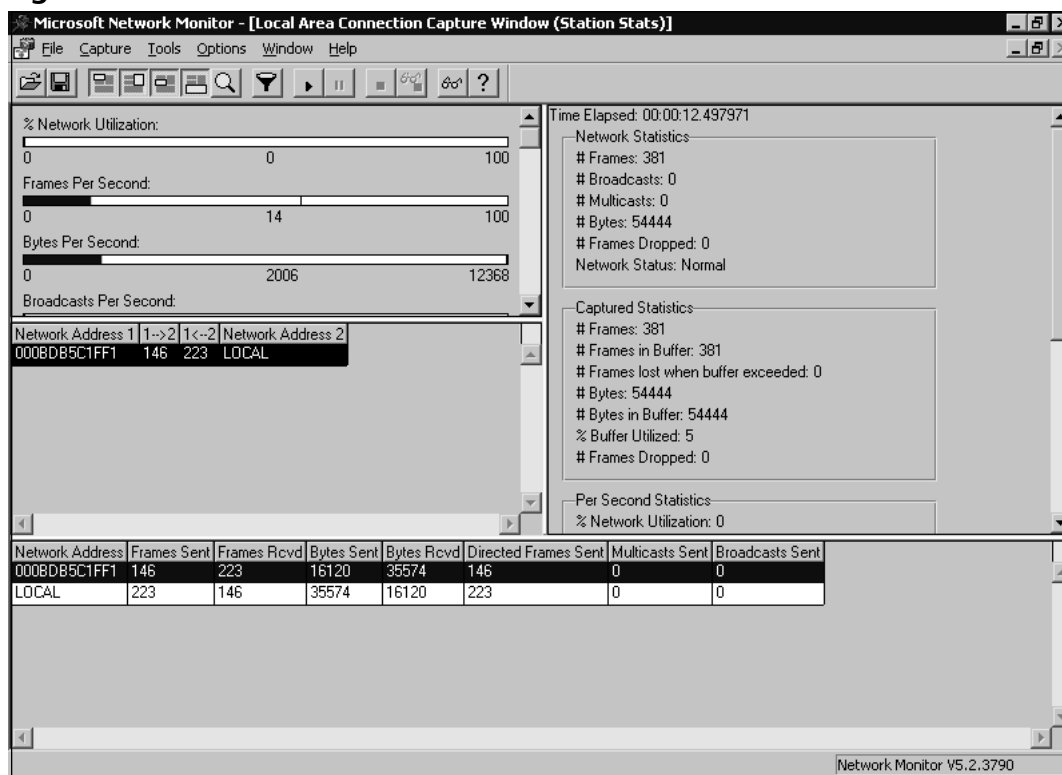
There are two versions of Network Monitor: one is part of the Windows Server 2003 operating system, and the other is part of Microsoft Systems Management Server (SMS). The version that ships with Windows Server 2003 can monitor only traffic inbound and outbound to the machine on which the utility is being run. The SMS version can monitor most network traffic from any machine to any other machine on the network, by placing the network card on the machine where it is running in promiscuous mode to capture all traffic.

Network Monitor is not installed by default. You can install it by following these steps:

1. From Control Panel, select **Add/Remove Programs**.
2. Click **Add/Remove Windows Components**.
3. Click **Management and Monitoring Tools**.
4. Click **Details**.
5. Click the check box next to **Network Monitor Tools**.
6. Click **OK**.
7. Click **Finish**.

After Network Monitor is installed, you can use the interface to monitor traffic, as shown Figure 3.19. When you want to view the results, you can view each frame of captured data. You can save the trace to a file, or you can start the trace over. You could then use the traces to find and filter traffic in order to analyze the data. You can also capture fragments into files for later analysis. You can even see some of the unencrypted data being transmitted on your network.

Figure 3.19 Network Monitor

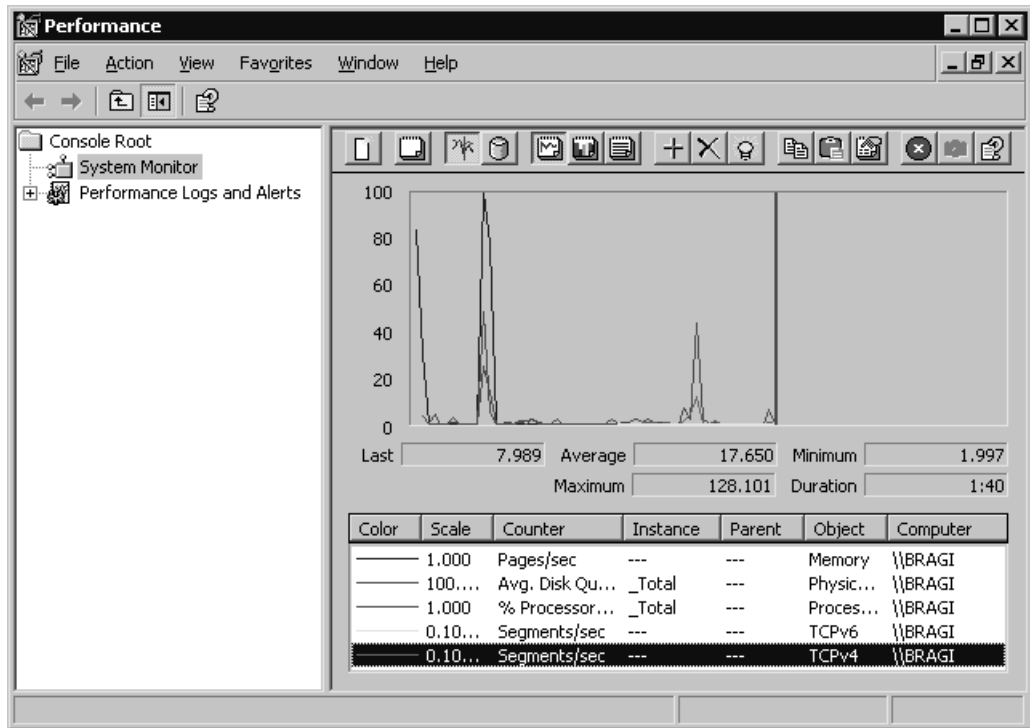


Network Monitor should be run during low-usage times or for short intervals to minimize the impact on performance of capturing all that data on your machine. It is also useful to identify the type of traffic you are concerned with and use the filters to capture only the data you need.

Using System Monitor

System Monitor is a Microsoft Management Console (MMC) snap-in tool that allows you to use counters to monitor the performance of hardware, applications, and operating system components on Windows Server 2003 machines.

A *counter* is basically a hook into a driver or application component that allows System Monitor to gather statistics. System Monitor can capture these statistics and display them in a graph, as shown in Figure 3.20, or in a report. It can also send administrative alerts when specified conditions are met, and even launch an application to allow you to correct the situation or send an e-mail or a page to an administrator. You can save the logs to different file formats to allow you to analyze them in other applications or tools.

Figure 3.20 System Monitor


NOTE

Windows Server 2003 includes command-line tools to help control the scheduling of performance counter and event trace logs. System Monitor is no longer required to gather performance data from remote computers (although it can still be used for that purpose). `Typeperf` allows you to write performance counter data directly to the command window.

System Monitor also allows you to view more than one log file at the same time, so that you can compare baseline logs with the current data. The Performance Logs and Alerts service can gather data and store it in a Microsoft SQL Server database that can be viewed by System Monitor. You can also save portions of log files or SQL Server data to a new file. This can help save space, simplify comparisons of data, and reduce analysis time.

Determining Bandwidth Requirements

When you have captured performance statistics and viewed the network traffic during various times of the day, you can identify the different sources of traffic on your network. You will need to analyze how name resolution occurs, where the requests for name resolution initiate, and the server-to-server traffic when replicating the information.

You will need to identify the following:

- Any slow connections and the quantity of data transmitted over those connections. This will help you to identify how often servers transmit replicated data to other servers.
- The cost of one client obtaining information from these servers. You can then use that information to calculate the cost of many users.
- Broadcast traffic, so that you can isolate that to certain networks. You will be able to identify areas where clients communicate heavily with other clients, such as file servers, and locate those resources on the same segment as the heavy users.

Optimizing Network Performance

TCP traffic uses a *sliding window* method of transmitting data. As data is successfully transmitted to the destination, the window slides over the remaining data and transmits the next packets of data. Window size is basically the maximum number of packets that can be sent without waiting for positive acknowledgment. If you transmit large amounts of TCP data, then larger TCP windows will improve TCP/IP performance. The maximum window size is limited to 64 kilobytes by default and is determined by the windows size setting of the destination host machine. It is possible to increase the size of the TCP window dynamically on Windows Server 2003 to accommodate this by enabling large TCP window support. Client computers can be set to request large windows by editing their Registries. These are then called *TCP1323Opts-enabled* computers. The window size is negotiated during the TCP three-way handshake process. TCP1323 is a TCP extension defined in RFC 1323.

With Windows Server 2003, it is possible to disable NetBIOS encapsulation over TCP/IP (disable NetBT). This can significantly reduce the overhead of data transfer and eliminate the need for WINS and any other NetBIOS name resolution. It will also reduce the browser master traffic. The drawback to disabling NetBIOS encapsulation is that you can no longer browse network resources. In addition, some applications depend on NetBIOS and will not work without it. If you are using NetBIOS name resolution, you should have WINS servers to allow for directed send requests for name resolution, rather than broadcasting for that information. WINS servers share data with each other at a regular intervals. You might wish to reduce that traffic by modifying the replication intervals to increase the time between synchronizations. You should minimize the number of WINS servers used on your network. It is not necessary to have a WINS server on every LAN. The more WINS servers you implement, the more network traffic is generated by WINS database replication.

The placement of other servers that provide network services is also important. DHCP servers must have an interface on the same segment as the clients that will use the DHCP server, or you must provide a means for DHCP requests to cross routers (such as a DHCP relay or using routers that allow DHCP and BOOTP requests). Place DNS servers on each LAN to minimize the amount of traffic generated when performing host name resolution. You can also designate which DNS servers can act as forwarders to control which machines can perform iterative DNS queries over the Internet.

Summary of Exam Objectives

In this chapter, we examined the factors associated with how to identify network protocols that are best suited to your needs. After we identified the different factors, we evaluated the advantages of using the TCP/IP protocol suite over other protocols, as well as how the Windows Server 2003 platform allows the flexibility to use multiple protocols to communicate on your network, and when it might be necessary to do so. We reviewed how to configure TCP/IP manually and summarized some of the new features and enhancements of the Windows Server 2003 networking components.

We reviewed how the TCP/IP network model (actually the DoD model) maps to the OSI reference model and leverages each layer of the TCP/IP model to provide a robust and stable platform for network communications. We took a more in-depth look at the new TCP/IP enhancements in Windows Server 2003, including many of the improvements that will reduce administrative workload such as the new alternate configuration feature for TCP/IP. You also discovered that TCP/IP can now determine the routing metric for the default gateway dynamically, which will help improve the performance of TCP/IP connections to other subnets.

We defined the criteria for addressing TCP/IP networks and how subnetting works. You learned how to subnet networks and convert binary numbers to decimal and back to help implement the addressing schemes you design. We reviewed how to troubleshoot TCP/IP connections and the issues with manual configuration of clients versus automatic configurations using DHCP. We identified your options for DHCP lease duration and how to decide how the duration is set.

After explaining how to install IPv6, we provided you with an overview of the utilities and software that uses IPv6, and how to configure and troubleshoot IPv6 using netsh, ipsec6, ping, and tracert commands. We also looked at the 6to4 router and hosts and how they can assist you in making the transition from IPv4 to IPv6 by encapsulation of IPv6 data in IPv4 packets.

Finally, we examined the tools that are included in Windows Server 2003 to help you monitor, maintain, and plan your network infrastructure. Using those tools, you can identify areas for performance tuning and improving resource availability to minimize network bandwidth utilization and improve network performance.

Exam Objectives Fast Track

Understanding Windows 2003 Server Network Protocols

- ☑ Windows Server 2003 supports multiple protocols at the same time using NDIS, allowing better integration and flexibility for network operations.
- ☑ Considerations for choosing the best protocol also help define why TCP/IP is best suited to enterprise environments.

- ☑ TCP/IP is a suite of protocols that includes applications and network protocols that can be used to access and share information with the world or to use the Internet as a means for implementing WANs.
- ☑ There are many enhancements to the TCP/IP protocol suite included in Windows Server 2003 that will improve your overall experience and reduce network load.

Planning an IP Addressing Strategy

- ☑ The number of hosts and the number of networks required define the basis for your addressing strategy.
- ☑ Planning for growth is critical for your networking address structure, but it is also beneficial to implement the addressing scheme in an efficient manner.
- ☑ CIDR can reduce the number of static routes and simplify your network implementation.

Planning the Network Topology

- ☑ Servers should be placed close to the clients that will be using the resources provided.
- ☑ DHCP provides automatic addressing and other IP address configuration settings to network machines, which prevents errors typically encountered when manually configuring IP address settings.
- ☑ DHCP servers must have an interface on the same segment as the DHCP clients, or you must implement a DHCP relay.
- ☑ DNS is used for host name resolution.
- ☑ You should have one DNS server for each LAN and define which DNS servers are forwarders and perform iterative queries over the Internet.
- ☑ WINS is used for NetBIOS name resolution, and it is not necessary if you do not use NetBIOS to access network resources and have only Windows 2000/XP/2003 machines on the network.
- ☑ You should minimize the number of WINS servers on your network. WINS replication uses a lot of network bandwidth.

Planning Network Traffic Management

- ☑ Network Monitor can be used to examine data transmissions sent over the network. It provides a means for tracking down network issues.
- ☑ System Monitor is a local or remote performance utility that you can use to identify bottlenecks and issue alerts when undesirable situations occur.
- ☑ Bandwidth requirements vary, but by using the tools provided, you can allocate resources appropriately and optimize your system's performance by reducing and perfecting data delivery.

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: Will I need to learn how to subnet networks as a LAN administrator?

A: Yes, the ability to design and implement and support networks using TCP/IP depends on your ability to understand IP addressing practices. It is also important to understand subnetting for troubleshooting problems and expanding your network.

Q: Is it necessary to memorize all the options for Netsh to manage my network effectively?

A: You should be familiar with the various functions provided by Netsh and understand its importance in configuring IPv6 and other networking components. You may find useful functionality that can simplify repetitive tasks, since netsh is a command-line tool and provides you with a means to automate tasks. You can even use it to back up configurations for services such as DHCP and DNS to simplify building similar machines on your network.

Q: Is everything I need to know about TCP/IP to do my job in this chapter?

A: No, volumes of data exist on TCP/IP, including many valuable Internet resources such as IPv6.org and IETF.org. Every day, new information about the development of TCP/IP protocols is available. In addition, there are books dedicated solely to TCP/IP and still others that talk about security on networks that use TCP/IP.

- Q:** Do I need to know all the port numbers for the different protocols to manage my network?
- A:** You should be familiar with the common port numbers, such as those for FTP, HTTP, and SMTP, but it is not necessary to memorize every single one. Understanding how to determine which port does what can help you identify which services are in use on machine, as well as provide better security for your network. You can learn to use and identify different ports to do other tasks, such as testing SMTP on port 25 using telnet.exe (the Telnet port defaults to 22).
- Q:** Can I use IPv6 exclusively on my network?
- A:** Yes, however, due to the limited application support, it would be very difficult at this point to eliminate IPv4 and still function efficiently. For instance, there is no IPv6 implementation of DHCPv6, so it is difficult to manage configuration settings for networks that have many clients. Other common protocols such as SMTP, POP, and NNTP do not currently support IPv6. In addition, the majority of Internet resources are using IPv4, and you would require some implementation of IPv4 on your network to access those resources.
- Q:** Can I use CIDR notation on any router?
- A:** No, only certain versions of the routing protocols RIPv2 and OSPF support CIDR notation. Routers using RIPv1 do not support CIDR notation, and thus require the full routing information to be provided. This could present issues if you are using CIDR notation for routers that will interface with RIPv1 routers and router discovery. Most hardware routers can use CIDR notation to define routes. CIDR notation can help reduce the number of route entries that must be added to the routing table.
- Q:** Do I need a public class IP address block for my network if I have 200 hosts that need Internet access?
- A:** No, it would be very costly and difficult to obtain an entire block of class C addresses. You should implement a firewall. Then you will be provided with either a single IP address or a small subnet of six or fewer public addresses that will provide the external interface to the world. Instead, you should use a private IP addressing scheme internally to allow for outbound traffic to the Internet via NAT. Public addresses would be necessary for Web servers, VPN over the Internet, and other interfaces that need to be accessible over the Internet. E-mail servers must have a public IP address to allow delivery of Internet messages. You may be hosting a DNS server that provides host name resolution for your public Web servers. The DNS server would require a public interface to allow other clients to perform lookups, to update and receive updates for a zone, and so on.

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 Windows 2003 Server Network Protocols

1. You are implementing a network that will include UNIX workstations that will share files and information with the Windows users. What protocols will you need to implement to provide integration with UNIX machines?
 - A. IPX/SPX
 - B. NetBEUI
 - C. TCP/IP
 - D. NetBIOS over TCP/IP
2. You purchased a new desktop computer running Windows XP for your small office and a server running Windows Server 2003. Your old desktop is running Windows 95. It has a network adapter and can access files on another Windows 95 machine. The Windows XP machine has not arrived, but you want to back up the data from the Windows 95 computer to the Windows Server 2003 machine. However, from the Windows Server 2003 computer, you are unable to see the shares on the Windows 95 computer. What should you do to allow the Windows Server 2003 machine to access the Windows 95 machine?
 - A. Install NetBEUI on Windows Server 2003 computer.
 - B. Install NWLink on the Windows 95 client.
 - C. Install TCP/IP on the Windows 95 client.
 - D. Ensure the server has a valid IP address and implement a DHCP server on the Windows Server 2003 machine with a valid scope.

Planning an IP Addressing Strategy

3. You are implementing a test lab that contains three Windows Server 2003 machines, twenty Windows XP Professional machines, and two IP-based printers. You have been given the network address of 155.1.50.0 and a subnet mask of 255.255.255.224. What is the CIDR notation for your subnet?

- A. 155.1.50.0/27
 - B. 155.1.50.0/5
 - C. 155.1.50.0/24
 - D. 155.1.50.0/3
4. You are given a task to create eight subnets on your LAN, and you have been assigned the address space 172.16.128.0/23. How many hosts will you have and what is the CIDR notation for the new subnet's address space?
- A. 2032 hosts on 172.16.128.0/24
 - B. 240 hosts on 172.16.128.0/27
 - C. 496 hosts on 172.16.128.0/26
 - D. 48 hosts on 172.16.128.0/29
5. Which of the following addresses is suitable for dividing into at least nine subnets, each with the ability to support 200 hosts per network?
- A. 10.1.1.0/24
 - B. 10.1.1.0/20
 - C. 10.1.1.0/19
 - D. 10.1.1.0/22
6. You are having trouble accessing Microsoft's Web site. When you ping `www.microsoft.com`, the request times out. How should you proceed in troubleshooting this problem?
- A. Ping the loopback adapter, the IP address of this machine, then the default gateway and determine if your connectivity is valid. If there are no issues, run `tracert` and identify where the communications stop.
 - B. Ping the default gateway, the IP address of a remote host other than Microsoft, such as Yahoo, then ping the IP address of this machine and then the loopback adapter.
 - C. Use Network Monitor to analyze the traffic to `www.microsoft.com`.
 - D. Use System Monitor to look at counters on the local machine to determine the error.
7. You implement a Windows Server 2003 machine that is functioning as a file server on your LAN. The server name is `FileServer01`. Users attempting to browse the shares on `\\FileServer01\` are unable to see any of the shares you created. What is likely the problem?

- A. You do not have DNS installed on the LAN.
 - B. DHCP is unavailable.
 - C. NetBIOS encapsulation is not enabled on the Windows Server 2003 machine.
 - D. FileServer01 FTP service is stopped.
8. A client computer configured as a DHCP client was unable to obtain an address from the DHCP server. Upon investigation, you discovered that the DHCP scope was not activated, so you activated it. The client computer has an APIPA address of 169.254.0.1. What actions are required for the client to obtain an IP address from the DHCP server?
- A. Run `ipconfig /all` from a command prompt.
 - B. Use Netsh to assign an address to the network adapter.
 - C. Log off Windows XP and log on again.
 - D. Take no action.

Planning the Network Topology

9. Your company is merging with another organization, and you have been tasked with merging the corporate networks. You have determined that the other company has between 50 and 125 hosts on 7 networks. Your company has 25 to 50 hosts on 12 networks. You want the integration to provide room for five percent growth over the next two years. Your routers do not support variable-length subnet masks. You decide to use the private address 192.168.0.0. What is the best subnet mask for your new corporate LAN?
- A. 255.255.0.0
 - B. 255.255.255.0
 - C. 255.255.255.192
 - D. 255.255.224.0
10. You want to simplify the configuration and management of TCP/IP clients on your network, which consists of 300 Windows XP Professional machines, 12 Windows Server 2003 machines, and 23 printers on four subnets. Which of the following solutions best suits your needs?

- A. Implement WINS using APIPA. Provide at least one DNS server for each WINS server.
 - B. Implement DHCP to provide assigned IP address leases and scope properties that contain the necessary host resolution methods, the IP address of the default gateway, and the DNS servers.
 - C. Implement AD integrated DNS and WINS and configure WINS to do reverse lookups.
 - D. Provide thorough documentation for each client to manually configure its IP address with a valid subnet mask and DNS server.
11. All of the clients on your network are configured to use DHCP for their TCP/IP configuration. You upgrade Internet access to use a T1 line that is connected to a different router than the current router that is being used by the Digital Subscriber Line (DSL) connection. What actions are required to allow the executive staff to access the Internet using the new default gateway, by configuring each executive's machine only one time, while not allowing the other company employees to use the T1?
- A. Create a logon script for the Executives Group that uses the route add -d command to add the new router information. Set the script to run every time members of the Executive Group log on.
 - B. Create a logon script for the Executives Group that uses the route add -p command to add the new router information. Set the script to run once the next time members of the Executive Group log on.
 - C. Create a new property for the router in the DHCP scope options. Set up reservations for each of the executive's machines.
 - D. Run the command route add with the information for the new router on each executive's machine.
12. You have integrated a smaller LAN into your network that contains a Novell NetWare server using IPX/SPX. You want to be able to access it from a Windows Server 2003 machine, so you install NWLink. You notice that after you installed NWLink, the Windows XP client machines that connect to Windows Server 2003 are taking longer to connect and read information. What can you do to ensure the best performance for the Windows XP clients?
- A. Install NWLink on the Windows XP machines.
 - B. Install the Novell NetWare Client on the Windows XP machines.
 - C. Move TCP/IP up in the binding order on the Windows Server 2003 machine.
 - D. Install the Novel NetWare Client on the Windows Server 2003 machine.

13. You are network administrator for a new company. Your LAN is connected to the Internet by a single T1 line. You obtain a single public IP address from your ISP. Your firewall services are outsourced to the ISP. The LAN includes five Windows XP Professional computers and one Windows Server 2003 computer named Server01. All Windows XP client computers are configured to use DHCP to obtain their IP configurations. Server01 is configured as a DHCP server and contains two network adapters. You connect one network adapter to the hardware for the ISP connection and connect the other network adapter to the LAN. You want client computers to access the Internet, including browsing the Web and file transfers via FTP. Which of the following configuration tasks must you complete?
- A. Install the DNS Server service.
 - B. Install WINS Services.
 - C. Install Routing and Remote Access Services (RRAS).
 - D. Assign the public IP address to the external adapter.

Planning Network Traffic Management

14. Users are complaining about slow network performance. Using Network Monitor, you have identified the source of the excessive traffic is inbound and outbound traffic from your DNS server. How would you identify the source of the excessive DNS traffic?
- A. Using the host IP addresses from Network Monitor, perform a tracert command to each host and determine the time it takes to get to each requested destination.
 - B. Use System Monitor to watch performance counters on the DNS server and identify the cause of the slow performance.
 - C. Use System Monitor to watch performance counters on the client machines to identify the machine that is using the DNS server heavily.
 - D. Ping the DNS server using the `-t` option from different host machines to identify the subnet that is causing the increase in network traffic.
15. You are using Network Monitor to analyze traffic on your Windows Server 2003 machine. You have a lot of data that has been captured, but you are looking for specific information. How do you accomplish this?
- A. Define a filter for the captured data.
 - B. Open the trace in Notepad and do a global search for the information you are seeking.
 - C. Export the data to a .cap file and view the reports in Excel.
 - D. Set up the counters for the appropriate data.

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

MCSE 70-293

Planning, Implementing, and Maintaining a Routing Strategy

Exam Objectives in this chapter:

- 2 Planning, Implementing, and Maintaining a Network Infrastructure
 - 2.1.2 Plan an IP routing solution.
- 3 Planning, Implementing, and Maintaining Routing and Remote Access
 - 3.1.1 Identify routing protocols to use in a specified environment.
 - 3.1.2 Plan routing for IP multicast traffic.
 - 3.1 Plan a routing strategy.
 - 5.3.1 Specify the required ports and protocols for specified services.
- 3.4 Troubleshoot TCP/IP routing. Tools might include the route, tracert, ping, pathping, and netsh commands and Network Monitor.
- 2.5.3 Diagnose and resolve issues related to client configuration.

Introduction

In the preceding chapter, you learned about the TCP/IP protocols and how to set up a TCP/IP infrastructure. One of the biggest advantages of TCP/IP as a network and transport protocol stack is its capability to route packets between different networks or subnets. Dealing with routing issues is an important part of the job of a Windows Server 2003 network administrator for a typical medium-to-large size network. In this chapter, we first review the basics of IP routing, including the role of routing tables, static and dynamic routing, and routing protocols such as Routing Information Protocol (RIP) and Open Shortest Path First (OSPF).

You'll learn to use the netsh commands related to routing, and then we'll show you how to evaluate routing options. This includes selecting the proper connectivity devices; we'll discuss hubs, bridges, switches (Layer 2, 3, and 4 varieties), and routers. We'll look at how you can use a Windows Server 2003 machine as a router, and how to configure the Routing and Remote Access Service (RRAS) to do so.

Next, we look at security considerations related to routing. We'll show you how to analyze requirements for routing components from a security-conscious point of view, and we'll discuss methods of simplifying the network topology to provide fewer attack points. This includes minimizing the number of network interfaces, the number of routes, and the number of routing protocols. We will also discuss router-to-router virtual private networks (VPNs), packet filtering, firewalls, and logging levels.

Finally, we cover how to troubleshoot IP routing issues. We'll identify troubleshooting tools and take a look at some common routing problems, including those related to interface configuration, RRAS configuration, routing protocols, TCP/IP configuration, and routing table configuration.

EXAM 70-293
OBJECTIVE
2
2.1.2
3

Understanding IP Routing

The basic concept of routing is that each packet on a network has a source address and a destination address. These two addresses are stored in the packet's header information. That means that any device on the network that receives this packet can inspect the header to find out where the packet came from and where it's going. If we provide our device with a little more information, such as details concerning the network's design and implementation, that device can also change the routing for the packet in an intelligent manner to help lower the total cost of the traffic.

So that we're all on the same page, we need to start by reviewing the basics of routing. Keep in mind as we go through the following material that it is mainly review and not intended as the final word on these topics.

Reviewing Routing Basics

Understanding the concepts concerning IP addressing is critical to understanding how IP routing works. A good understanding of IP addressing, and subsequently the art of subnetting, requires that you be comfortable with binary notation and math.

You already know that an IP address is a numeric identifier assigned to every machine on a network. This address tells where the device is located on the specific network.



EXAM WARNING

Keep in mind that an IP address is a software address. Don't confuse it with a hardware address. The hardware address is hard-coded into the machine itself or in the network interface card (NIC). Also keep in mind that starting with Windows 2000, Microsoft began listing IP address ranges in the same manner that Cisco does. This method, Classless Interdomain Routing (CIDR), lists the IP address followed by the number of ones in the subnet mask. For instance, 192.168.1.0 with a subnet of 255.255.255.0 is written as 192.168.1.0/24.

As a quick review, IP addresses are currently made up of 32 bits of information. These bits are divided into four sections (octets) that each contains 1 byte (6 bits). You will see IP addresses specified in three basic formats:

- Binary such as in 11000000.10101000.00000000.00000001
- Dotted-decimal such as in 192.168.0.1
- Hexadecimal such as in C0 A8 00 01

All three of these examples represent the same IP address. In reality, the computer can use only the binary version. The other two formats are provided because they are easier for people to understand and use.

There are three basic types of IP addresses:

- **Unicast addresses** IP addresses assigned to a single network interface that is attached on the network. Unicast IP addresses are used for one-to-one communications between hosts.
- **Broadcast addresses** IP addresses designed to be received and processed by every IP address located on a given network. They're basically one-to-many communications.
- **Multicast addresses** IP addresses where one or more IP nodes can listen in on the same network segment. Multicast IP addresses are also one-to-many communications.

Next, you should also understand the differences between routed and Network Address Translation (NAT) connections. NAT is the process of switching back and forth between the IP addresses used on an internal network, sometimes referred to as *private addresses*, and Internet IP addresses, sometimes known as *public addresses*.

There are three address blocks set aside and defined as private address space:

- **10.0.0.0 with a subnet mask of 255.0.0.0, or 10.0.0.0/8** This network is a private address space that has 24 host bits that can be used.
- **172.16.0.0 with a subnet mask of 255.240.0.0, or 172.16.0.0/12** This network is a private address space that has 20 host bits that can be used. This provides a range of 16 class B network IDs from 172.0.0.0/16 through 172.31.0.0/16.
- **192.168.0.0 with a subnet mask of 255.255.0.0, 192.168.0.0./16** This network is a private address space that has 16 host bits that can be used. This provides a range of 256 class C network IDs from 192.168.0.0/24 through 192.168.255.0/24.

Remember that private and public spaces do not overlap. Machines on an intranet with a private IP address cannot directly connect to the Internet. Instead, they must be connected indirectly via either a proxy server or NAT. Essentially, all of the computers on your intranet are masquerading behind a single public IP address.



EXAM WARNING

Understand the ranges and subnet masks used with private addressing. Know how NAT translates and connects for them.

Routed connections require a single public IP address for each connection to the Internet. Using NAT allows you to connect multiple private addresses to a single public IP address. This is done by translating and modifying packets to reflect the changed addressing information.

There are three basic components that make up NAT:

- **Translation** This component maintains the NAT table for inbound and outbound connections.
- **Addressing** This component is handled by a stripped-down version of a Dynamic Host Configuration Protocol (DHCP) server that assigns the IP address, subnet mask, default gateway, and IP address of the Domain Name System (DNS) server.
- **Name resolution** This component forwards all name-resolution requests to the DNS server defined on the Internet-connected adapter, and then returns the reply. It can be thought of as a DNS proxy.



EXAM WARNING

Understand the three components of NAT and how they interact with other Windows Server 2003 components such as DNS and DHCP.

Keep in mind that NAT is not always the solution. It is extremely limited when it comes to security. You cannot encrypt anything that is carrying or that has been derived from an IP address. Tracking hackers and other problems is also extremely difficult, because the source IP address is stripped away in the NAT process. Another problem arises when you try to use NAT with large networks that have many hosts attempting to communicate with the Internet at the same time. The size of the mapping tables in this kind of environment is overwhelming and can cause performance problems.

Another basic concept related to IP routing is how the Internet Control Message Protocol (ICMP) works. ICMP is a maintenance protocol used to create and maintain routing tables. It supports router discovery and advertisements to hosts on a network. Very simply, its designed to pass control and status information between TCP/IP devices. When a client computer starts up on your network, it usually has only a few entries in its routing table. When that host sends data out to a specific destination on a network, the host first checks its routing table to see if there is already an entry matching the destination's IP address. If no match is found, the packet is sent to the default gateway. When the default gateway receives the packet, it will check to see if it has a matching entry in its routing table. If it does, it forwards the packet to the destination. At the same time, it sends an ICMP message back to the originating host, telling that host about the better route available. ICMP can also let hosts on a network know if a specific router is still active by sending out periodic messages with this kind of information.

IP version 6

The Internet that we have all come to know and love uses IP version 4 (IPv4) and is based on 32-bit addressing. Because of the numerous disadvantages of IPv4, including the problem of limited address space that NAT addresses, a new proposal was put forth in 1995. Originally called Internet Protocol Next Generation (IPng), this proposal offered several improvements, including 128-bit addressing, global addressing, automatic configuration, built-in security, improved quality of service (QoS) support, and built-in mobility. The new version of IP became known as IP version 6 (IPv6). IP version 5 was reserved for a different proposal that was never adopted or implemented.

Because of the differences between IPv4 and IPv6, IPv6 is not backward-compatible with IPv4. The address syntax is just one example. IPv4 addresses can be expressed in the traditional 192.168.0.0/20 format. IPv6 has been forced to settle

Continued

on the colon-hexadecimal notation. The 128-bit block is divided into eight 16-bit blocks and delimited by colons. The 32-bit block of IPv4 is divided into four 8-bit blocks.

An example of an IPv6 unicast address is 3FFE:FFFF:2A:41CD:2AA:FF:FE5F:47D1. Leading zeros within a block are suppressed, but each block must contain at least one hexadecimal digit. Another example is FE80:0:0:0:2AA:FF:FE5F:47D1. Notice the 0 blocks. IPv6 allows for the compression of IPv6 addresses using double colons. The above address then becomes FE80::2AA:FF:FE5F:47D1. A multicast address such as FF02:0:0:0:0:0:1 would then become FF02::1.

IPv6 doesn't use subnet masks, but rather continues to use the CIDR notation. Using this notation, 3FFE:FFFF:2A:41CD::/64 would be a subnet identifier; 3FFE:FFFF:2A::/48 would be a route; and FF::/8 would be an address range.

Just remember that IPv6 is actually a suite of protocols. It replaces IP, ICMP, Internet Group Management Protocol (IGMP), and Address Resolution Protocol (ARP) in the TCP/IP protocol suite.

Routing Tables

A routing table is basically a list, a huge list sometimes, that is used to direct traffic on a network. The table includes information about what other networks are reachable from a given network by providing the network address and subnet mask, as well as the metric, or cost, for that specific network route. Another way to think of it is as a database of routes to other locations.

The way this works is simple. When a packet arrives at the routing device (which could be a dedicated router or a Windows Server 2003 computer), the routing table is queried to discover the lowest cost route to the intended destination. Sometimes when there is no specific information concerning that network in the routing table, the packet will be forwarded to the default gateway, assuming that the default gateway will get the packet where it needs to go.

The level of detail, or the number of routes in the table, depends on whether the IP node is a host or a router. Usually, a host will have fewer entries in this table than a router has in its table. For instance, it would be normal to find an IP host configured with a default gateway. Creating a default route in the table allows for the effective summarization of all destinations. Routing tables on a router, on the other hand, will normally contain an entry for each and every reachable network on the IP network system.

Let's turn our attention back to the table itself. Each of the rows in this list, or entries in this database, is commonly referred to as a *route*. There are three basic types of routes:

- **Host route** A route to a specific IP address in the network. A host is a particular computer, or more specifically, an interface on a computer or device. In these cases, the network mask is always 255.255.255.255 (/32). Host routes are typically

used for custom routes to specific hosts. This helps in the optimization and control of a network.

- **Network ID route** A route for classful, classless, subnet, and supernetted destinations. The network mask in these cases will be somewhere between 129.0.0.0 (/1) and 255.255.255.254 (/31).
- **Default route** A route to all other destinations. This route is used when the routing table cannot find a host or network ID route that matches the destination in the packet's header. The default route has a destination of 0.0.0.0 and a network mask of 0.0.0.0 (/0), and it is sometimes expressed as 0/0. All destinations not found in the routing table are simply forwarded to this destination, where the specific destination address will be found.

Each route in the routing table contains the necessary forwarding information for a range of destination IP addresses. This information includes two values for the destination IP address: the next-hop interface and the next-hop IP address. The *next-hop interface* is just a representation of the next physical or logical device over which the IP packet will be forwarded. The *next-hop IP address* is the IP address of the node to which the IP packet is being forwarded. In an indirect delivery, the next-hop IP address is the IP address of a directly reachable intermediate router to which the packet is being forwarded.

So, from this discussion, we glean that there is enough information contained in the route entry of a routing table to identify the destination, the next-hop interface, and the next-hop IP address, and to determine which route is the best when there is more than one route available to the intended destination. Let's break down the route entry into its component parts:

- **Destination** Sometimes referred to as the *network destination*, this value is usually a representation of the IP address that is reachable with this route. It is usually used in conjunction with the Network Mask field. This can be a network ID (classful, subnet, or supernet) or an IP address. Other terms that are sometimes used to represent the destination include *destination host*, *subnet address*, *network address*, and *default route*. The destination for a default route is 0.0.0.0. The destination for a limited broadcast is 255.255.255.255.
- **Network Mask** Sometimes referred to as the *netmask*, this value is a bit mask that is used to determine the significant bits in the Destination field. The 1 bit in a network mask identifies those bits that must match the Destination field for this route. The 0 bit indicates the bits that don't need to match the Destination field. This field is usually a string of contiguous 1 bits followed by a string of contiguous 0 bits. The combination of the destination and the network mask defines a range of IP addresses. A host route has a network mask of 255.255.255.255. With this mask, only an exact match with the destination would be able to use this route. On the other end of the spectrum, a default route has a network mask of 0.0.0.0. A mask of 0.0.0.0 allows any destination to use this route. A subnet or network route has a mask that exists somewhere between these two extremes.

- **Next-hop IP Address** This points to the IP address where the packet is to be forwarded using this route. It's sometimes also referred to as the *forwarding address* and most often called the *gateway*. This gateway must be directly reachable by this router by using the interface defined in the Interface field. This can be a hardware address, a network address, or sometimes even the address of the interface attached to the network.

NOTE



When working with routes of directly attached network segments, the Next-hop IP Address field can be set to the IP address of the network segment's interface. This is the default behavior of the IP routing table for the Windows 2003 Server family.

- **Interface** This is the logical or physical interface used when forwarding the packet using this specific route. It indicates the local area network (LAN) or demand-dial interface needed to reach the next router. The value here can be either a logical name or the IP address assigned to the interface. This can be the port number or some other logical identifier.

NOTE



Windows Server 2003 family uses the IP address assigned to the interface.

- **Metric** This field is where the route's cost is maintained. It's commonly used to store the *hop count*, or the number of routers between the host and the destination. It is also used by the route-determination process to choose among the many routes to the same location that might be possible. When there are multiple routes with the same destination and network mask, the route with the lowest metric value is used. Anything on the local subnet is always considered one hop. Each router crossed is counted as an additional hop. The lowest metric is usually the preferred one.
- **Protocol** This field shows how the route was learned. This column will normally list RIP, OSPF, or other routing protocols. If it lists Local, the router is not receiving routes.

Viewing Routing Tables

Viewing your routing tables in Windows Server 2003 is a simple procedure, but you must be logged on as an Administrator or, as a security best practice, using the Run As command. Follow these steps:

1. Select **Start | Control Panel | Administrative Tools | Routing and Remote Access**.
2. In the console tree on the left side of the **Routing and Remote Access** window, click the plus sign to the left of **Routing and Remote Access**.
3. Under that, you will see the name of the server. Click the plus sign there, and you'll see **IP Routing**.
4. Click the plus sign next to **IP Routing**, and you should see **Static Routes**.
5. Right-click **Static Routes** and choose **Show IP Routing Table** from the context menu.

You can also use a command-line utility to view the routing table. (Speed is one of the most important reasons for choosing to use the command line over a GUI tool.) To view the routing table from the command prompt, click **Start | All Programs | Accessories | Command Prompt**. This opens the command prompt window. At the prompt, type **route print** and press the **Enter** key. You'll now see a screen resembling the one shown in Figure 4.1.

Figure 4.1 Viewing the Routing Table from the Command Prompt

```

E:\WINDOWS\system32\cmd.exe
E:\>route print

IP4 Route Table
=====
Interface List
=====
0x1 ..... MS TCP Loopback interface
0x2 ...00 04 ac 1b 3f f2 ..... Intel 825x-based PCI Ethernet Adapter (10/100)
=====

Active Routes:
Network Destination    Netmask          Gateway         Interface       Metric
-----
0.0.0.0                0.0.0.0         192.168.0.1    192.168.0.5     1
127.0.0.0              255.0.0.0       127.0.0.1     127.0.0.1      1
192.168.0.0            255.255.255.0   192.168.0.5   192.168.0.5    30
192.168.0.5            255.255.255.255 127.0.0.1     127.0.0.1      30
192.168.0.255         255.255.255.255 192.168.0.5   192.168.0.5    30
224.0.0.0              240.0.0.0       192.168.0.5   192.168.0.5    30
255.255.255.255       255.255.255.255 192.168.0.5   192.168.0.5    1
Default Gateway:      192.168.0.1
=====

Persistent Routes:
None

E:\>
    
```

The routing table shown in Figure 4.2 (viewed from the Windows Server 2003 **Routing and Remote Access** utility) is for a computer running Windows Server 2003 Enterprise Edition with one 10MB network adapter, an IP address of 192.168.0.13, a subnet mask of 255.255.255.0, and a default gateway of 192.168.0.1.

Let's look at the individual rows more closely:

- The first row in the table, beginning with 0.0.0.0, is the default route.
- The second and third rows, beginning with 127.0.0.0 and 127.0.0.1, are the loop-back network.
- The fourth row, beginning with 192.168.0.0, is the local network.
- The fifth row, beginning with 192.168.0.13, is the local IP address.
- The second-to-last row, beginning with 224.0.0.0, is the multicast address.
- The final row, beginning with 255.255.255.255, is the limited broadcast address.

We'll now turn our attention to the upkeep of these tables. You can perform the main-

Figure 4.2 IP Routing Table

Destination	Network mask	Gateway	Interface	Metric	Protocol
0.0.0.0	0.0.0.0	192.168.0.1	Local Area Connection	1	Static (non demand-dial)
0.0.0.0	0.0.0.0	192.168.0.1	Local Area Connection	30	Network management
127.0.0.0	255.0.0.0	127.0.0.1	Loopback	1	Local
127.0.0.1	255.255.255.255	127.0.0.1	Loopback	1	Local
192.168.0.0	255.255.255.0	192.168.0.5	Local Area Connection	30	Local
192.168.0.5	255.255.255.255	127.0.0.1	Loopback	30	Local
192.168.0.255	255.255.255.255	192.168.0.5	Local Area Connection	30	Local
224.0.0.0	240.0.0.0	192.168.0.5	Local Area Connection	30	Local
255.255.255.255	255.255.255.255	192.168.0.5	Local Area Connection	1	Local


tenance of the routing tables manually or automatically. If you do it manually, you'll be using *static routing*. If you do it automatically, you'll be using *dynamic routing*. Let's take a closer look at these two concepts.

Static versus Dynamic Routing

Remember that the basic idea of routing is that each packet you find on your network has a source and a destination. That means that any device that receives the packet inspects the packet's headers to determine where it came from and where it's going. When the device has information about the network, such as how long it would take a packet to go from one point to another, that device can change the routing intelligently to improve the performance of the network.


Static routing uses manually configured routes. Here, there is no attempt to discover other routers or systems on a network. All entries into the routing table are entered by hand, and the routing table is used to get information to other networks. This type of routing works well with classless routing, because each route must be added with a network mask. It works well for small networks, but it doesn't scale well. Static routes are often used to connect to the Internet. Static routing is, however, not fault tolerant. Figure 4.3 shows a simple network using static routing.

Figure 4.3 Simple Network Using Static Routing



Dynamic routing doesn't depend on fixed, unchangeable routes to remote networks being added to the routing tables. In other words, you don't need to enter the routes by hand. Dynamic routing uses routing protocols to maintain the routing tables. Dynamic routing allows for the discovery of the networks surrounding the router by finding and communicating with other nearby routers in the network. Routes are discovered using routing protocol traffic and are then added or removed from IP routing tables as required. Dynamic routing can provide fault tolerance. When a route is unreachable, the route is removed from the routing table. Figure 4.4 shows a more complex network using dynamic routing.

Figure 4.4 A More Complex Network Using Dynamic Routing



In summary, static routing has two main advantages:

- It works well with classless routing.
- It works well with small networks.

Static routing also has two main disadvantages:

- It doesn't scale well.
- It is not fault tolerant.

For more complex networks, dynamic routing offers several advantages:

- It scales well with larger organizations.
- It is fault tolerant.
- It requires less administration than static routing.

Gateways

Although we've mentioned the term *default gateway* earlier in this chapter, we have not really gone into much detail about what a gateway is. Basically, a *gateway* is a device that connects networks using different communication protocols in a way that allows for information to pass from one network to the other. It both transfers and converts the information into a form that can be used by the protocols on the receiving network. Think of it as a TCP/IP node that has routing capabilities. In other words, a gateway is a kind of router. A router, by definition, is a device or computer that sends packets between two or more network segments as necessary, using logical network addresses, most often IP addresses. The default gateway is the path used to pass information when the device doesn't know where the destination is. More directly, a default gateway is a router that connects your host to remote network segments. It's the exit point for all the packets in your network that have destinations outside your network.

Planning a Routing Strategy for IP Multicast Traffic

Multicast traffic involves sending a message to multiple devices using a single (multicast) IP address. Multicasting is referred to as point-to-multipoint communication because the sender only has to send the message to one address to a group of computers that share a multicast group ID, which is an address from the Class D range.

Planning a Windows Server 2003 routing strategy in which multicast messages are sent involves the following steps:

1. Planning for the deployment of MADCAP servers (Multicast Address Dynamic Client Allocation Protocol). MADCAP is part of the Windows Server 2003 DHCP service, but works independently of DHCP.
2. Planning for deployment of routers that support IP multicasting. The routers need to be configured to use multicast routing protocols. Windows Server 2003 does not include multicast routing protocols, but RRAS supports multicast routing protocols such as Protocol Independent Multicast (PIM), Multicast Extensions to OSPF (MOSPF) and Distance Vector Multicast Routing Protocol (DVMRP).
3. Configuring the Internet Group Management Protocol (IGMP).
4. Configuring Multicast scopes on the MADCAP server, using administrative scoping for multicast addresses that are used on the internal network and global scoping for multicast addresses that are used on the Internet.
5. Configuring client computers to be MADCAP clients.

Multiple IP Addresses

Computers running Windows Server 2003 can have multiple IP addresses, even if the computer has only one NIC. In this case, if your network is divided into multiple logical IP network subnets, you can set up the single NIC to have multiple IP addresses. Then the address 192.168.0.10 could be used to communicate with the workstations and computers you have on the 192.168.0.0 subnet, and the address 192.168.1.10 could be used to communicate with the workstations and computers you have on the 192.168.1.0 subnet.


Keep in mind that if you are using a single NIC, the IP addresses must be assigned to either the same network segment or to segments that are part of the same single logical network. If your network is divided into multiple physical networks, you will need to use multiple NICs, with each card assigned an IP address from the different physical network segments.

Configuring Multiple Gateways



To install multiple gateways, follow these steps:

1. Select **Start | Control Panel | Network Connections**, and then select the connection you want to configure.
2. Click **Properties** and double-click **Internet Protocol (TCP/IP)** to open the **Internet Protocol (TCP/IP) Properties** dialog box, shown in Figure 4.5.
3. Click the **Advanced** button to open the **Advanced TCP/IP Settings** dialog box, shown in Figure 4.6.
4. On the **IP Settings** tab, you can add default gateways as you deem necessary. Click the **Add** button, and then type the gateway address in the **Gateway** text box, as shown in Figure 4.7.
5. The metric, as we have discussed previously, provides a relative cost of using this gateway, or route. When multiple gateways are available for a particular IP address, the gateway with the lowest metric will be used. If for some reason the Windows Server 2003 computer cannot communicate with the first gateway, it will try to use the gateway with the next lowest metric. By default, Windows Server 2003 assigns the metric to the gateway automatically. If you want to do so manually, uncheck the **Automatic metric** check box and enter a metric in the text box.

Figure 4.5 Internet Protocol (TCP/IP) Properties



Continued

Figure 4.6 The IP Settings Tab of the Advanced TCP/IP Settings

Figure 4.7 Enter the Gateway Address


EXAM
70-291
OBJECTIVE
3.1.1

Routing Protocols

Router discovery enables new, or rebooted, routers to configure themselves automatically. The two major and most common dynamic-routing protocols are RIP and OSPF. Both of these protocols are supported by the Windows Server 2003 family. Both are interior gateway protocols (IGPs) that use routers to communicate (not to be confused with the proprietary Cisco IGRP). But before we discuss these two protocols, we need to explore how protocols make routing decisions.

In general, routing protocols can use one of two different approaches to making routing decisions:

- **Distance vectors** A distance-vector protocol makes its decision based on a measurement of the distance between the source and the destination addresses.
- **Link states** A link-state protocol bases its decisions on various states of the links that connect the source and the destination addresses.

Distance-vector algorithms, also known as Bellman-Ford algorithms, periodically pass copies of their routing tables to their immediate network neighbors. The recipient adds what is called a distance vector, which is little more than a distance value, to the routing table it has just received, and then forwards it on to its immediate neighbors. The process

results in each router learning about the other routers and thereby developing a cumulative table of network distances to other routers. This table is then used to update the router's own routing table. Keep in mind that the only thing the router learns about is distance.

The main drawback to distance-vector routing is that it requires time for the changes in a network to propagate across the network. This makes distance-vector routing inappropriate for larger, more complex networks. The advantages of distance-vector routing are its ease of configuration, use, and maintenance. As we will discuss shortly, RIP is the epitome of distance-vector routing.

Link-state routing algorithms are usually known cumulatively as shortest path first (SPF) protocols. OSPF, which will be discussed shortly, is an example of this protocol group. These protocols maintain a complex database that describes the network's topology. Link-state protocols develop and maintain extensive information concerning the network's routers and how they interconnect. They do this by exchanging link-state advertisements (LSAs) with each other. Any change in the network will trigger the exchange of LSAs. Each router then constructs an extensive database using these received LSAs, so it can compute different routes and determine how reachable the networked destinations really are. This information is then used to update the routing table. Component failures and growth of the network are easily documented.

The main drawbacks to using link-state protocols involve the heavy use of bandwidth, memory, and processor time. Especially during the initial discovery process, link-state protocols flood the network with messages, thereby lowering the overall network efficiency. Also, overall, link-state protocols require more memory and higher processor speeds than distance-vector protocols need for efficient operation.

The main advantage of link-state protocols comes into play with large and complicated networks. A well-designed network will be more able to withstand the effects of unexpected changes using link-state protocols. Overhead caused by the frequent, time-driven updates required for distance-vector protocols can be avoided. Networks using a link-state protocol are also more scalable. For most large networks, the advantages of using link-state protocols will outweigh the disadvantages.

RIP

RIP is simple and easy to configure and is used widely in small and medium-sized networks. RIP is an IGP used to route data within autonomous networks. RIP does have performance limitations, however, that restrict its usefulness on medium-sized to large networks. RIP is a distance-vector routing protocol. This means that it distributes routing information in the form of a network ID and the number of hops (or the distance) from the destination. RIP has a maximum distance of 15 hops. Anything over that is considered unreachable.

There are two versions of RIP: version 1 described in RFC 1058 and version 2 described in RFC 1723. Windows Server 2003 supports both RIP versions.

RIP version 1 is a class-based routing protocol. Only the network ID is announced here. The message format for RIP version 1 is shown in Figure 4.8.

Figure 4.8 RIP Version 1 Message Format

Command	Version 01	Must be Zero	Family Identifier 00x02	Must be Zero	IP Address	Must be Zero	Must be Zero	Metric
1 byte	1 byte	2 byte	2 byte	2 byte	4 byte	4 byte	4 byte	4 byte

RIP version 2 is a classless routing protocol. This version includes both a network ID and a subnet mask in its announcement. It also provides more information, allowing for both authentication and a measure of security. The message format for RIP version 2 is shown in Figure 4.9.

Figure 4.9 RIP Version 2 Message Format

Command	Version 02	Must be zero	Family Identifier 00x02	Route Tag	IP Address	Subnet Mask	Next Hop	Metric
1 byte	1 byte	2 byte	2 byte	2 byte	4 byte	4 byte	4 byte	4 byte

There are several shortcomings to RIP version 1:

- RIP version 1 uses MAC-level broadcasting, requiring all hosts on a network to process all packets.
- RIP version 1 doesn't support sending a subnet address with the route announcement. This can be a problem when there is a shortage of available IP addresses.
- Because RIP version 1 route announcements are being addressed to the IP subnet and MAC-level broadcast, non-RIP hosts may also be receiving the RIP announcements, contributing to the broadcast clutter and possibly lowering the efficiency and performance of your network.
- By default, every 30 seconds, RIP routers broadcast lists of networks they can reach to every other adjacent router. Again, this can contribute to lower network performance.
- RIP version 1 does not handle subnetted addresses well, since it doesn't send the subnet address along with the broadcast.
- RIP version 1 provides no defense from a rogue router. A *rogue router* is an RIP router that advertises false or erroneous route information.
- RIP version 1 is difficult to troubleshoot. In general, most problems in RIP routing stem from incorrect configuration or from the propagation of bad routing information.

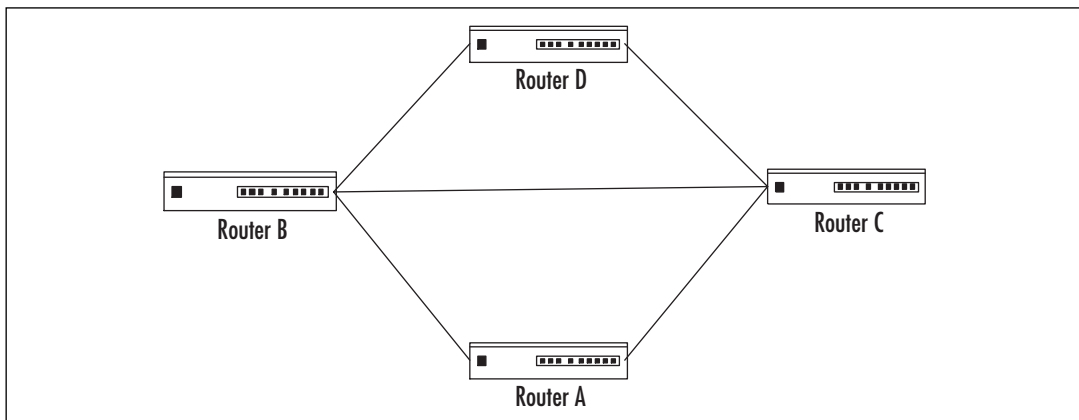
So, what does RIP version 2 do to attempt to correct the problems with RIP version 1?

- RIP version 2 advertisements include the subnet mask with the network ID.
- RIP version 2 sends multicast announcements to the multicast IP address 224.0.0.9 with a time to live (TTL) of 1 instead of broadcasting announcements, so it does not require IGMP.
- RIP version 2 allows for authentication to substantiate the source of the incoming routing announcements.
- RIP version 2 is compatible with RIP version 1.

RIP routers begin with a basically empty routing table and start sending out announcements to the networks to which they're connected. These announcements include the appropriate routes listed for all interfaces in the router's routing table. The router also sends out a RIP General Request message asking for information from any router receiving the message. These announcements can be broadcast or multicast. Other routers on other networks hear these announcements and add the original router and its information to their own routing tables. They then respond to the new router's request for information. The new router hears the announcements from these other routers on the network and adds them and their information to its own routing table.

After the initial setup, the RIP router will send out information based on its routing table. The default time period is 30 seconds. Over time, the routers of the network develop a consensus of what the network looks like. The process of developing this consensual perspective of the network's topology is known as *convergence*. Basically, this means that the network's routers individually agree on what the network looks like as a group. It is this very process of convergence, however, that can sometimes lead to problems. A typical network using convergence is shown in Figure 4.10. One of the occasional problems that occurs is called *counting to infinity*. Let's look at how that happens.

Figure 4.10 Typical Network Using Convergence



In our example, we will assume that Router A has failed. With its failure, all the hosts on the A network will no longer be accessible from the other three networks. After missing six

updates from Router A, Router B will invalidate its B–A route and advertise its unavailability. Routers C and D remain ignorant of the failure of Router A until notified by Router B. At this point, both Router B and Router D still think they can get to Router A through Router C, and they raise the metric of this route accordingly. So, Routers B and D send their next updates to Router C. Router C, having timed out its route to Router A, still thinks it has access through Router B or Router D. Thus, a loop is formed between Routers B, C, and D, based on the mistaken belief that both Routers B and C can still access Router A. With each iteration of updates, the metrics are incremented an extra hop for each route. This count speeds up the process by which the router approaches its definition of infinity—the point where the router says the destination is unreachable.

There are two methods of preventing this counting to infinity loop: split horizon and triggered updates. If the router is implementing split horizon, routes will not be announced back over the interfaces by which they were learned. The limitation of the split-horizon approach is that a route will not timeout until it has been unreachable for six tries, so each router has five opportunities to transmit incorrect information to the neighboring routers. If the router is implementing split horizon with poison reverse, routes learned on interfaces are announced back as unreachable. Split horizon with poison reverse is much more dependable than simple split horizon. However, although split horizon with poison reverse will stop loops in small networks, loops are still possible on larger, multipath networks.

Fault tolerance in RIP networks is based on the timeout of RIP-learned routes. When changes happen in the network, RIP routers send out triggered updates, rather than waiting for a scheduled time for routing announcements. These triggered updates contain the routing update and are sent immediately. Triggered updates are nothing more than a method of speeding up split horizon with poison reverse. However, triggered updates are not foolproof. While the triggered updates are being propagated around the network, routers that have not received the triggered update are still sending out the incorrect information. It's possible that a router could receive the triggered update and then receive an update from another router reintroducing the incorrect information, so the count-to-infinity problem, though not as likely, is still possible.

OSPF

Because OSPF is designed to work inside the network area, it belongs to a group of protocols called IGRPs. OSPF is defined in RFC 2328 and its purpose is to overcome the shortcomings of both versions of RIP when they are used for large organizations. OSPF is designed for use on large or very large networks. OSPF is much more efficient than RIP, and it also requires much more knowledge and experience to set up and administer.

There are many reasons why OSPF is a better choice for large networks than either version of RIP, including the following:

- Faster detection and changes of the network topology. This means less chance of encountering the count-to-infinity problem.
- OSPF routes are loop-free.

- In OSPF, large networks can be broken down into smaller contiguous groups of networks, called *areas*. (RIP does not allow for the subdivision of a network into smaller components.) Routing table entries can then be minimized by using the technique called *summarizing*. Summarizing allows for the creation of default routes for routes outside the area.
- The subnet mask is advertised with OSPF. This provides support for disjointed subnets and supernetting.
- Route exchanges between OSPF routers can be authenticated.
- Because external routes can be advertised internally, OSPF routers can calculate least-cost routes to external destinations.

The packet header structure for OSPF is shown in Figure 4.11.

Figure 4.11 The OSPF Packet Header Structure

Version Number	Type	Packet Length	Router ID	Area ID	Checksum	Authentication Type	Authentication
1 byte	1 byte	2 byte	4 byte	4 byte	2 byte	2 byte	8 byte

There are five basic messages that are attached to this header structure:

- **Hello packet** Used to discover and maintain information about neighboring routers.
- **Database Description packet** Used to summarize database contents.
- **Link-State Request packet** Used to initialize the database download from another router.
- **Link-State Update packet** Used to update other routers with the information contained in the local router's database.
- **Link-State Acknowledgment packet** Used to acknowledge flooding of information from other routers.

OSPF is a link-state routing protocol that uses LSAs to send information to other routers in the same area, known as *adjacencies*. Included in the LSA is information about interfaces, gateways, and metrics. OSPF routers collect this information into a link-state database (LSDB) that is shared and synchronized among the various routers. Using this database, the various routers are able to calculate the shortest path to other routers using the SPF algorithm. The cost of each router interface is assigned by the network administrator. This number can include the delay, the bandwidth, and any monetary cost factors. The accumulated cost of any OSPF network can never be more than 65,535. So, the way OSPF works can be divided into three main phases:

- The LSDB is put together from neighboring routers.
- The shortest path to each node is then calculated.
- The router creates the routing table entries containing the information about the routes.

When the router initializes, it sends out an LSA that contains only its own configuration. Each router has its own unique ID that it sends out with the LSA. This ID is not, however, the destination address of that router. Usually, it is the highest IP address assigned to that router, thereby ensuring that each router ID is unique. Over time, the router receives LSAs from other routers. The original router includes these routes in its own LSA and eventually will again send out its LSA, now containing the information it received. This process is called *flooding*. Every router in the area will soon have the information from all other routers in the area.

After the LSDB is compiled, the router determines the lowest cost path to each destination using the Dijkstra algorithm. Now, every other router and network reachable from that router will have a shortest, least-cost path calculated. The resulting data structure is called the *SPF tree*. The SPF tree is different for each router in the network, because the routes are calculated based on each router as the root of the tree. After the SPF tree is calculated, the routing table is created from the information it contains. An entry will be created for each network in the area of the router. The routing table will contain the network ID, the subnet mask, the IP address of the appropriate router for traffic to be directed to for that network, the interface over which the router is reachable, and the OSPF-calculated cost to that network. This cost is the metric unit, not the hop count as it would be in an RIP-routed network.



NOTE

The Dijkstra algorithm is part of a branch of mathematics called graph theory. This algorithm was developed to ascertain the least-cost path between a single vertex and the other vertices in a graph. If you're interested in the computations that go into working with Dijkstra's algorithm, you can find more information at www-b2.is.tokushima-u.ac.jp/~ikedasuuri/dijkstra/Dijkstra.shtml and <http://ciips.ee.uwa.edu.au/~morris/Year2/PLDS210/dijkstra.html>.

OSPF router interfaces must be configured for an appropriate network type because the OSPF message address will be set for the network type specified. There are three network types supported by OSPF:

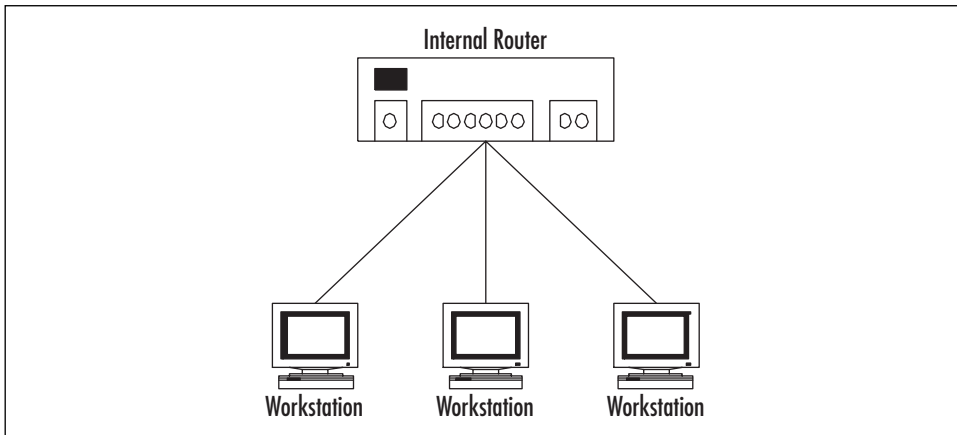
- **Broadcast** This type of network is connected by two or more routers and broadcast traffic is passed between them. Examples of broadcast networks include Ethernet and FDDI.
- **Non-broadcast multiple access (NBMA)** Broadcast traffic doesn't pass on this network, even though it is connected by two or more routers. OSPF must be configured to use IP unicasting instead of multicasting. Examples of this type of network include Asynchronous Transfer Mode (ATM) and Frame Relay.

- **Point-to-Point** Only two routers can be connected using this type of network. Examples of Point-to-Point networks include WAN links like Digital Subscriber Line (DSL) or Integrated Services Digital Network (ISDN).

Your network is divided into areas by placing routers in specific locations to join or divide the network in the manner you want. What the router does and what designation it is given are determined by its location and role in the network area. The roles that an OSPF router might file include the following:

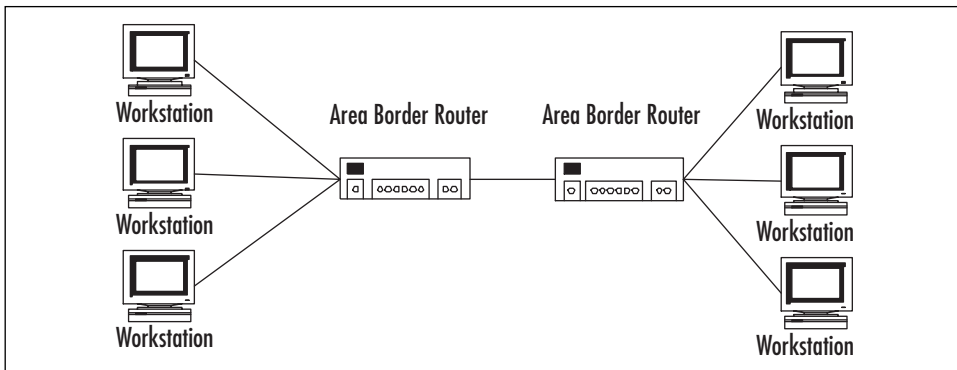
- **Internal router** All interfaces of the router are connected to the same area, as illustrated in Figure 4.12. An internal router will have only one LSDB because it is connected to only one area.

Figure 4.12 An Internal Router



- **Area border router (ABR)** When a router's interfaces are connected to different areas, that router is an ABR. An ABR has one LSDB for each area it's connected to, as illustrated in Figure 4.13.

Figure 4.13 An Area Border Router



- **Backbone router** If one of a router's interfaces is on the backbone area, that router is considered a backbone router. This applies to both ABRs and internal routers.
- **Autonomous system boundary router (ASBR)** If a router exchanges routes with sources outside the network area, it is known as an ASBR. These special routers announce external routes throughout the area network.

Using netsh Commands

Administering your routing server through the Routing and Remote Access console is easy, but in order to pass the exam, as well as get by in the real world, you need to know how to use the command-line utility netsh, introduced in Chapter 3. You might wonder why anyone would want to use the command line when a perfectly acceptable and easy-to-use console is available. There are two main reasons:

- You can administer a routing server much more quickly from the command line. This might be especially important over slow network links.
- You can administer multiple routing servers more efficiently and consistently by creating scripts using these commands, which can then be run on many servers.

The Netsh utility is available in the Windows 2000 Resource Kit and is a standard command in Windows XP and Windows Server 2003. This utility displays and allows you to manage the configuration of your network, including both local and remote computers. It is designed to simplify the process of creating command-line scripts such as batch files. The utility itself is little more than a command interpreter that connects and interfaces with a number of services and protocols through the aid of a number of dynamic link libraries (DLLs). Each of these DLLs provides the utility with an extensive set of commands that applies specifically to that DLL's service or protocol. These DLLs are referred to as *helper files*, and sometimes helper files are used to extend other helper files.

You can use the Netsh utility to perform the following tasks:

- Configure interfaces
- Configure routing protocols
- Configure filters
- Configure routes
- Configure remote access behavior for Windows 2000 and Windows Server 2003-based remote access routers that are running RRAS
- Display the configuration of a currently running router on any computer
- Use the scripting feature to run a collection of commands in batch mode against a specific router

The syntax for the Netsh utility is as follows:

```
netsh [-r router name] [-a AliasFile] [-c Context] [Command |
-f ScriptFile]
```

Context strings are appended to a command and passed to the associated helper file. The helper file can have one or more entry points that are mapped to contexts. The context can be any of the following: **DHCP**, **ip**, **ipx**, **netbeui**, **ras**, **routing**, **autodhcp**, **dnsproxy**, **igmp**, **mib**, **nat**, **ospf**, **relay**, **rip**, and **wins**. Under Windows XP, the available contexts include **AAAA**, **DHCP**, **DIAG**, **IP**, **RAS**, **ROUTING**, and **WINS**. Appending a specific context to the input string makes a whole different set of commands available that are specific to that context.

The easiest way to learn how the Netsh utility works is by viewing its help information. Open a command prompt window on your Windows Server 2003 computer and enter the **netsh** command at the prompt. The command prompt changes to the **netsh** prompt. Enter a **?** to display a list of available commands, as shown in Figure 4.14. To see the subcontexts and commands that are available to use with the routing context, type **routing ?** at the **netsh** prompt (or simply type **netsh routing ?** at the command prompt), and then press **Enter**. You can get command-line help for each command by typing **netsh**, followed by the command, followed by **?**.

Figure 4.14 Type **?** at the netsh Command Prompt to View Available Commands

```
c:\E:\WINDOWS\system32\cmd.exe
E:\>netsh ?

Usage: netsh [-a AliasFile] [-c Context] [-r RemoteMachine] [-u [DomainName\]Use
rName] [-p Password ! *]
           [Command ! -f ScriptFile]

The following commands are available:

Commands in this context:
?          - Displays a list of commands.
aaaa      - Changes to the 'netsh aaaa' context.
add       - Adds a configuration entry to a list of entries.
delete    - Deletes a configuration entry from a list of entries.
dhcp      - Changes to the 'netsh dhcp' context.
diag      - Changes to the 'netsh diag' context.
dump      - Displays a configuration script.
exec      - Runs a script file.
help      - Displays a list of commands.
interface - Changes to the 'netsh interface' context.
ipsec     - Changes to the 'netsh ipsec' context.
ras       - Changes to the 'netsh ras' context.
routing   - Changes to the 'netsh routing' context.
rpc       - Changes to the 'netsh rpc' context.
set       - Updates configuration settings.
show     - Displays information.
wins     - Changes to the 'netsh wins' context.

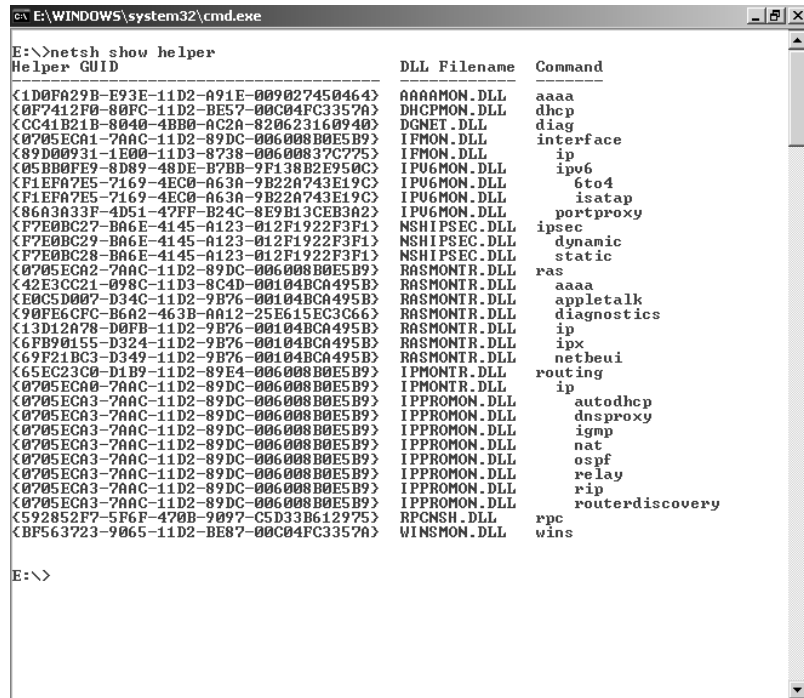
The following sub-contexts are available:
aaaa dhcp diag interface ipsec ras routing rpc wins

To view help for a command, type the command, followed by a space, and then
type ?.

E:\>
```

Rather than entering commands through the netsh utility as shown in Figure 4.14, it is more efficient to use the DLLs without needing to load the Netsh shell. This reduces the amount of coding time required, and you can use multiple DLLs within a single script. To use Netsh commands this way, follow the **netsh** command with the name of the DLL and the command string. For example, to use the **show helper** command to see a complete list of the available DLLs, type **netsh show helper**, as shown in Figure 4.15.

Figure 4.15 Type **netsh show helper** at the Command Prompt to View Available DLLs



```

c:\E:\WINDOWS\system32\cmd.exe
E:\>netsh show helper
Helper GUID                                     DLL Filename  Command
-----
<1D0FA29B-E93E-11D2-A91E-009027450464>        AAAAMON.DLL   aaaa
<0F7412F0-80FC-11D2-BE57-00C04FC3357A>        DHCPMON.DLL  dhcp
<CC41B21B-8040-4BB0-AC2A-820623160940>        DGNET.DLL    diag
<0705ECA1-7AAC-11D2-89DC-006008B0E5B9>        IFMON.DLL    interface
<89D00931-1E00-11D3-8738-00600837C775>        IFMON.DLL    ip
<05BB0FE9-8D89-48DE-B7BB-9F138B2E950C>        IP6MON.DLL   ipv6
<F1EFA7E5-7169-4EC0-A63A-9B22A743E19C>        IP6MON.DLL   6to4
<F1EFA7E5-7169-4EC0-A63A-9B22A743E19C>        IP6MON.DLL   isatap
<86A3A33F-4D51-47FF-B24C-8E9B13CEB302>        IP6MON.DLL   portproxy
<F7E0BC27-BA6E-4145-A123-012F1922F3F1>        NSHIPSEC.DLL ipsec
<F7E0BC29-BA6E-4145-A123-012F1922F3F1>        NSHIPSEC.DLL dynamic
<F7E0BC28-BA6E-4145-A123-012F1922F3F1>        NSHIPSEC.DLL static
<0705ECA2-7AAC-11D2-89DC-006008B0E5B9>        RASMONTR.DLL ras
<42E3CC21-098C-11D3-8C4D-00104BCA495B>        RASMONTR.DLL aaaa
<E0C5D007-D34C-11D2-9B76-00104BCA495B>        RASMONTR.DLL appletalk
<90FE6CFC-B6A2-463B-AA12-25E615EC3C66>        RASMONTR.DLL diagnostics
<13D12A78-D0FB-11D2-9B76-00104BCA495B>        RASMONTR.DLL ip
<6FB90155-D324-11D2-9B76-00104BCA495B>        RASMONTR.DLL ipx
<69F21BC3-D349-11D2-9B76-00104BCA495B>        RASMONTR.DLL netbeui
<65EC23C8-D1B9-11D2-89E4-006008B0E5B9>        IPMONTR.DLL routing
<0705ECA0-7AAC-11D2-89DC-006008B0E5B9>        IPMONTR.DLL ip
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL autodhcp
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL dnstproxy
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL igmp
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL nat
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL ospf
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL relay
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL rip
<0705ECA3-7AAC-11D2-89DC-006008B0E5B9>        IPPROMON.DLL routerdiscovery
<592852F7-5F6F-470B-9897-C5D33B612975>        RPCSSH.DLL  rpc
<BF563723-9065-11D2-BE87-00C04FC3357A>        WINSMON.DLL wins

E:\>

```

As you can see in Figure 4.15, when the script is processed, you see the results of the script and then are returned to the command prompt, from which you can execute your next script.

Using Netsh with Nested Contexts

There are times when using Netsh with simple commands is not sufficient for the tasks you want to accomplish. Sometimes, you will need to create scripts with nested contexts. Let's take look at an example to add an interface to the network. The syntax of the command is as follows:

```
Add interface [[InterfaceName=]][InterfaceName=]InterfaceName
[[IcmpPrototype=]{icmptrtrv1 | icmptrtrv2 | icmptrtrv3 |
    igmpproxy}]
[[IfEnabled=]{enable | disable}] [[RobustVar=]Integer]
[[GenQueryInterval=]Integer] [[GenQueryRespTime=]Integer]
[[StartupQueryCount=]Integer]
[[StartupQueryInterval=]Integer]
[[LastMemQueryCount=]Integer]
[[LastMemQueryInterval=]Integer] [[AccNonRtrAlertPkts=]{yes
    | no}]
```

For our example, we'll use this command to configure IGMP on a specified device. We type in the following command:

```
netsh routing ip igmp add interface "Local Area Connection"
    startupqueryinterval = 21
```

This command modifies a default startup query interval to 21 seconds with IGMP configuration of the interface named Local Area Connection.

EXAM
70-293
OBJECTIVE
3.1

Evaluating Routing Options

In order to make good decisions about routing in your network, you need to evaluate potential network traffic, as well as the number and types of hardware devices and applications used in your environment. For the most part, the heavier the routing demand, the higher the need for dedicated hardware routers. Lighter routing demands can be met sufficiently by less expensive software routers. Your routing decisions should be based on your knowledge and understanding of both options.

Selecting Connectivity Devices

For small, segmented networks with relatively light traffic between subnets, a software-based routing solution such as the Windows Server 2003 RRAS might be ideal. On the other hand, a large number of network segments with a wide range of performance requirements would probably necessitate some kind of hardware-based routing solution. Evaluating your

routing options includes selecting the proper connectivity devices: hubs, bridges, switches, or routers. You also should understand where these devices fit in the OSI reference model.

Head of the Class...

A Review of the OSI Model

The Open System Interconnection (OSI) reference model is an International Organization for Standardization (ISO) standard for worldwide communications. OSI defines a network framework for implementing an agreed-upon format for communicating between vendors. The model identifies and defines all the functionality required to establish, use, define, and dismantle a communication session between two network devices, no matter what the device is or who manufactured it.

All communication processes are defined in seven distinct layers with specific functionality. Microsoft and other proprietary systems may combine multiple-layer functionality into one layer in their particular version, but most, if not all, of the functionality of the original OSI model layers are incorporated. It is for this reason that most discussions of computer-to-computer communication begin with a discussion of this model. Table 4.1 shows the layers in the OSI reference model.

Table 4.1 The OSI Reference Model Layers

Layer	Description
7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

Layer 1 of the OSI reference model is often referred to as the bottom layer. This is the Physical layer, which is actually responsible for the transmission of the data. As a result, the Physical layer operates with only ones and zeros. It receives incoming streams of data, one bit at a time, and passes them up to the Data Link layer. Examples of transmission media associated with Layer 1 include coaxial cabling, twisted-pair wiring, and fiber-optic cabling.

Layer 2 is the Data Link layer, which is responsible for providing end-to-end validity of the data being transmitted. This layer deals with frames. The frame contains the data and local destination instructions. This means that the Physical and Data Link layers provide all the information required for communication on the local LAN. Figure 4.16 illustrates a Data Link layer domain.

At Layer 3, the Network layer, internetworking is enabled and the route to be used between the source and the destination is determined. There is, however, no

Continued

native transmission error detection/correction method. Some manufacturers' Data Link layer technologies support reliable delivery, but the OSI reference model does not make this assumption. For this reason, Layer 3 protocols such as IP assume that Layer 4 protocols such as TCP will provide this functionality. Figure 4.17 illustrates a network similar to the one shown in Figure 4.16, but with a second, identical network connected via a router. The router effectively isolates the two Data Link layer domains. The only way the two domains can communicate is via the use of Network layer addressing.

The Network layer implements a protocol that can transport data across the LAN segments or even across the Internet. These protocols are known as *routable protocols* because their data can be forwarded by routers beyond the local network. These protocols include IP, Novell's Internetwork Packet Exchange (IPX), and AppleTalk. Each of these protocols has its own Layer 3 addressing architecture. IP has emerged as the dominant routable protocol. Unlike the first two layers, which are required for all applications, the use of the Network layer is required only if the two communicating systems reside on different networks or if the two communicating applications require its service.

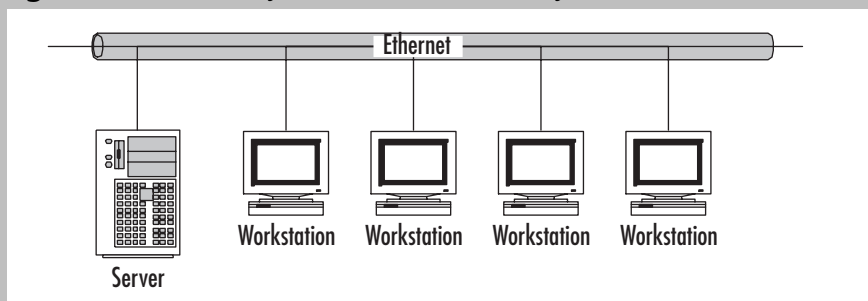
As with the Data Link layer, the fourth layer, the Transport layer, is responsible for the end-to-end integrity of data transmissions. The main difference is that the Transport layer can provide this function beyond the local LAN. The layer detects if packets are damaged or lost in transmission and automatically requests the data to be retransmitted. This layer is also responsible for resequencing any data packets that arrived out of order.

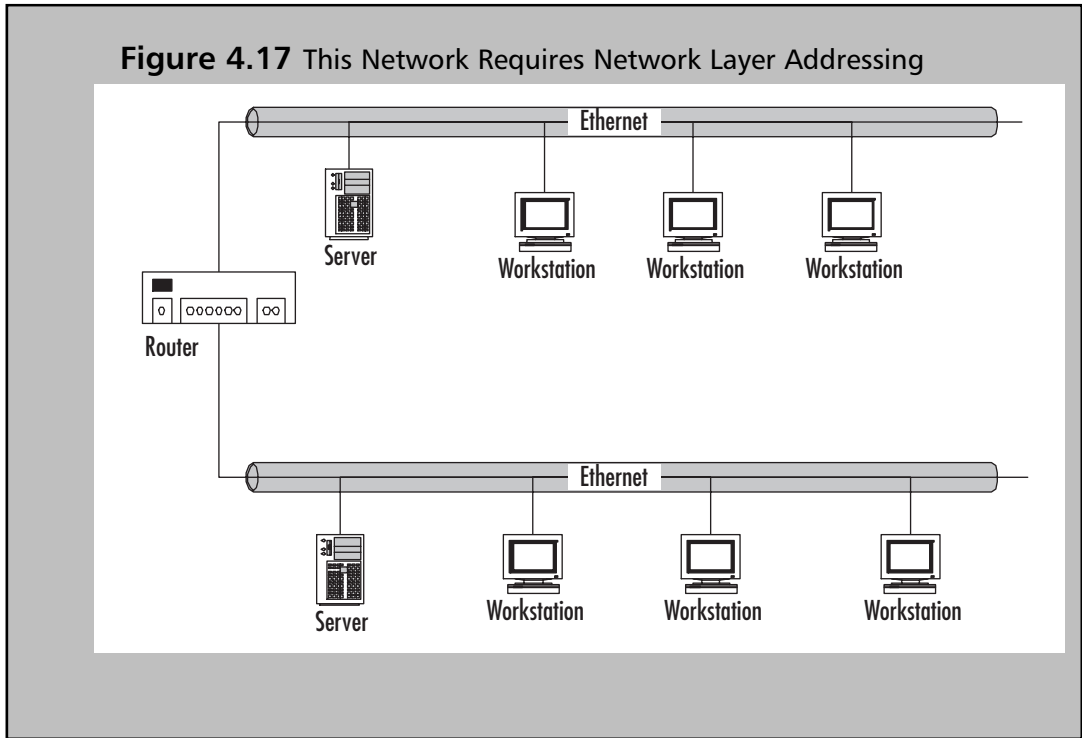
Layer 5 of the OSI model is the Session layer. Many protocols handle the functionality of this layer in the same layer they handle the functionality of the Transport layer. Examples of Session layer services include Remote Procedure Calls (RPCs) and quality of service (QoS) protocols such as RSVP, the bandwidth reservation protocol.

Layer 6, the Presentation layer, is responsible for how the data is encoded. Not every computer uses the same data-encoding scheme. This layer is responsible for translating data between otherwise incompatible encoding schemes. This layer can also be used to provide encryption and decryption services.

Layer 7 is the Application layer. This layer provides the interface between user applications and network services.

Figure 4.16 The Physical and Data Link Layers





Hubs

Hubs, sometimes referred to as *repeaters*, are devices used to connect communication lines in a central location and help provide common connections to all other devices on the network. A hub usually has one input and several outputs. These outputs are known as *ports*, but don't confuse them with TCP/IP ports (as in port 80, the one used for HTTP traffic). These ports are just connections and nothing more. They generally accept RJ-45 connectors. Think of a hub as like the center of an old wagon wheel with all the spokes radiating out to the other part of the wheel.

A hub simply takes the data that comes into its ports and sends it out on the other ports of the hub. For this reason, it is sometimes referred to as a *repeater*. It doesn't provide

or perform any filtering or redirection of the data from the various sources plugged into it. Hubs are commonly used to connect various network segments of a LAN.

Hubs generally come in three flavors:

- **Passive** Serves simply as a pipeline allowing data to move from one device, or network segment, to another.
- **Intelligent** Sometimes referred to as an *active, managed, or manageable* hub, it includes additional features that allow you to monitor the traffic passing through the hub and configure each port for specific purposes.
- **Switching** Reads the destination address of each packet and forwards that packet to the correct port. Most hubs of this variety also support load balancing.

Bridges

There are several definitions for a *bridge*, each carrying a specific meaning when used in a particular context. In one context, a bridge can be thought of as a gateway, connecting one network to another using the same communication protocols and allowing the information to be passed from one to the other. In another context, a bridge can be used to connect two networks with dissimilar communication protocols at the Data Link layer (Layer 2), in much the same manner as a router itself. There is also a bridge called a *bridge router*, which supports the functions of both the bridge and the router using Layer 2 addresses for routing.

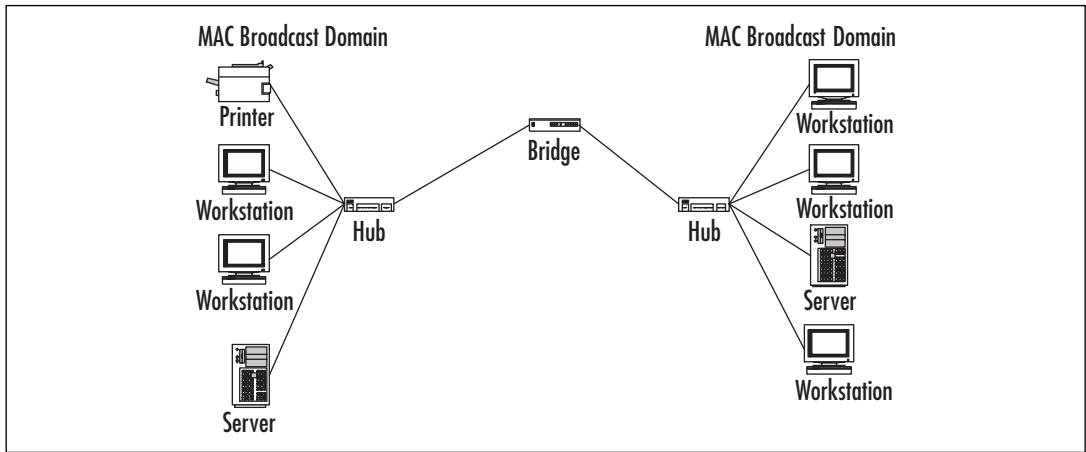
Here, we'll look at the traditional bridge and the context that is most often associated with this device. Bridges work at both the Physical (Layer 1) and Data Link (Layer 2) layers of the OSI reference model. That means that a bridge knows nothing about protocols but forwards data depending on the destination address found in the data packet. This destination address is not an IP address, but rather a Media Access Control (MAC) address that is unique to each network adapter card. For this reason, bridges are often referred to as *MAC bridges*.

Basically, all bridges work by building and maintaining an address table. This table includes information such as an up-to-date listing of every MAC address on the LAN, as well as the physical bridge port connected to the segment on which that address is located.

There are three basic types of bridges:

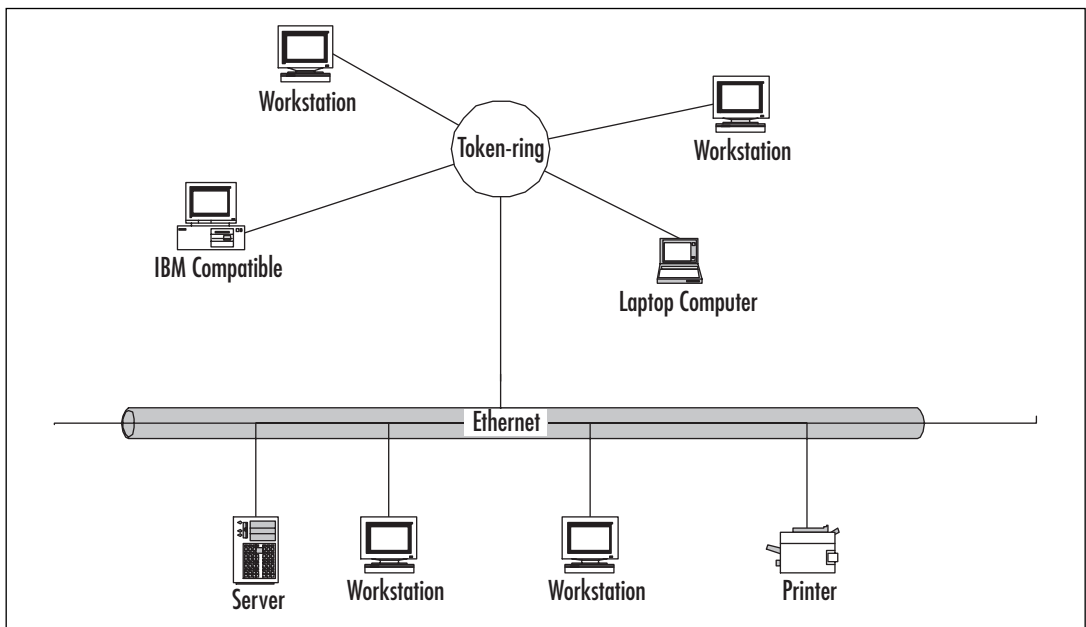
- **Transparent bridge** Links together segments of the same type of LAN. A transparent bridge effectively isolates the traffic from one LAN segment from the traffic of another LAN segment, as shown in Figure 4.18.

Figure 4.18 Transparent Bridge



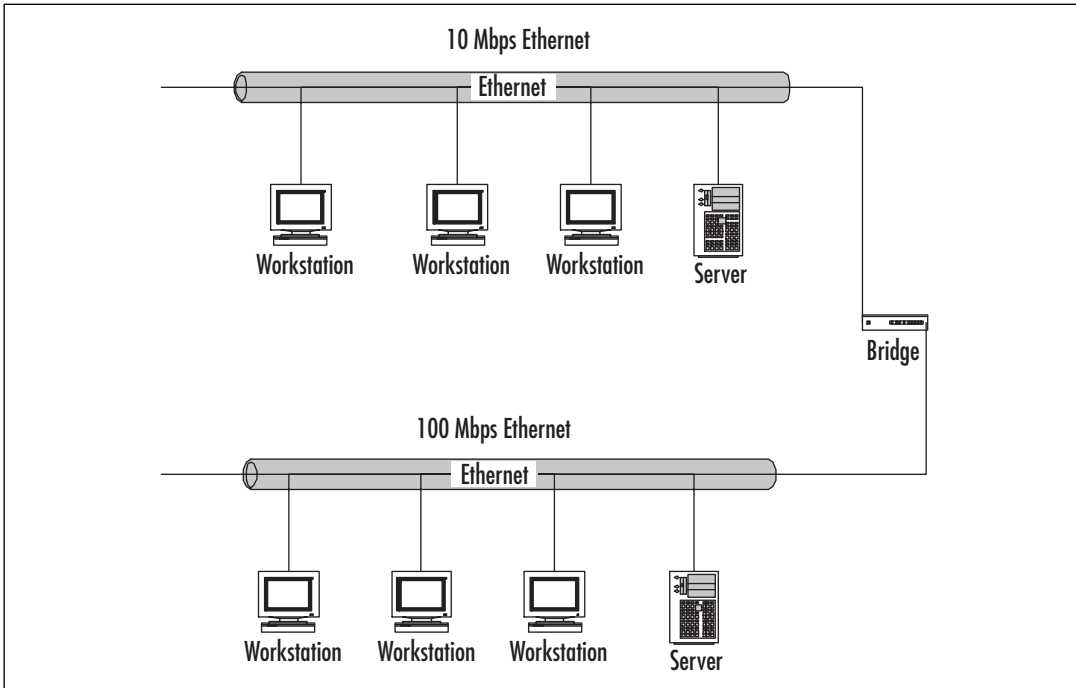
- Translating (or translational) bridge** Like a transparent bridge, links together segments of the same type of LAN, but also can provide conversion processes needed between different LAN architectures. This allows you to connect a Token Ring LAN to an Ethernet LAN, as shown in Figure 4.19.

Figure 4.19 Translating Bridge



- **Speed-buffering bridge** Used to connect LANs that have similar architectures but different transmission rates. Figure 4.20 shows how you might use a speed-buffering bridge to connect a 10-Mbps Ethernet network to a 100-Mbps Ethernet network.

Figure 4.20 Speed-buffering Bridge



Bridges are self-learning, so the administrative overhead is small. The functionality of bridges has been built into routers, hubs, and switches.

Switches

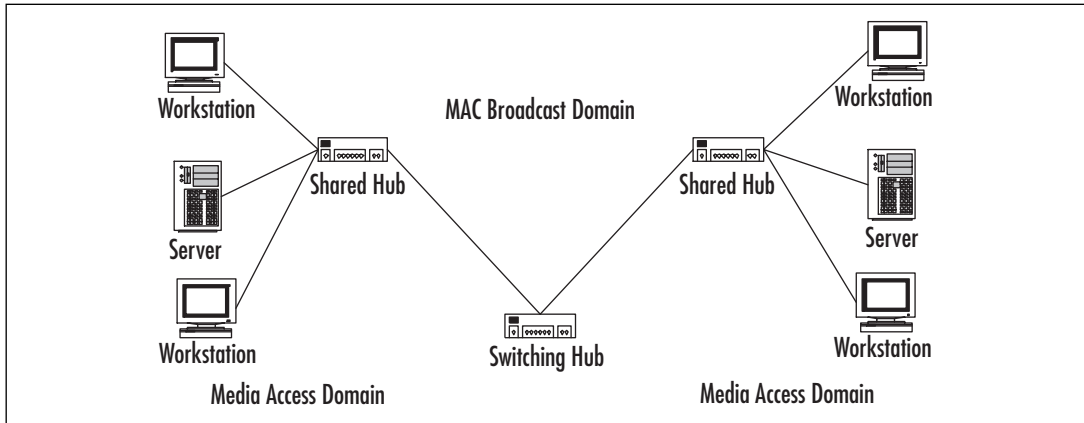
Switches are like bridges, except that they have multiple ports with the same type of connection (bridges generally have only two ports) and have been described as nothing more than fast bridges. Switches are used on heavily loaded networks to isolate data flow and improve the network performance. In most cases, most users get little, if any, advantage from using a switch rather than a hub.

That's not to oversimplify and suggest that a switch doesn't have many benefits. Switches can be used to connect both hubs and individual devices. These approaches are known as *segment switching* and *port switching*, respectively.

Segment switching implies that each port on the switch functions as its own segment. This process tends to increase the available bandwidth, while decreasing the number of devices sharing each segment's bandwidth, but at the same time maintaining the Layer 2

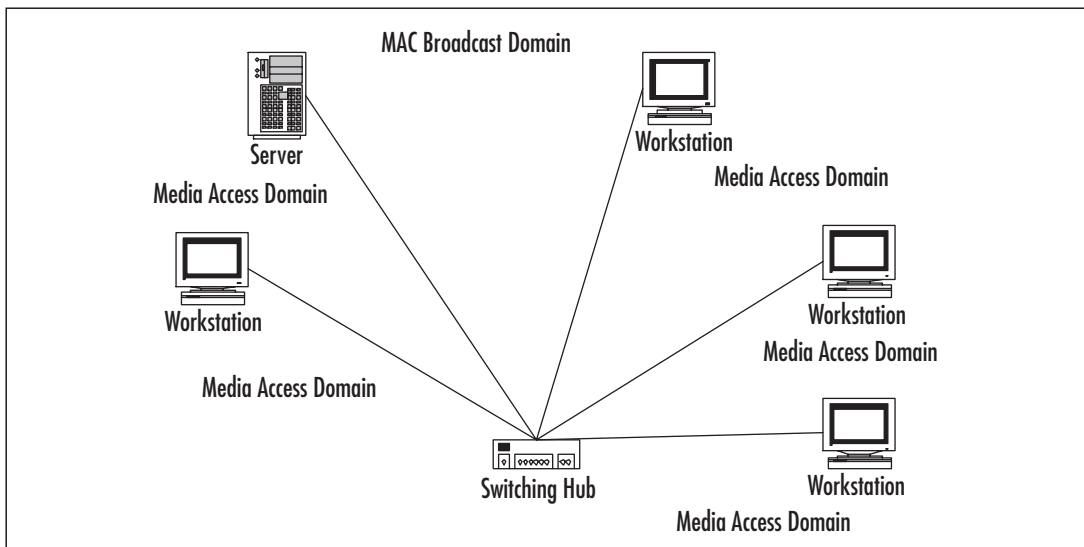
connectivity. Each shared hub and the devices that are connected to it make up their own media access domain, while all devices in both domains remain part of the same MAC broadcast domain. Figure 4.21 illustrates how a segment-switched LAN can be divided to improve performance.

Figure 4.21 Segment Switching



Port switching implies that each port on the switching hub is directly connected to an individual device. This makes the port and the device their own self-contained media access domain. All of the devices in the network still remain part of the same MAC broadcast domain. Figure 4.22 illustrates how the media access and MAC broadcast domains are configured in a port-switched LAN.

Figure 4.22 A Port-switched LAN



Layer 2 Switches

Layer 2 switches, operating at the Data Link layer, can be programmed to respond automatically to a wide range of circuit conditions. By monitoring control and data events, these switches automatically reroute circuits or switch to backup equipment, as the need requires. These switches operate using physical network, or MAC, addresses. These switches will be fast but not terribly smart. They only look at the data packet to find out where it's headed.

Layer 3 Switches

Layer 3 switches, operating at the Network layer, are designed for disaster recovery service (or, more importantly, for disaster avoidance). These network backup units are usually designed specifically to provide high levels of automation, intelligence, and security. Layer 3 switches use routing protocols such as RIP or OSPF to calculate routes and build their own routing tables.

Layer 3 switches use network or IP addresses to identify locations on the network, identifying the network location as well as the physical device. These switches are smarter than Layer 2 switches. They incorporate routing functions to actively calculate the best way to get a packet to its destination. Unless their algorithms and processor support high speeds, though, these switches are slower.

Layer 4 Switches

Layer 4 switches, operating at the Transport layer, allow network managers to choose the best method of communicating for each switching application. Because Layer 4 coordinates communication between systems, these switches are able to identify which application protocols (HTTP, SMTP, FTP, and so forth) are included in the packets, and they use this information to hand off the packet to the appropriate higher layer software. This means that Layer 4 switches make their packet-forwarding decisions based not just on the MAC and IP addresses, but also on the application to which the packet belongs.

Because these devices allow you to set up priorities for your network traffic based on applications, you can assign a high priority for your vital in-house applications and use different forwarding rules for low-priority packets, such as generic HTTP-based traffic. Layer 4 switches can also provide security, because company protocols can be confined to only authorized switched ports or users. This feature can be reinforced using traffic filtering and forwarding features.

All these devices can be used to segment your network, but segmentation does not create separate LANs. LANs exist at only the first two layers of the OSI reference model. There's another way to segment your network into separate LANs: use a router.

Routers

Routers are Layer 3 devices that forward data depending on the network address, not the MAC address. Since we are dealing with TCP/IP here, this means they use the IP address. Routers read the header information from each packet and determine the most efficient

route by which to send that packet on its way. Think of the router as providing the link between the various networks that make up the Internet, or any other network that consists of multiple subnets. Routers isolate each LAN into separate subnets.

Like bridges, routers control bandwidth by keeping data out of subnets where it doesn't belong. Routers, however, need to be set up before they can be used. Once they are set up, they can communicate with other routers and learn the topology of the network.

Windows Server 2003 As a Router

So, can Windows Server 2003 be used to provide routing services within your network? The answer is yes. Any computer running a member of the Windows Server 2003 family can act as a dynamic router supporting RIP, OSPF, or both. To have Windows Server 2003 provide routing services, you install multiple network interface adapters, and then enable and configure RRAS. Each network interface adapter is assigned its own IP address and subnet mask to define the directly attached network ID routes. Because you will probably use dynamic routing, default routes won't be used, so you do not need to configure a default gateway for either network adapter.

Static IP routing will be enabled by default when the RRAS is enabled. Your next step should be to use the Routing and Remote Access administration tool to install RIP for IP or OSPF routing protocols. Next, enable the protocols on your installed network adapters by adding them to the appropriate routing protocol.

But we're getting ahead of ourselves. Let's start by building a checklist to follow when setting up Windows Server 2003 as a router:

- Install and configure any necessary network adapters.
- Install RRAS.
- Configure RIP or OSPF.
- Configure the remote access devices.
- Install and configure the DHCP Relay Agent.
- Install a WINS or DNS name server.

Because you're setting up this Windows Server 2003 machine as a router, you'll need to install two network adapters in it. You'll also need to make sure that the necessary drivers are installed, that the TCP/IP protocol is installed, and that IP addresses have been configured on both of the network adapters. Table 4.2 shows how you might set up the IP addresses for this router.

Table 4.2 Typical Network Adapter Setup

Network Card	Connected to	IP Address
1	Backbone	192.168.0.1
2	Subnet	192.168.1.1

Your next step will be to enable RRAS on your Windows Server 2003 machine. The following exercise will walk you through this process.


EXERCISE 4.01

CONFIGURING WINDOWS SERVER 2003 AS A STATIC ROUTER

Configuring a Windows Server 2003 as a static router is simple. To follow these steps, you'll need to be a member of the Administrators group. For security, you may want to consider using the Run As command rather than logging in with Administrator credentials.

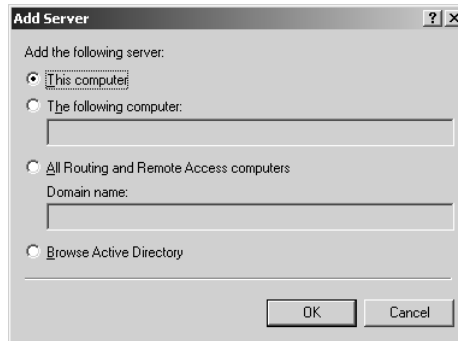
1. If this server is a member of an Active Directory (AD) domain and you're not a domain administrator, you'll need to get your domain administrator to add the computer account of this server to the RAS and IAS Servers security group in the domain that this server is a member of. There's two ways this can be accomplished.
 - Add the computer account to the **RAS and IAS Servers** security group using Active Directory Users and Computers.
 - Use the **netsh ras add registeredserver** command.
2. Select **Start | Administrative Tools | Routing and Remote Access**. The Welcome window appears, as shown in Figure 4.23.

Figure 4.23 Routing and Remote Access Welcome



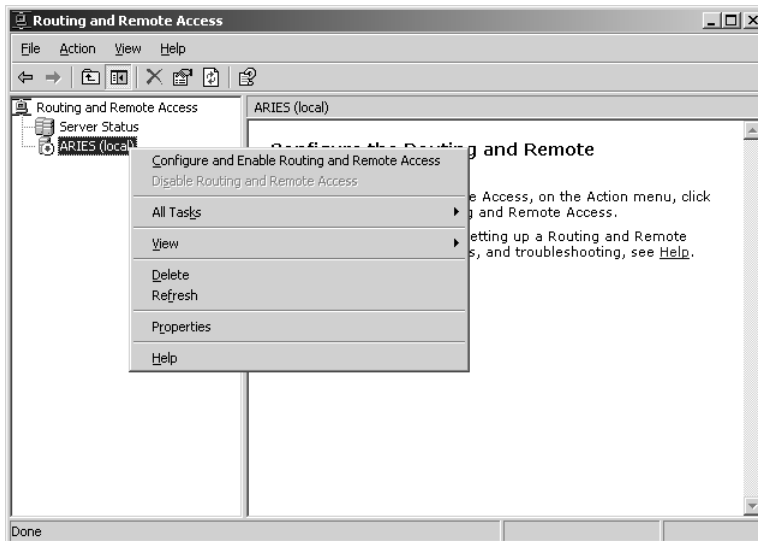
3. The default is that the local computer will be listed as a server. If you want to add another server, right-click **Server Status** in the console tree on the left, and then click **Add Server**.
4. Click the appropriate option in the **Add Server** dialog box, as shown in Figure 4.24, and then click **OK**.

Figure 4.24 Add a Server



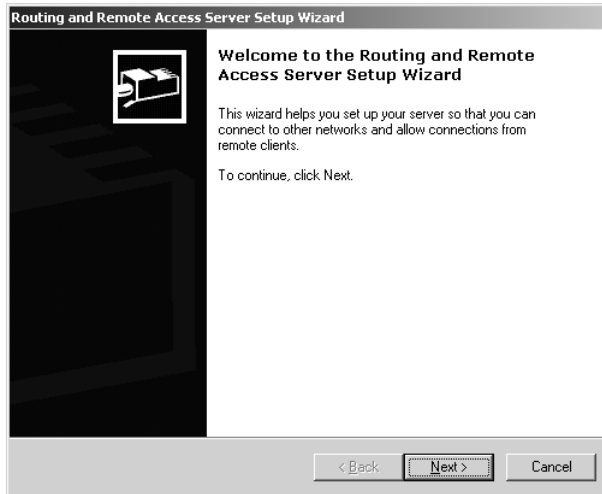
5. In the console tree on the left side of the **Routing and Remote Access** window, right-click the server you want to enable, as shown in Figure 4.25, and then click **Configure and Enable Routing and Remote Access**.

Figure 4.25 Click Configure and Enable Routing and Remote Access




- You've now started the Routing and Remote Access Server Setup Wizard, as shown in Figure 4.26. Click the **Next** button.

Figure 4.26 The RRAS Setup Wizard



- In the next window, choose the **Custom configuration** option, as shown in Figure 4.27. Then click the **Next** button.

Figure 4.27 Choose Custom Configuration




- In the **Custom Configuration** window, choose **LAN routing**, as shown in Figure 4.28, and click the **Next** button.

Figure 4.28 Choose the LAN Routing Option




9. A summary of your selections will now be presented, as shown in Figure 4.29. Verify that the selections you made are correct, and then click the **Finish** button.


Figure 4.29 Finish the RRAS Setup Wizard




10. A dialog box will appear, telling you that the Routing and Remote Access Service has been installed and asking you if you want to start the service, as shown in Figure 4.30. Click **Yes**.

Figure 4.30 Start the Routing and Remote Access Service

11. You should still have the **Routing and Remote Access** window open, and it should now look something like Figure 4.31. To add a static default route to the server, right-click **Static Routes** and then click **New Static Route**.

Figure 4.31 Routing and Remote Access Window after RRAS Installation

12. Choose the interface you want to use for the default route, as shown in Figure 4.32. In the **Destination** text box, type **0.0.0.0**. Do the same in the **Network mask** text box.

Figure 4.32 Choose Your Interface


13. If this is a demand-dial interface, the Gateway text box will be unavailable. Select the **Use this route to initiate demand-dial connections** check box. This will initiate a demand-dial connection when any traffic matching this route occurs.
14. If this interface is an Ethernet or Token Ring LAN connection, in the **Gateway** text box, type the IP address of the interface that is on the same network segment as the LAN interface.
15. In the **Metric** box, type 1. Then click **OK**. You've now added a default static IP route to your router. Follow the same process (steps 11 through 15) for any other route that you want to add to the router.

After you've enabled RRAS, you can also add a static IP route from the command prompt using the **route add** command, which has the following form:

```
route add destination mask subnet-mask gateway metric costmetric if interface
```

Where:

- **Destination** Specifies either an IP address or host name for the network or the host.
- **Subnet-mask** Specifies the subnet mask that is to be associated with this route entry. This entry defaults to 255.255.255.255.
- **Gateway** Specifies either an IP address or host name for the gateway or router to use when forwarding.
- **Costmetric** Assigns a metric cost ranging from 1 to 9,999 to use in calculating the fastest, most reliable route. This defaults to 1.

- **Interface** Specifies the interface you want used for the route. If you don't specify the interface, it will be determined from the gateway IP address.

For example, to add a static route to the 192.168.1.0 network that uses a subnet mask of 255.255.255.0, a gateway of 192.168.0.1, and a cost metric of 2, type this command at the command prompt:

```
route add 192.168.1.0 mask 255.255.255.0 192.168.0.1 metric 2
```

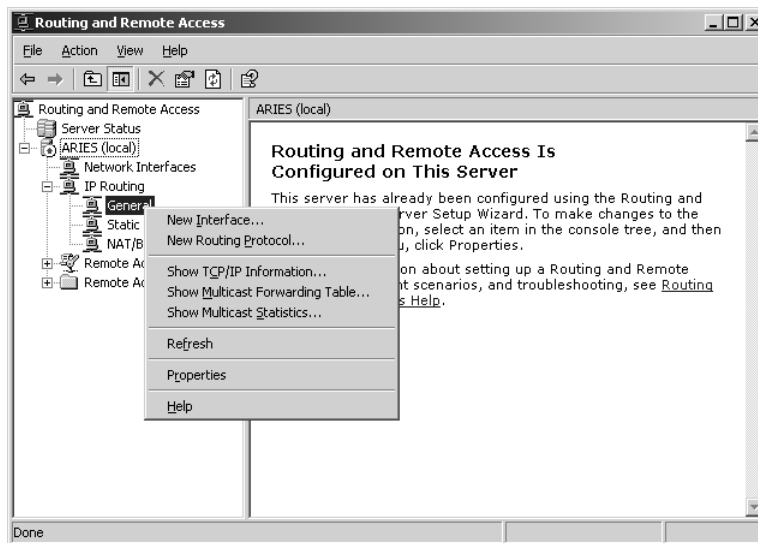
EXERCISE 4.02

CONFIGURING RIP VERSION 2

After you have enabled RRAS and configured a default static route, you need to enable and configure RIP on your router. This is an easy process using the Routing and Remote Access console. Follow these steps:

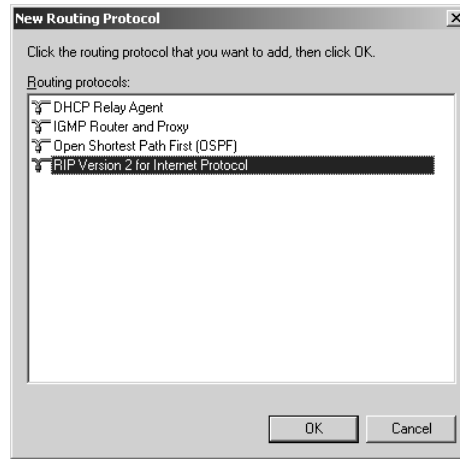
1. Open the **Routing and Remote Access** window.
2. In the console tree on the left side of the window, right-click **General**, and then select **New Routing Protocol**, as shown in Figure 4.33.

Figure 4.33 Add a New Routing Protocol



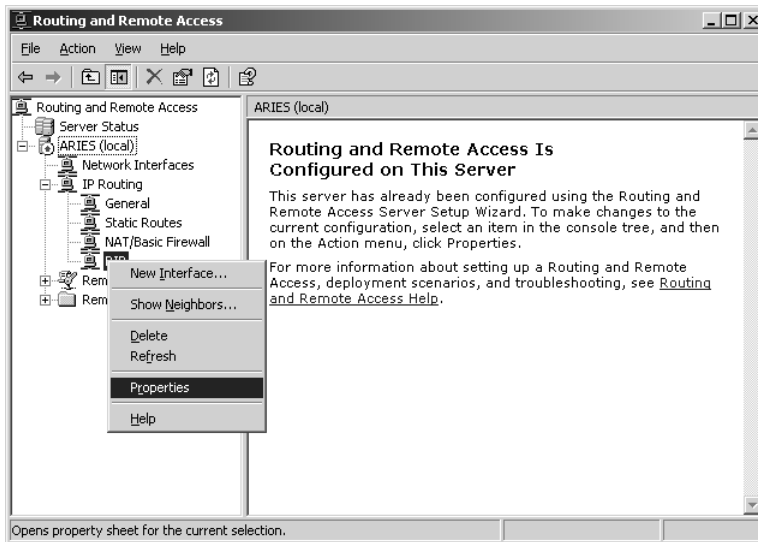
3. From the **New Routing Protocol** dialog box, choose **RIP Version 2 for Internet Protocol**, as shown in Figure 4.34, and then click the **OK** button.

Figure 4.34 Choose RIP Version 2 for Internet Protocol



4. **RIP** now appears under your server and IP Routing. Right-click **RIP** and choose **Properties** from the context menu, as shown in Figure 4.35.

Figure 4.35 Choose RIP Properties



5. On the **General** tab of the **RIP Properties** dialog box, shown in Figure 4.36, you can set the maximum amount of time you want this router to wait before it sends out triggered updates, as well as the level of logging you wish to have performed. Remember that triggered updates

occur when the network topology changes. Updated routing information is sent out immediately reflecting that change. The **General** tab of the **RIP Properties** dialog box lets you set an interval that these triggered updates will wait before being sent. The default is five seconds. There are four levels of logging you can choose from:


- **Log errors only**
- **Log errors and warnings (the default)**
- **Log the maximum amount of information**
- **Disable event logging**



NOTE

Keep in mind that logging consumes system resources so use it sparingly when you are not having network problems. When you are having a problem and you are in the process of identifying and correcting the problem, you'll want to use the **Log the maximum amount of information** option, but after the problem is cleared, immediately reset logging to the default level.

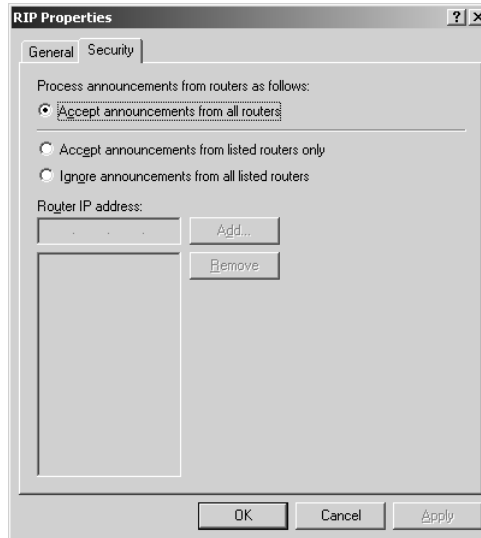
Figure 4.36 The General Tab of the RIP Properties



6. Choose the **Security** tab, shown in Figure 4.37. On this tab, you can designate if this router will process announcements from routers. You can accept all announcements from all routers; you can accept announcements from the listed routers only; or you can ignore announcements from those routers listed.

7. After you've made your choice, click **OK**.

Figure 4.37 The Security Tab of the RIP Properties



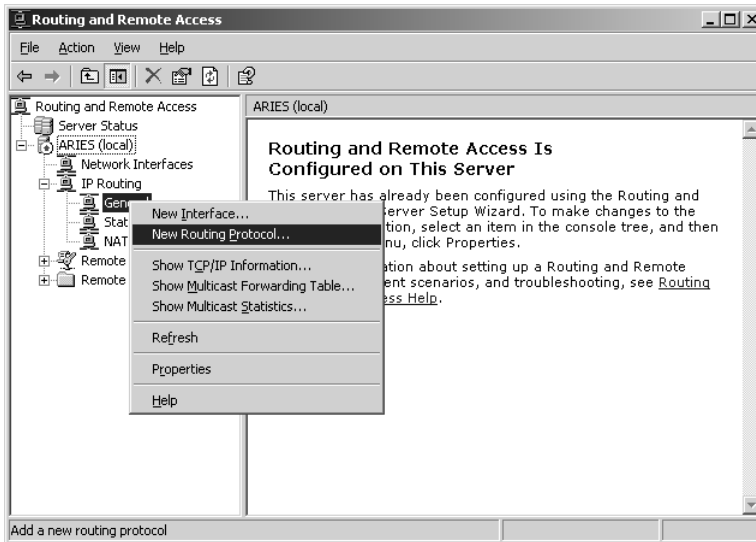
EXERCISE 4.03

CONFIGURING OSPF

You can also configure your RRAS for OSPF. Again, using the Routing and Remote Access console to configure this protocol is easy.

1. Open the **Routing and Remote Access** window.
2. In the console tree on the left side of the window, right-click **General**, as shown in Figure 4.38, and then click **New Routing Protocol**.

Figure 4.38 Add a New Routing Protocol



3. In the **New Routing Protocol** dialog box, choose the **Open Shortest Path First (OSPF)** option, as shown in Figure 4.39, and then click the **OK** button.

Figure 4.39 Choose Open Shortest Path First (OSPF)



4. As with RIP, this action has now added **OSPF** under your server and IP Routing. Right-click **OSPF** and choose **Properties**.

5. You're offered similar choices to those that are available when you configure RIP (see Exercise 4.2). After you've made your choices, click **OK**.

EXAM
70-293

OBJECTIVE
2
2.1.2
3
3.1
5.3.1

Security Considerations for Routing

Keep in mind that IPv4 has no default security mechanism. Unless you take security into consideration, your network will be susceptible to unauthorized monitoring and access. To prevent this, develop a strategy for your IP deployment. The following are two methods that you can use to help you enhance security when deploying IP:

- **Secure your IP packets** End-to-end security requires that you not use address translation (NAT). Internet Protocol Security (IPSec) is the most efficient method of providing for a secure data stream.
- **Set up a perimeter network** Use perimeter networks to help secure your internal network.

Let's talk first about using IPSec to secure your data stream. The Windows Server 2003 IPSec protocol provides end-to-end security of your data stream using encryption, digital signatures, and hashing algorithms. IPSec resides at the Transport layer of the OSI reference model and protects the individual packets before they reach your network, removing the protection on receipt. Even data passed through from applications not having any security features can be protected using IPSec.

Keep in mind that IPSec protects the actual packets of data, not the link. Because of this, IPSec provides security even on insecure networks, and only the computers actually involved in the communication are even aware of it. IPSec provides a number of security features, including the following:

- Authentication by using digital signatures to identify the sender
- Integrity through the use of hash algorithms ensuring that the data has not been altered
- Privacy through encryption that protects the data from being read
- Anti-replay prevents unauthorized access by an attacker who resends packets
- Nonrepudiation through the use of public-key digital signatures that prove the message's origin
- Dynamic rekeying to allow keys to be generated during communication, so that the different transmissions are protected with different keys
- Key generation using the Diffie-Hellman key agreement algorithm, allowing computers to agree on a key without exposing it

- Configurable key lengths, allowing for export restrictions or highly sensitive transmissions

The way that IPSec works is relatively simple. In order for data to be transmitted and protected between two IPSec-enabled computers, the computers must agree on which keys, mechanisms, and security policies will be used to protect the data. This agreement, or negotiation, produces a security association (SA).


The first SA established between the two computers, called Internet Security Association and Key Management Protocol (ISAKMP), provides the method of key exchange. Using ISAKMP to provide protection, the two computers negotiate the production of a pair of IPSec SAs and keys: one for inbound transmissions and one for outbound transmissions. These SAs include the agreed-upon algorithm for encryption and integrity and the agreed-upon IPSec protocol to use. Two IPSec protocols can be used:

- **Authentication Header (AH)** Provides data authentication, integrity, and anti-replay to IP packets.
- **Encapsulating Security Payload (ESP)** Provides confidentiality, along with data authentication, integrity, and anti-replay to IP packets.

Using the IPSec SAs and keys, the two computers protect the data during transmissions.

The second method that you can use to enhance security is a perimeter network. These are also sometimes called a *demilitarized zone* (DMZ) or a *screened subnet*. This type of network is generally an additional network between the protected network and the unprotected network. These types of networks are usually small LANs connecting border routers with internal routers. Servers that are required to be exposed to the Internet, like your Web server or mail server, can be placed in the DMZ and be protected by a firewall. Then additional firewalls are placed between the DMZ and your network. Figure 4.40 demonstrates how this type of configuration might look.

Figure 4.40 A Perimeter Network or DMZ



Analyzing Requirements for Routing Components

A router is nothing more than a very specialized computer. It's made up of the following elements:

- A central processing unit (CPU)
- Random access memory (RAM)
- Input/output system (BIOS)
- Operating system (OS)
- A motherboard
- Input/output (I/O) ports
- A power supply
- A case to hold all of this

Most of these parts remain hidden, but that's okay because these components are generally extremely reliable. Most of the time, you won't need to worry about them at all. The components that you will have the most interaction with are the operating system and the I/O ports.

As you know, the operating system is the software that controls the various hardware components and makes the computer usable. The router usually has a configuration file that includes the number, location, and type of each I/O port, as well as details about bandwidth, addressing, and security.

The I/O ports are the one component that you will get to know on a personal basis. These ports function like NICs, in that they define the medium and framing mechanisms and provide the appropriate physical interfaces.

Simplifying Network Topology to Provide Fewer Attack Points

Attacks on your network can come in a variety of ways, in both active and passive forms. An active form of attack is launched with the purpose of damaging or destroying your data and/or your network infrastructure. Passive attacks, on the other hand, can be thought of more along the lines of "fishing expeditions." In these situations, the attackers are mostly snooping—just looking around.

One of the best defensive postures against both forms of attacks is to limit the paths to your network an attack can take. You can accomplish this by implementing three simple tactics:

- Minimize the number of network interfaces through which the attack may come
- Minimize the number of routes over which the attack may come

- Minimize the number of routing protocols through which the attack may come

Most router attacks involve the manipulation of the routing table entries so that service to legitimate systems or networks is denied. RIP version 1 and Border Gateway Protocol (BGP) offer no or little authentication, and what little they do offer usually isn't implemented. This offers the perfect target for attackers to alter legitimate routes, often by spoofing their source IP address and creating a denial-of-service (DoS) condition. The easiest remedy is to use whatever tools you have available: if your routing protocol offers authentication, implement it. If it doesn't, consider changing to one that does.

Minimizing the Number of Network Interfaces and Routes

You want to limit the number of network interfaces through which an attacker could gain entrance. Every NIC you have exposed to the Internet is a potential doorway through which someone could enter. The fewer interfaces exposed, the less work for you in preventing someone coming through an open port and wrecking havoc on your network.

Minimizing the number of routes an attacker might take to your network is similar to minimizing the interfaces. You are restricting the paths through which an attack may come.

Minimizing the Number of Routing Protocols

You also want to limit the options of attackers if they do manage to gain access to your network. By reducing the number of routing protocols, you reduce the options available to the attacker.

Demand-dial routing allows you to use impermanent, dial-up WAN lines to exchange data between two networks. It allows for the effective use of these impermanent connection methods, such as analog modems and ISDN, to mimic dedicated Internet connections. Demand-dial routing brings up the connection only when outbound traffic is addressed to an associated link. With a demand-dial connection, you can use additional leased lines to add needed bandwidth at peak use times. However, you should check all the potential costs before you choose this alternative, to avoid any unexpected and unpleasant surprises when the telephone bill arrives.

Demand-dial routing concepts are relatively simple. A link is created when needed, and the connection is dropped when it's no longer needed. There are three basic phases of demand-dial connection setup:

- Configure the first router to initiate and receive demand-dial connections from the second router.
- Configure the second router to initiate and receive demand-dial connections from the first router.
- Initiate the demand-dial connection from the first router to the second router.

Adding a Demand-Dial Interface

Although the three phases for setting up a demand-dial connection are relatively straightforward, actually setting up a demand-dial interface can be a complex and lengthy process. Make sure you double-check your work as you go along, because troubleshooting at a later phase could be extremely difficult and complicated. To set up the interface, follow these instructions:

1. Select **Start | Administrative Tools | Routing and Remote Access**.
2. In the console tree on the left side of the window, click the appropriate server or router.
3. Right-click **Network Interfaces**, as shown in Figure 4.41, and choose **New Demand-dial Interface** from the context menu.
4. The Demand-Dial Interface Wizard starts. Click **Next** in the Wizard's first window.
5. The next window asks for a name for this demand-dial interface. The default name is **Remote Router**, as shown in Figure 4.42. You might want to use a more descriptive name, such as the name of the branch office or the name of the network to which you are connecting. When you've named the interface, click **Next** again.
6. You're now confronted with three choices of connection type. For our purposes of adding a demand-dial interface, the first two choices are the only ones we will deal with. If your computer doesn't have one of these, that specific option will be grayed out and unavailable.
 - **Connect using a modem, ISDN adapter, or other physical device**
Choose this option, and then click the **Next** button. Choose which modem you want to use, and then enter the telephone number you want to be dialed. Notice that in addition to the primary number, you can also click **Alternatives** and enter other numbers to be tried automatically if the primary number cannot be reached.
 - **Connect using virtual private networking (VPN)** If you select this option and click **Next**, the **VPN Type** window opens. Choose the tunneling protocol you want to use, and click **Next** again. Finally, in the **Destination Address** window, provide either the host name or the IP address for the remote router and click **Next** again.
7. Under **Protocols And Security**, choose all the conditions that will apply to the connections. If you have chosen to connect using a modem, ISDN device, or other physical device, you will have two options here. This second option will not be available if you have chosen the VPN option earlier. If you choose both options, the wizard will present you with a window to configure each of the items.

Continued

- Add a User Account So A Remote Router Can Dial In
 - Use Scripting to Complete The Connection With The Remote Router
8. In the next window, fill in the IP address of the network or networks you want to access.
 9. The next window asks you to provide the user account and password as your **Dial Out Credentials**. This will complete the Wizard, and a new routing interface will be added in the Routing and Remote Access window.

Figure 4.41 Choose New Demand-dial Interface


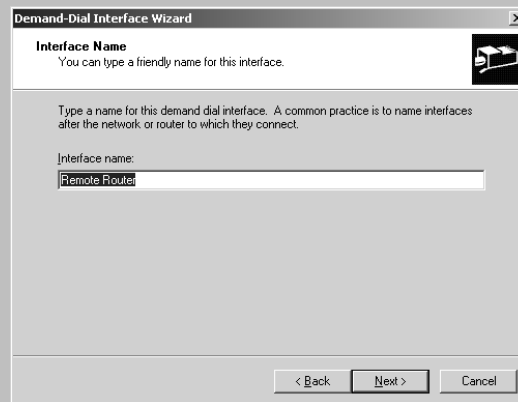


Figure 4.42 Choose an Appropriate Interface Name



Router-to-Router VPNs

Take two separate networks and put the Internet between them. Now, connect them using a tunnel through the Internet. You create this tunnel using the Point-to-Point Tunneling Protocol (PPTP) or Layer Two Tunneling Protocol (L2TP), so that the data being exchanged between the two networks is encrypted. But what's the difference between a normal client VPN connection and this type of VPN? What can you use to connect the two networks together? You can use routers.

You can use a router-to-router VPN to connect two separate networks together over the Internet and still maintain security. Before we get into the specifics of setting up a router-to-router VPN, let's look briefly at how to set up a client VPN connection first. That way, you will understand the difference between the two and why you might want to use one over the other. The first step is to turn on the Windows Server 2003 VPN Server.


EXERCISE 4.04

INSTALLING AND ENABLING WINDOWS SERVER 2003 VPN SERVER

Installing and setting up a Windows Server 2003 VPN Server is simple. Just follow these steps:

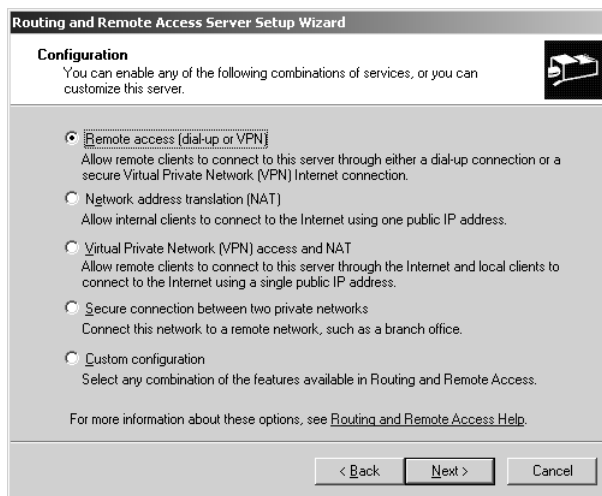
1. Select **Start | Administrative Tools | Routing and Remote Access**. If you have not set up RRAS, you'll see a red circle in the server icon. If you have set up your server to be a VPN server when you were installing the Windows Server 2003 software, you will see a green arrow, as shown in Figure 4.43.

Figure 4.43 RRAS Has Already Been Turned On



2. If the service has already been turned on, you may want to reconfigure your server. You can reconfigure it by right-clicking the server icon and choosing **Disable Routing and Remote Access**. Click **Yes** to continue when you are prompted. Your server icon should now have the red circle rather than the green arrow.
3. Right-click your server's icon and choose **Configure and Enable Routing and Remote Access** to start the Setup Wizard. Click **Next** to continue.
4. Select the **Remote Access (dial-up or VPN)** option, as shown in Figure 4.44, and then click the **Next** button.

Figure 4.44 Choose Remote Access



5. Check the **VPN** check box, and then click the **Next** button.
6. In the **VPN Connection** window, shown in Figure 4.45, select the network interface that is connected to the Internet, and then click the **Next** button.

Figure 4.45 Choose the Interface Connected to the Internet

7. In the **IP Address Assignment window**, you have two choices:
 - **Automatically** Choose this option if you have a DHCP server you can use to automatically assign IP addresses to the remote clients. This setup will be easier to administer than assigning addresses manually. (However, if you do not have a DHCP server, you must specify a range of static addresses.) Click **Next** to continue.
 - **From a specified range of addresses** Choose the option if the remote clients can only be given an address from a specified pool of addresses. Click **Next** to continue. In the **Address Range Assignment** window, click the **New** button. In the **Start IP address** box, type the first IP address in the range of addresses you want to use. Then type in the last IP address in the range you've chosen. Windows Server 2003 will automatically calculate the number of addresses for you. Click the **OK** button to return to the **Address Range Assignment** window, and then click the **Next** button to continue.
8. In the next window, accept the default value of **No, use Routing and Remote Access to authenticate connection requests**, and click the **Next** button to continue.
9. Click **Finish** to turn on RRAS and to configure the server as a remote-access server.

Once you have your server set up to provide VPN service (completed Exercise 4.04), you can allow client machines to connect to it over the Internet.

Configuring a VPN Connection from a Client Computer

To configure a VPN connection from a client computer, you must first be logged on as the Administrator or as a member of the Administrators group. The following steps will vary depending on which version of Windows the client computer has installed.

1. Make sure that you have a correctly configured Internet connection on the client computer.
2. Select **Start | Control Panel | Network Connections | Create a New Connection**. This opens the New Connection Wizard. Click the **Next** button to continue.
3. Click the **Connect To The Network At My Workplace** option, and then click the **Next** button.
4. Choose **Virtual Private Network Connection**, and then click **Next**.
5. Type in a description name in the **Company Name** text box and click **Next**.
6. Choose **Do Not Dial The Initial Connection**. If the computer isn't always connected to the Internet, you should probably choose **Automatically Dial This Initial Connection**, click the name of the connection to the ISP, and click **Next**.
7. Type in the IP address or the host name of the VPN server computer to which you are connecting.
8. Depending on if you want anyone to be able to have access to this VPN connection of just yourself, choose **Anyone's Use** or **My Use Only**, and then click **Next**.
9. Click the **Finish** button and save the connection information.
10. Choose **Start | Control Panel | Network Connections** again and double-click the new connection you just created.
11. Go to **Properties** and configure the options for this connection you want. If you're connecting to a domain, click the **Options** tab and select the **Include Windows Logon Domain** check box, so you can specify that you want to request Windows Server 2003 logon domain information before trying to connect. Another option you'll probably want to select is the **Redial If Line Is Dropped** check box on the **Options** tab.

Using your new VPN connection is simple: click **Start | Connect To** and choose your new connection. If you don't already have a current connection to the Internet, you'll be offered the opportunity to connect. When the connection is made, the VPN server will prompt you for your name and password. Enter the necessary information and click the **Connect** button. All of the same resources available when you are directly connected to the network are available now. When you're ready to disconnect, simply right-click the connection and choose **Disconnect**.

Now that you know how to create and use a client VPN connection, what are the differences in setting up a router-to-router VPN? There are actually not very many differences.

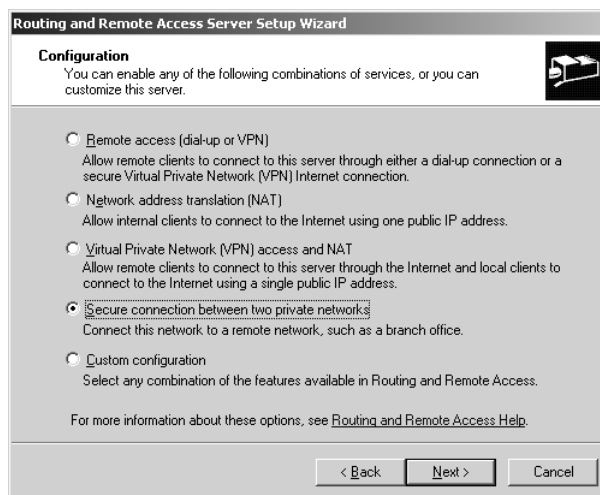
EXERCISE 4.05

SETTING UP WINDOWS SERVER 2003 AS A ROUTER-TO-ROUTER VPN SERVER

The differences in the setup of Windows Server 2003 as a router-to-router VPN server and as a static router (Exercise 4.1) are minimal. Follow these steps:

1. Select **Start | Administrative Tools | Routing and Remote Access**.
2. Right-click your server's icon and choose **Configure and Enable Routing and Remote Access** to start the Setup Wizard. Click **Next** to continue.
3. Select the **Secure connection between two private networks** option, as shown in Figure 4.46, and then click the **Next** button.

Figure 4.46 Choose Secure Connection between Two Private Networks



4. Choose the **No** option when you are asked if you want to use demand-dial connections, unless you need to use them, and then click the **Next** button again. If you choose **Yes** to use demand-dial connections, you'll have the opportunity to set up the demand-dial connections when this Wizard is finished. If you are using a full-time connection, you don't need the demand-dial connection.
 5. Click **Finish** to turn on RRAS and to configure the server as a router-to-router VPN server.
-

Make sure you have addresses assigned to all the installed interfaces and that you've installed and set up your routing protocols on each interface. Then you should be able to use this router.

Packet Filtering and Firewalls

One of the best features available in RRAS is the ability to filter TCP/IP packets traveling in either direction. For all practical purposes, enabling packet filtering creates a firewall on your server. You can build filters that can either allow or deny packet traffic into or out of your network. You do this by specifying rules that designate source and destination addresses and ports.

Normally, you set up these filters to block information that the machines in your network should not receive. The filters are set up on a specific interface. This means that the filters on one interface are completely independent of the filters on another. Incoming and outgoing filters are independent of one another also.

Simply put, you have two choices with input filters: accept all traffic over the interface except the traffic you specify, or drop all traffic except the traffic you specify. Output filters are configured in the same manner. Which choice you should make most often depends on the context and purpose of the filter. The second option is the most secure. If you are attempting to keep all but very specific traffic out of your network, this would be the correct choice. The first choice is appropriate if you are just trying to stop specific traffic.

For instance, say you have a Web server and the only traffic you want to allow on this server is traffic traveling to and from the Web server service. All you need to do is configure an input filter for the destination IP address of the Web server and the TCP destination port 80. At the same time, you will want to configure an output filter for the source IP address of the Web server and the TCP source port 80. If these two filters are the only two filters operational on this server, the only traffic that will be allowed across the interface is TCP traffic to and from the Web server service on your Windows Server 2003 machine.

You need to be careful about how you implement these filters, so that you don't make them too restrictive, which would impair the functionality of the other protocols operating on the server. For instance, given our example of a Web server, we can't use PING or any

other basic IP troubleshooting tool on that computer now, because we've restricted it to only Web traffic on port 80. We'll talk more about troubleshooting shortly.



TEST DAY TIP

Know how to set up both inbound and outbound TCP/IP packet filters. Understand that you can accept all but those IP addresses you want to reject, or you can deny all except those IP addresses you wish to accept.

It's a good idea to use packet filtering to block unwanted traffic from your VPN servers. There are two basic sets of rules for this process: PPTP packet filters and L2TP packet filters.

For PPTP, there are at least two filters that are required to block non-PPTP traffic. You need to allow Generic Routing Encapsulation (GRE) packets to pass. You also need to allow inbound traffic on TCP port 1723. If the PPTP server is also acting as a PPTP client, you can add a third filter to allow outbound traffic on TCP port 1723 also. After these packets are established, choose the **Drop All Packets Except Those That Meet The Criteria Below** radio button. Then close the dialog box. Repeat the process on the output side.

For L2TP packet filters, you will need four filters: two for input and two for output, as follows:

- A filter with the VPN interface address and a network mask of 255.255.255.255, filtering the User Datagram Protocol (UDP) with a source and destination port of 500
- An input filter with a destination of the VPN address and a network mask of 255.255.255.255, filtering UDP traffic with a source and destination port of 1701
- An output filter with a source of the VPN interface address and a network mask of 255.255.255.255, filtering UDP traffic with a source destination of 500
- An output filter with a source of the VPN interface address and a network mask of 255.255.255.255 filtering UDP with a source and destination port of 1701



TEST DAY TIP

Make sure you know how to filter all packets except VPN traffic on a PPTP or L2TP server. Make sure you understand the process and the number of filters each protocol requires.

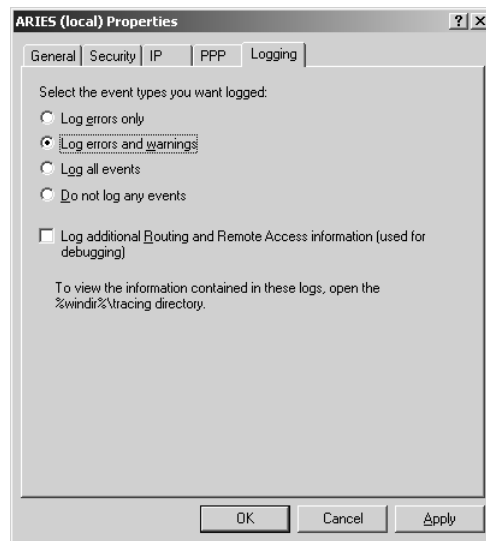
Logging Level

Coming up with a good logging strategy is important for the proper maintenance of your network and the devices that are used on it. Deciding what to log is probably one of the

most important questions you will consider. If you have too much logging, the performance of your server and the network will decline sharply. If you have too little logging, when you have a problem, you won't have the information you need to determine the source and cause. The best choice is to log only those options you really need, and when you don't need a particular type of log data anymore, stop recording it.

In order to set the logging levels, open the RRAS module, right-click the server you wish to administer, choose **Properties**, and then click the **Logging** tab. As shown in Figure 4.47, the **Logging** tab contains several options for the various types of events that you can log. The default is to log all errors and warnings. You can also check the **Log additional Routing and Remote Access information (used for debugging)** check box, which, as its name implies, will assist you in debugging.

Figure 4.47 Set the Logging Level



EXAM
70-293

OBJECTIVE
2
2.1.2
3
3.4

Troubleshooting IP Routing

Here, we will look at the two main tools you might use in troubleshooting IP routing and the common problems that occur with IP routing, which you will be expected to know how to deal with in the exam.



NOTE

There are entire books devoted to troubleshooting IP routing. Also, Microsoft's online help system is fairly good at suggesting probable causes and solutions for many common routing problems.

Identifying Troubleshooting Tools

Your best troubleshooting tools are those tools you should be using on a daily basis for network management and monitoring. Windows Server 2003 ships with the Network Monitor tool (NETMON.exe), which is an excellent protocol analyzer that you can use to monitor your network. As discussed in Chapter 3, this tool captures and displays information about the IP packets moving in your network and can tell you about traffic patterns, broadcast rates, how the network is being used, what kinds of errors you might be experiencing, and many other aspects concerning the behavior of your network.

The Routing and Remote Access console is another excellent troubleshooting tool. Using this tool, you can show your network's TCP/IP information, your IP routing table, the router's RIP neighbors, its OSPF area, the LSDB, the router's OSPF neighbors, and the OSPF virtual interface.

Other familiar tools that you can use for troubleshooting include PING, pathping, tracert, mrinfo, and netsh. Let's take a look at how you can use these tools to verify and troubleshoot your connections.

Testing your TCP/IP Connections with PING

To use PING to test your TCP/IP connections, follow these steps:

1. Click **Start | Run**, type **cmd**, and press the **Enter** key to bring up the command prompt.
2. Using the **ipconfig** command, discussed in Chapter 3, determine the IP addresses of your computer and your default gateway.
3. Making sure that TCP/IP is installed and working on your local computer. Then type **ping 127.0.0.1** at the command prompt and press the **Enter** key. You should receive a response in the command prompt window displaying four replies from the 127.0.0.1 loopback address. If not, you will need to reset the TCP/IP configuration on your machine.
4. If you received the proper replies, test the IP address of your local machine that you obtained from the **ipconfig** command by pinging it. If you receive the correct four replies, you know that your computer was added to the network correctly.
5. Ping the default gateway address to verify that it is up and running. This also lets you know if you are able to connect to a local host on your local network.
6. Ping the IP address or hostname of another remote host. You can ping a hostname by typing **ping www.microsoft.com**. This will let you know that you are able to communicate through a router.

Another useful troubleshooting tool is the pathping command. This command combines aspects of PING and tracer, and adds in some additional features that make it an excellent troubleshooting tool. This tool works by measuring the packet loss across each router between the source machine and the destination. This information can help you determine where your network reliability problems may be coming from. The syntax for the pathping command is as follows:

```
pathping [-n] [-h maximum_hops value] [-g host-list] [-p value]
        [-q value] [-w value] final_destination
```

Where:

- **-n** Tells pathping not to resolve addresses to host names.
- **-h maximum_hops value** Sets the maximum number of hops you want the command to search for the target. The default is 30 hops.
- **-g host-list** Provides a loose source route along the host list.
- **-p period** Sets the wait period in milliseconds between pings. The default is 250 milliseconds.
- **-q num_queries** Sets the number of queries per hop. The default is 100 queries.
- **-w timeout** Sets the time length in milliseconds for each reply before the command times out on that hop. The default is 3000 milliseconds.
- **-T** Tests the connectivity to each hop with Layer-2 priority tags.
- **-R** Tests to see if each hop is RSVP-aware.
- **final_destination** The host name or IP address of the network, domain, or machine that you are testing the route to.

The tool will first trace the route to the destination, and then analyze the traffic running through each hop. Keep in mind that one test is not sufficient to give you a good idea about what is going on. There is no specific number of lost packets that signify that a link is causing you problems. If the number is in double digits, though, you should probably examine that route carefully. To get a realistic picture of what is going on in your network, test a router over time and test in both peak and off-peak usage.

If you're using multicast routing, another useful troubleshooting command is mrrinfo. This command displays multicast router configuration information. The syntax is as follows:

```
mrrinfo [-n] [-?] [-i address] [-t secs] [-r retries] destination
```

Where:

- **-n** Displays the IP addresses in numeric format.
- **-?** Prints usage information.

Using Tracert to Test TCP/IP Connections

You can also use Tracert to test your TCP/IP connections. Just follow these steps:

1. Click **Start | Run**, type **cmd**, and press the **Enter** key to bring up the command prompt window.
2. At the prompt, type **tracert target-ipaddress** and press the **Enter** key. Replace **target-ipaddress** with the IP address of the remote network host you are attempting to connect with. This can also be a host name.

The display will now include a list of the routers the packets have successfully crossed, along with the length of time the packet took to reach that network segment.

- **-i** Specifies the IP address of the local interface from which the query was sent.
- **-r** Specifies how many times an SNMP query is to be resent. The default value is 0.
- **-t** Specifies how long to wait for an IGMP neighbor query reply. The default is three seconds.

The `mrinfo` command displays the interfaces for both the multicast router and its neighbors on each interface. It also provides the names of the neighboring domains, the multicast routing metric, and the TTL.

Also, the `netsh` utility, discussed in the “Using netsh Commands” section earlier in this chapter, can display the configurations of protocols, filters, and routes. It also allows you to reconfigure interfaces. Don’t overlook this valuable tool as an option for troubleshooting IP routing.

Common Routing Problems

If you suspect that your RRAS server isn't functioning properly, start by making sure the RRAS server is running. You might be surprised how many times the cause of the problem turns out in fact to be that RRAS is not turned on.

Most TCP/IP administrators spend much of their time troubleshooting the hardware. Connectors go bad, NICs die, and cables break or are cut. You need to troubleshoot and repair these elements before you start looking at the software. Consider these potential trouble spots first:

- Check for basic communication between systems first. Broken cables, loose connections, and so on can cause what might look like much more complex problems.
- Make sure that your systems are in compliance with the standards you've chosen. This means you need to verify all devices on your Ethernet are broadcasting Ethernet and not something else. Make sure you have the correct types of cables. An example of this is the common mistake beginners sometimes make using RG59A/U cable instead of RG58A/U. The former cable type is used in broadcasting specifically with video; the latter is used with IEEE 802.3 10Base2 networks.
- Carefully isolate your problem to a single LAN, MAN, or WAN segment by going through each individually. Keep in mind it is extremely rare for two segments to go down at the same time.

Interface Configuration Problems

Make sure that the RRAS server is configured to perform as an IP router. Open the RRAS Microsoft Management Console (MMC) and verify all your settings. Make sure that you have enabled RRAS on the Windows Server 2003 machine you are expecting to perform as a router. It could be that you have the wrong server configured. Also, keep in mind that the system must first make the physical connection to the network. After that, it must make the logical connections.

The router also might not be receiving routed data from other routers. Take a look at the routing table to see that the router is receiving routes from the other routers. If there is anything there other than **Local** in the **Protocol** column, the router is receiving routes via the routing protocols. If not, double-click the rest of the settings in this section and pay particular attention to the appropriate protocol.

RRAS Configuration Problems

Routing for the correct LAN protocol may not be enabled. If you're using IP routing, make sure that IP routing is enabled on the IP tab of the server's property sheet. Also, make

sure that you have IP routing protocols attached to each of the interfaces where they are needed.

The wrong protocol could be installed, or the right protocol could have been installed on the wrong interface. The correct protocol must be installed on the appropriate interface for this to work correctly.

Routing Protocol Problems

One of the most common problems you'll face with RIP for IP is incorrect routing table entries. If you're seeing wrong or inconsistent routes in the routing tables, or if routes are totally missing, you should look at the following possibilities:

- The wrong version of RIP could be in use.
- Silent RIP hosts might not be receiving updates.
- The subnetting scheme on your network could be incompatible with your routing infrastructure.
- A router might be using the wrong password.
- Routing filters might be too restrictive.
- Packet filters might be too restrictive.
- Neighbors might be incorrectly configured.
- Default routes might not be being propagated.

If your router is using OSPF, make sure that the **Enable OSPF on this interface** check box is selected. This option is in the interface's **OSPF Properties** dialog box.

Also make sure that your router is receiving routing information from the other routers on the network. Do this by opening the routing table and looking at the **Protocol** column. One of the following might be the problem with OSPF:

- OSPF might not be enabled on the desired interface.
- The neighboring router might be unreachable.
- The OSPF settings may not match on each of the neighboring routers.
- The stub area configuration or area ID on neighboring routers may not match.
- Interfaces may not be configured with OSPF neighbor IP addresses.
- There may not be a designated router (DR) for the network.
- Packet filtering may be too restrictive.
- Summarized routes may be configured improperly.
- ASBR source or route filtering may be too restrictive.
- Virtual links may be incorrectly configured.

If a routing table entry is marked as being either OSPF or RIP, then information from some of the other routers on your network is getting through. If you do not see any OSPF or RIP entries in the table, you have a problem.

EXAM
70-293
OBJECTIVE
2.5.3

TCP/IP Configuration Problems

Verifying that the router's TCP/IP configuration is correct first may save you a lot of time. You must use the correct IP address and subnet mask.

Routing Table Configuration Problems

You'll need to have a static default route defined and enabled so that your router will forward any packets when there is no specific route designated for them. If the default route is incorrect or missing, you will have problems. If you're using default routing, the default route must be learned through the routing protocols or statically configured on the router over the correct interface.



TEST DAY TIP

You will need to know extremely basic problems and their solutions for the exam.

Summary of Exam Objectives

In this chapter, we discussed three main topics: understanding IP routing, security considerations for routing, and troubleshooting IP routing. We've looked at routing basics, including how devices are identified on the network, how NAT works, and the differences between IPv4 and IPv6. We've also looked at creating, viewing, and updating routing tables and the differences between static and dynamic routing.

The chapter continued with a discussion of the various routing protocols and the differences between distance-vector and link-state routing algorithms. This naturally led to a discussion of the two primary examples of both algorithms, RIP and OSPF. This discussion also examined the differences between RIP version 1 and RIP version 2 and what sample networks using these protocols might look like. We also examined OSPF and how sample networks might work using this protocol.

Later in the chapter, we looked at the OSI reference model and routing devices, as well as how to use utilities such as Netsh, PING, pingpath, and Tracert for troubleshooting. The chapter continued by looking at how to use Windows Server 2003 as a router and as a VPN server.

Exam Objectives Fast Track

Understanding IP Routing

- ☑ You must understand the concepts underlying IP addressing in order to understand how IP routing works. Understand the three IP address formats: hexadecimal, binary, and dotted-decimal. Have a firm grasp of how IP addresses are structured and how the network and node information is contained in the various address classes.
- ☑ Know that an IP address is a software address, not a hardware address.
- ☑ Know how to view the routing tables of your servers.
- ☑ Understand the differences between static and dynamic routing. Make sure you are familiar with the various configurations that enable and disable both static and dynamic routing, as well as which protocols are associated with each type of routing.
- ☑ Know the differences between RIP and OSPF. Understand why RIP is best used for smaller networks and OSPF is best for large networks.
- ☑ Know how to use the netsh utility and why.

Security Considerations for Routing

- ☑ Remember that IP has no default security mechanisms. In order to add security, you must add other protocols.
- ☑ Secure your IP packets with IPSec or other encryption protocols, depending on the routing strategy you choose.
- ☑ Set up a perimeter network to defend your inner network.
- ☑ Minimize the number of network interfaces.
- ☑ Minimize the number of routes.
- ☑ Minimize the number of routing protocols.

Troubleshooting IP Routing

- ☑ Understand how to use the available troubleshooting tools, including Network Monitor, the Routing and Remote Access console, and the netsh utility.
- ☑ Know the most common routing problems and their solutions.
- ☑ The pathping command and Tracert are excellent ways to troubleshoot your network.
- ☑ Know how to use logging and where to change the parameters.

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 a metric and how is it used in choosing the best routes?

A: A metric is that value assigned to an IP route for a particular network interface that tells you how much it's going to cost to use that specific route. Metrics can be assigned values based on link speed, number of hops, or even the time delay that might be associated with that particular route.

Q: What are the two principal packet-filtering methods supported by RRAS?

A: The two principal methods used in packet filtering are inbound and outbound filters. You basically accept all inbound packets except those expressly denied, or you deny all inbound packets except those expressly allowed. The same principle works for outbound traffic.

Q: How do I check my TCP/IP configuration in Windows Server 2003?

A: If you are troubleshooting your TCP/IP network, the first thing you want to do is check your TCP/IP configuration on the machine having the problem. You can do this by clicking **Start | Run** and typing **cmd** in the **Run** text box. Now press the **Enter** key. This brings up the command prompt window. At the command prompt, type **ipconfig /all**, and then press the **Enter** key. This command will display a detailed configuration report containing all of the information concerning your network interfaces, including DNS suffix, IP address, subnet mask, and default gateway. Make sure that your computer has all the correct settings for the DNS and WINS servers, a correct and available IP address, a correct subnet mask, a correct default gateway, and the correct host name.

Q: Which tools can I use to test my TCP/IP connections?

A: If you are having problems connecting to a remote server, you'll want to test your connections. There are two common tools that are used for this task: PING and Tracert. The **ping** command is used to verify if a host computer can connect to network resources or not. The **tracert** command is used to examine the route being used from your computer to the destination. The Tracert utility shows the series of IP routers that are used to deliver packets from your computer to the destination and how long it takes

for each hop. If the packets cannot reach their destination, the name of the last router that successfully forwarded the packets is listed. Two other commands that can also be used to test functionality are **route** and **pathping**.

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 IP Routing

1. Your IT Director has decided the new internal network needs to use private addressing. Which of the following IP addresses are private addresses?
 - A. 193.168.0.1
 - B. 171.17.0.1
 - C. 10.0.0.1
 - D. 172.16.0.15
2. Your IT Director has determined that your network should use dynamic routing. You've determined that a route is now being considered unreachable. What has happened to that route in the routing table?
 - A. It has been marked as unreachable in the routing table.
 - B. Nothing has happened to that route in the routing table.
 - C. It has been removed from the routing table.
 - D. You must manually go into the routing table and remove the entry.
3. Your newest hire has been assigned the task of configuring a Windows Server 2003 computer as a router and has asked you how to determine if a machine address or an IP address is being used at the router. You explain that routers use IP addresses, while bridges and hubs use machine addresses. You continue to explain that the OSI reference model has seven layers and that IP, or the Internet Protocol, operates at what layer?

- A. The Physical layer
 - B. The Data Link layer
 - C. The Network layer
 - D. The Transport layer
4. Your IT Director has opened a command prompt window on your Windows Server 2003 computer and is trying to figure out what routes are available to this computer. Which of the following commands should you tell him to use to list the active routes from the command prompt?
- A. route list
 - B. route print
 - C. show route
 - D. dump
5. Your IT Director is determined to use static routing on your large corporate network. You need to convince him that static routing probably is not the best choice, and you want him to think that decision was his idea. You decide to do this by asking him which of the following is an advantage of using static routing?
- A. Fault tolerance
 - B. Scalability
 - C. Manual configuration
 - D. Classless routing
6. RRAS is enabled on your Windows Server 2003 computer, and you have three network adapter cards in the computer configured for subnet IDs of 192.168.32.0/20, 192.168.64.0/20, and 192.168.96.0/20. Which subnet ID can you use if you need to support another subnet with this RRAS server?
- A. 192.168.20.0/20
 - B. 192.168.40.0/20
 - C. 192.168.48.0/20
 - D. 192.168.60.0/20

7. You want to configure a multiple gateway on a Windows Server 2003 machine, but you have only one NIC installed. How do you accomplish this goal?
 - A. Assign the IP addresses 192.168.0.10 and 192.168.1.10 to the interface.
 - B. Assign the IP addresses 10.0.0.1 and 172.16.0.1 to the interface.
 - C. Assign the IP addresses 172.16.0.1 and 192.168.0.1 to the interface.
 - D. You cannot configure multiple gateways on a machine with one NIC.

8. Your IT Director has been reading again. He has decided that he wants to convert the network to OSPF, but he is having some difficulty with terminology. He knows that an OSPF router can serve one of four roles. His problem is that he can't remember which role exists when one of the router's interfaces is on the backbone area. Help him out. Which of the following is it?
 - A. Internal router
 - B. Area border router
 - C. Backbone router
 - D. Autonomous system boundary router

Security Considerations for Routing

9. As the network administrator, you are asked to set up network access so that a group of contract developers can work via a VPN connection connecting to your network's Windows Server 2003 VPN server. The contract developers are all using either Windows 2000 Professional or Windows XP Professional workstations. You must meet the following requirements:
 - The contract developers must be allowed to connect to the network via the Internet.
 - You must use PPP encryption.
 - You must use a protocol that provides tunnel authentication.
 - You must use a protocol that secures the data between the endpoints of the tunnel.

You configure a VPN using PPTP. Which of requirements are met? (Select all that apply.)

- A. The contract developers are able to connect to the network via the Internet.
- B. PPP encryption is used.
- C. Tunnel authentication is used.
- D. Data between the endpoints of the tunnel is secure.

10. You have enabled RRAS on your Windows Server 2003 computer. You want to set up IP packet filtering to help you manage access from remote clients. Where in the Routing and Remote Access console will you enable IP packet filters?
 - A. The properties of the remote access ports
 - B. The properties of the remote access server
 - C. The profile of a Remote Access Policy
 - D. The conditions of a Remote Access Policy

11. You have set up an isolated, secure subnet with only an RRAS server running on Windows Server 2003 connecting the two parts of your internal network. You are protecting your internal network against unauthorized access with your firewall, and authorized users on the intranet establish VPN tunnels to your secure subnet through the RRAS server. You do have a problem, however. It seems that remote VPN clients cannot access the secure subnet through your configuration. How should you reconfigure the system to allow remote VPN clients access to the secure subnet?
 - A. Ask your ISP to create the necessary filters to allow IPSec traffic to pass.
 - B. Create filters on the RRAS server to allow only VPN traffic to pass.
 - C. Define filters on the firewall to allow the VPN traffic to pass.
 - D. Configure the router in front of the firewall to allow IPSec traffic to pass.

12. You've been asked to provide Internet access for clients on your network. You decide to use NAT. You try to establish a secure VPN session from a remote site unsuccessfully. You try again using L2TP. Again the connection fails. You are able to successfully connect when in the same office. Why are you unable to make a connection from the remote location?
 - A. You haven't configured the NAT server to translate the IP Security packets.
 - B. You cannot establish an L2TP connection behind a computer running NAT. The L2TP session fails because the IP Security packets become corrupted.
 - C. L2PT does not work with Windows Server 2003 VPNs.
 - D. NAT does not allow for remote networking.

13. You've just been asked to set up things so that a group of developers can work from home and still connect to your office network. The developers are using either Windows 2000 Professional or Windows XP Professional. You must meet the following requirements:
 - Allow the developers to connect to the network through the Internet.
 - Use PPTP encryption.

- Use a protocol that provides tunnel authentication.
- Use a protocol that secures data between the endpoints of the tunnel.

You plan to configure a VPN that uses L2TP. Which requirement or requirements are met?

- A. The developers can connect to the network through the Internet.
- B. PPTP encryption is used.
- C. Tunnel authentication is provided.
- D. Data between the endpoints of the tunnel is secured.

Troubleshooting IP Routing

14. You've installed RRAS on a Windows Server 2003 computer in your network. The network is not connected directly to the Internet, and the private IP address range you are using is 192.168.0.0. When you dial in, you connect successfully, but you're unable to access any resources. Pinging other servers using their IP addresses results in the message "Request timed out." Running the `ipconfig` command shows you that your dial-up connection is being given the IP address 169.254.75.182. What should you do to resolve the problem?
- A. Configure the remote-access server to act as a DHCP Relay Agent.
 - B. Ensure that the remote-access server is able to connect to a DHCP server that has a scope for its subnet.
 - C. Configure the remote-access server with the address of a DHCP server.
 - D. Authorize the remote-access server to receive multiple addresses from a DHCP server.
15. You think you may have a problem on your network. You need to open a command line window and troubleshoot your network. Which of the following lists of commands represent the command-line utilities most often used in maintaining and testing routing functionality?
- A. `show helpers`, `Trace`, `PING`, `Route`
 - B. `pathping`, `Tracert`, `show helpers`, `show routing`
 - C. `pathping`, `PING`, `Route`, `Tracert`
 - D. `pathping`, `PING`, `Route`, `Trace`

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, D**

2. **C**

3. **C**

4. **B**

5. **D**

6. **C**

7. **A**

8. **C**

9. **A, B, D**

10. **C**

11. **C**

12. **B**

13. **A, C**

14. **B**

15. **C**

MCSE 70-293

Planning, Implementing, and Maintaining an Internet Connectivity Strategy

Exam Objectives in this chapter:

- 2 Planning, Implementing, and Maintaining a Network Infrastructure
 - 2.3 Plan an Internet connectivity strategy
 - 2.5 Troubleshoot connectivity to the Internet.
 - 2.5.1 Diagnose and resolve issues related to Network Address Translation (NAT).

- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

Internet connectivity is no longer a luxury for most businesses; it is a necessity. Employees use the Internet to exchange e-mail with clients, suppliers, and co-workers in other physical locations; to conduct research via the Web; and to remotely access the local area network (LAN) from home or when on the road. Creating an effective policy for implementing and managing the organization's Internet connections is an important part of the Windows Server 2003 network administrator's job.

This chapter is about how to develop the best strategy for connecting your company's Windows Server 2003 network to the Internet. We'll discuss connecting the LAN to the Internet using routed connections or translated connections (via Internet Connection Sharing or the Routing and Remote Access Service's Network Address Translation component). You'll learn how to use both Internet-based virtual private networks (VPNs) and router-to-router VPNs to provide connectivity to the company's LAN from remote locations or to connect two branch offices. We'll discuss the intricacies of demand-dial/on-demand connections and persistent connections, and explain the difference between one-way and two-way initiation. We'll also show you how to use Remote Access Policies to control VPN connections, and we'll discuss VPN protocols supported by Windows Server 2003 and how to make VPN connections using either the Point-to-Point Tunneling Protocol (PPTP) or the Layer 2 Tunneling Protocol (L2TP). You'll learn about VPN security and the authentication and encryption protocols that make your virtual network private.

Next, we'll take a look at the Internet Authentication Service (IAS) and how it can provide centralized user authentication and authorization, centralized auditing and accounting, and extensibility and scalability. You'll learn about IAS integration with Windows Server 2003 Remote Access and Routing Service (RRAS), and how to control authentication via Remote Access Policies. We'll show you how to use the IAS Microsoft Management Console (MMC) snap-in and how to implement monitoring of IAS, and we'll discuss the use of the IAS Software Development Kit (SDK). Then we'll delve a little deeper into the IAS authentication methods and discuss Remote Authentication Dial-In User Service (RADIUS) access server support, wireless access points (WAPs), and authenticating switches.

In the next section, we'll walk you through the process of using the Connection Manager Administration Kit (CMAK) to create service profiles, custom actions, and custom help files, as well as VPN support, to make it easier for nontechnical users to connect remotely without needing to do complex configuration. We'll talk about security issues pertaining to Connection Manager, and show you how to prevent editing of service profile files, how to prevent users from saving their passwords, and how to distribute service profiles securely.



Connecting the LAN to the Internet

You can connect a Windows Server 2003 network to the Internet in two basic ways:

- Using a router to directly route traffic to and from the Internet
- Using a translation service to convert traffic from an internal network to Internet traffic

The following sections discuss the advantages and disadvantages of these methods.

Routed Connections

The traditional method of connecting a network to the Internet is to use a router to route traffic between the external network and your local network. The advantages of this approach are that it is easy to configure, requiring only simple hardware setup, and that it allows full Internet access for all machines on the local network segment. It also allows all machines on the network to provide services to the Internet.

Routed connections have two chief disadvantages. First, every machine on the local network is reachable from anywhere on the Internet. This is rarely necessary and creates a large number of potential security problems. Second, a separate Internet IP address is required for each machine that can access the Internet. Since IP addresses are scarce and are issued only to networks that can prove a need for them, this is not the most efficient approach.

Advantages of Routed Connections

Although translated connections are becoming increasingly popular, routed connections do have a number of advantages:

- Since each client is connected to the Internet through the router, clients can connect even if the local network servers are not working.
- Some Internet clients, such as multimedia applications and games, do not work correctly over a translated connection.
- Each machine has a dedicated Internet IP address and can be used for services such as File Transfer Protocol (FTP) and Domain Name System (DNS) that require a unique IP address per host.

Hardware and Software Routers

A routed connection uses a *router*, a device that transmits data between the internal network and the Internet. There are two types of routers:

- A hardware router is a dedicated device. Hardware routers provide a simple “out-of-the-box” solution for Internet connections.

- A software router runs as a service on one of the computers on the network. The Routing and Remote Access Service (RRAS) in Windows Server 2003 allows a computer to act as a router.

In order to use a computer as a software router, it must have two network connections: one to the internal network (LAN) and one to the external network (the Internet). Microsoft sometimes refers to a computer with two network connections as a *multihomed computer*.

IP Addressing for Routed Connections

When you are using a routed connection to the Internet, each machine on the internal network will need a valid Internet IP address. IP addresses are managed by a central authority, the American Registry for Internet Numbers (ARIN). You will typically obtain IP addresses from an Internet Service Provider (ISP), which has obtained a block of addresses from ARIN for use by its clients.

Once you have been issued one or more IP addresses, you can assign them to the computers in the network. There are two basic ways to accomplish this:

- By manually configuring an IP address in each computer's network connection properties
- By using the Dynamic Host Configuration Protocol (DHCP) to assign addresses

Using DHCP, you can define the IP addresses you have been issued in the DHCP server, and clients are automatically assigned, or *leased*, an address when they are booted. If a client disconnects from the network, its lease is terminated after a timeout period and available to other computers.



TEST DAY TIP

Any Windows Server 2003 (or Windows 2000 Server) computer can act as a DHCP server. To configure DHCP, select **Start | Administrative Tools | Configure Your Server Wizard** and enable the DHCP Server role.


Translated Connections

The second strategy is to use a service that translates between internal IP addresses and external addresses used on the Internet. By using this technique, you can enable Internet access for many computers using a single Internet IP address. Along with conserving address space, address translation ensures that your computers are not accessible directly from the Internet, effectively preventing many types of network attacks.

Network Address Translation (NAT) is an Internet standard defined in RFC 1631 for systems that translate between internal and external network addresses. Windows networks support two types of NAT service:

- Network address translation (NAT) is a full-featured NAT implementation supported by Windows 2000 Server and Windows Server 2003.
- Internet Connection Sharing (ICS) is a simplified NAT implementation for small networks, and is supported by Windows 98 Second Edition, Windows Me, Windows XP, and Windows 2000 Professional.

When you configure the NAT or ICS service, the computer that acts as the NAT server must have at least two network connections: a connection to the Internet (typically a modem or broadband connection) and a connection to the LAN containing the computers that will share the Internet connection.



Network Address Translation (NAT)

NAT is Microsoft's full-featured address translation feature. When you access the Internet on a network that uses a NAT server, outgoing packets are sent to the NAT server, which changes their originating address and forwards them to the Internet. The returned packets are delivered to the NAT server. The server then translates the packets to internal IP addressing and sends them to the machine that made the original request.

The Windows Server 2003 NAT server actually supports three separate services:

- NAT, the address translation service
- DHCP for assigning IP addresses to clients that are sharing the Internet connection
- DNS for name resolution

Depending on your network configuration, you might not need the NAT server to handle address assignment or name resolution. You can choose whether to use these components when you configure the NAT server. If you have dedicated DHCP or DNS servers on the network, you can continue to use them with NAT. (The DNS service forwards requests to an Internet DNS server and returns the results to the appropriate client within the private network.)

Installing the NAT Service

NAT is part of the RRAS component of Windows Server 2003. RRAS is installed with Windows Server 2003 but is not enabled by default. You can enable this service using the Manage Your Server application that is launched when you install the operating system or by using the Routing and Remote Access MMC snap-in. Windows Server 2003 includes a wizard that can enable RRAS and set up a NAT server. Exercise 5.01 shows how to configure NAT using the wizard.



TEST DAY TIP

Remember that you need at least two network interfaces on the NAT server: one connected to the private network, usually a LAN adapter, and one connected to the Internet. You can configure a demand-dial Internet connection (if you're using a modem or ISDN dial-up instead of an "always-on" connection to the Internet) during the NAT server setup process.

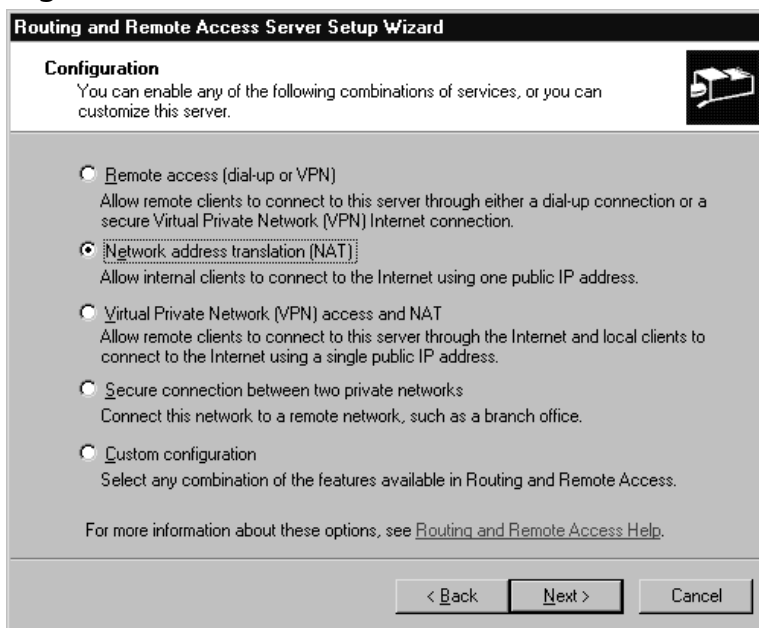
You can also configure NAT manually using the Routing and Remote Access MMC snap-in. This is the only way to configure a NAT server on a machine that already has RRAS enabled. RRAS can perform NAT along with its other functions, which include acting as a network router or accepting dial-up network connections.

EXERCISE 5.01

INSTALLING NAT USING THE WIZARD

You can install NAT on a Windows Server 2003 server that does not yet have RRAS enabled using the Routing and Remote Access Server Setup Wizard. This exercise guides you through the process of setting up a basic NAT server using the Wizard.

1. Select **Start | Administrative Tools | Routing and Remote Access** to start the RRAS MMC snap-in.
2. Click the RRAS server name (usually the current machine) in the left column to highlight it.
3. From the menu, select **Action | Configure and Enable Routing and Remote Access**.
4. The Wizard displays a Welcome window. Click **Next** to continue.
5. The **Configuration** window appears. Select the **Network address translation (NAT)** option, as shown in Figure 5.1, and click **Next**.

Figure 5.1 Select NAT from the RRAS Wizard

6. The **NAT Internet Connection** window is displayed. Here, you can choose how the NAT server will connect to the Internet. Choose either **Use this public interface to connect to the Internet** or **Create a new demand-dial interface to the Internet**.
7. You can optionally choose to enable basic security for the Internet interface by checking the **Enable security on the selected interface by setting up Basic Firewall** option. This option is enabled by default.
8. Click **Next** to continue.
9. The **Ready to Apply Selections** window is displayed. Click **Next** to start the RRAS service.

If you chose to create a new demand-dial interface in Step 6, the **Demand-Dial Interface Wizard** will guide you through this process. This Wizard is described in Exercise 5.04, later in this chapter. Otherwise, you are returned to the Routing and Remote Access MMC snap-in, and you can now manage the NAT service as described in the next section.

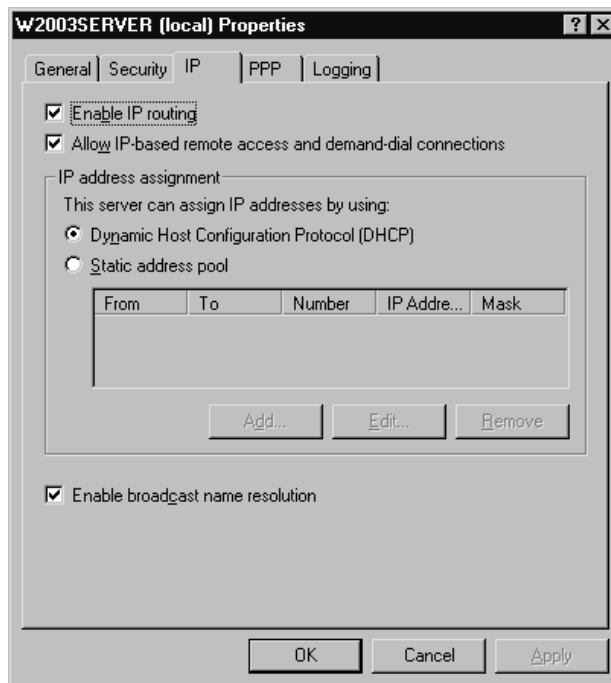
Managing NAT

After you have enabled RRAS and set up a NAT server, you can manage the server from the Routing and Remote Access MMC snap-in. Select the server and select **Action | Properties** to display the **Properties** dialog box. Select the **IP** tab within this dialog to display the IP properties, shown in Figure 5.2. This page allows you to manage the address assignment feature of NAT. The NAT server can assign IP addresses in one of two ways:

- Select **Dynamic Host Configuration Protocol (DHCP)** to use an existing DHCP server to handle addressing.
- Select **Static address pool** to explicitly list the IP addresses this server can assign to clients. Once you have selected this option, you can use the **Add**, **Edit**, and **Remove** options to create a list of one or more IP address ranges for the address pool.

The IP properties tab also include an option to manage the name resolution feature of NAT. Select the **Enable broadcast name resolution** option if you do not have a DNS or Windows Internet Name Service (WINS) server on the network to handle name resolution. If this option is selected, the RRAS server uses network broadcasts to resolve names. This eliminates the need for a dedicated name server on single-subnet Windows-based networks.

Figure 5.2 The IP Properties for an RRAS Server





TEST DAY TIP

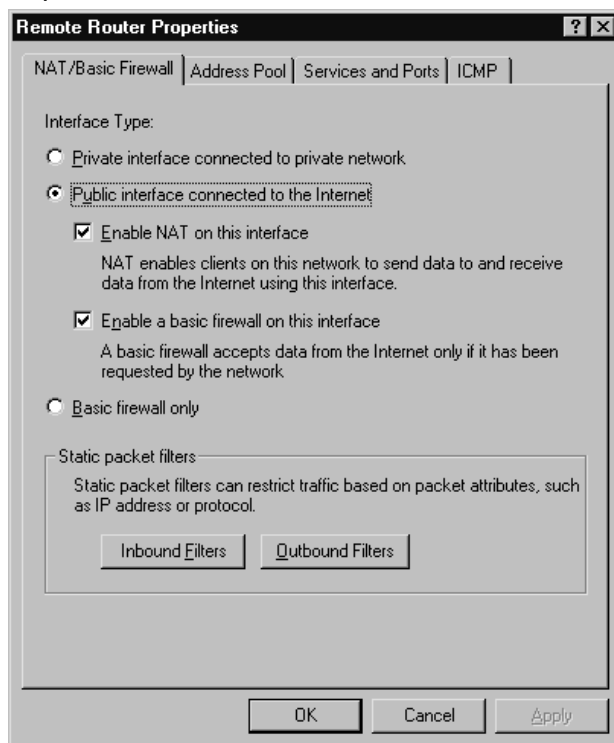
If you are not using broadcast name resolution, the NAT server needs to know the IP address of a DNS or WINS server to complete resolution requests. These server addresses are not part of the RRAS configuration. You must specify them using the **Properties** dialog box for the network interface.

Configuring a NAT Connection

You can also manage the settings for a NAT interface from the Routing and Remote Access console. To access these settings, select the **NAT/Basic Firewall** entry under **IP routing** in the left column, and then select **Action | Properties** from the menu. The **Properties** dialog box is divided into four tabbed sections:

- **NAT / Basic Firewall** On this tab, shown in Figure 5.3, you can enable or disable NAT for the connection. You can also enable a basic firewall, which prevents unauthorized traffic from the Internet from reaching the internal network. You can also use the **Inbound Filters** and **Outbound Filters** buttons to define IP filters to further secure the connection.
- **Address Pool** Allows you to define the Internet addresses that will be used by the NAT server. Don't confuse this with the pool of *private* addresses the server can assign to clients. At least one Internet address must be included here. You can also use the **Reservations** button to define an external address that always reaches the same internal client machine. This is useful if you need to run a Web server or other service and make it accessible over the Internet.
- **Services and Ports** Allows you to enable various services, such as FTP and Simple Mail Transfer Protocol (SMTP), that will be accessible to Internet users, and define the internal machines these packets will be routed to.
- **ICMP** Allows you to enable various types of diagnostic packets. These may be needed if you wish the NAT server to respond to PING or Traceroute diagnostics.

Figure 5.3 NAT Properties



How NAT Works

NAT transparently handles translation, so clients do not need to be aware that NAT is in use. Instead, they are configured with the NAT server's address as their default gateway. When a client sends an outgoing packet, it is sent to the NAT server. The NAT server receives the packet and performs the following tasks:

- The packet's destination address and port are stored in an entry in the NAT table, along with the internal address from which the packet originated.
- The packet's source address is changed to the NAT server's address, and a random port number is assigned.
- The packet is sent over the Internet.
- When the remote server responds, the response is sent to the NAT server at the port number previously assigned. The NAT server consults the NAT table to determine which client requested the response, edits the packet to use the client's internal IP address as its destination, and sends it to the internal network.

Some Internet protocols, such as FTP, store addressing information within the packet itself, which would not normally work with NAT. The NAT server uses a *NAT editor* to modify the addresses for these protocols. Windows Server 2003 includes editors for several protocols. Keep in mind that some protocols may not be supported across the NAT server.

Internet Connection Sharing (ICS)

Internet Connection Sharing (ICS) is a simple implementation of a NAT server and is included with all versions of Windows 2000, Windows XP, and Windows Server 2003, as well as Windows 98 Second Edition and Windows Me. It is much easier to configure and use than the full NAT service. Although ICS supports the basic translation features of NAT, it has a couple of limitations:

- ICS supports only a single Internet IP address and a single LAN connection. The full NAT service can connect any number of public IP addresses to multiple LANs.
- ICS cannot be used on networks that have a DHCP or DNS server implemented.



TEST DAY TIP

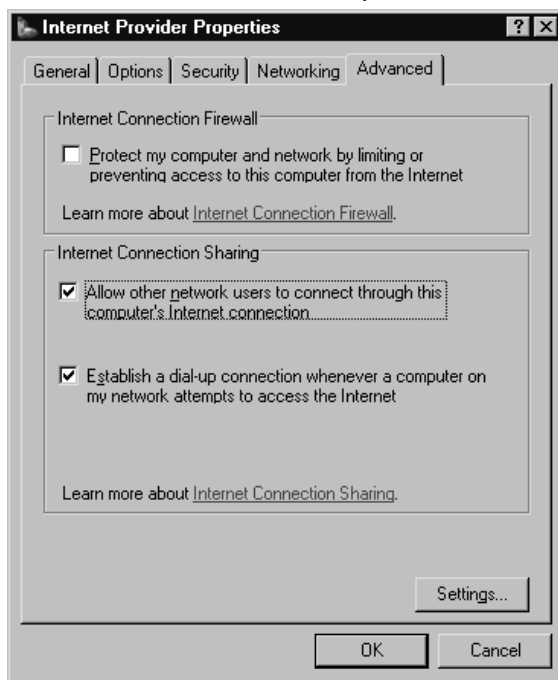
You should use ICS only when you are not using the NAT feature on the server, or when you are using an operating system for the NAT host, such as Windows XP, that supports ICS but not the full NAT service.

Activating the ICS Service

ICS is included and installed automatically with all versions of Windows Server 2003 and Windows 98 Second Edition and later. This feature is disabled by default, but enabling it is a simple process.

To enable ICS, open the **Properties** dialog box for the network adapter that connects to the Internet and select the **Advanced** tab. The **Advanced** properties are displayed, as shown in Figure 5.4. To enable ICS, simply check the **Allow other network users to connect through this computer's Internet connection** option. You can also optionally check the **Establish a dial-up connection whenever a computer on the network attempts to access the Internet** option for a dial-up Internet connection.

Figure 5.4 The Advanced Internet Provider Properties



TEST DAY TIP

The ICS options are included only in the **Advanced** tab of the **Properties** dialog box for Internet connections. LAN connections, such as the default Local Area Connection, do not include this option, since they connect only to the local network. You will, however, find the **Connection Sharing** option in the **Properties** dialog box for VPN connections.


Configuring Services

ICS is primarily a way for computers on your network to access Internet services, but it also allows you to configure services that are provided by a machine on your network and available via the Internet. When you use this option, incoming requests from the Internet are received by the ICS server and forwarded to whichever local machine is providing the service.

When ICS is enabled, you can click the **Settings** button in the **Advanced** tab of the **Properties** dialog box to configure the services available on your network and specify which client machines provide them. No services are enabled by default. The **Services**

dialog box, shown in Figure 5.5, lists a number of common services and allows you to configure them or add additional services.

Figure 5.5 The Network Services That Internet Users Can Access



Whether you use one of the predefined services, such as an FTP server or a Telnet server, or configure a custom service, you need to specify which computer on the local network will provide the service. Exercise 5.02 demonstrates the process of adding a new service.

EXERCISE 5.02

ADDING A CUSTOM SERVICE

You need to add an entry for any service on your network that should be accessible from outside the network. For example, the Network News Transfer Protocol (NNTP) service is not included as one of the default options, so you can add an entry for it. Follow these steps to add a custom service:

1. From the **Network Connections** window, right-click the Internet connection you are sharing and click **Properties**.
2. Select the **Advanced** tab.

3. Ensure that the **Allow other network users to connect through this computer's Internet connection** is enabled and click **Settings**.
4. The **Services** dialog box is displayed. Click **Add**.
5. The **Service Settings** dialog is displayed. In the **Description of service** text box, enter **Net News Transfer Protocol**, as shown in Figure 5.6.

Figure 5.6 Service Settings

6. In the **Name or IP address** text box, enter the machine name or IP address for the local machine providing the service.
7. In the **External port number for this service** text box, enter **119**.
8. In the **Internal port number for this service** text box, also enter **119**.
9. Click **OK**.
10. You are returned to the **Services** dialog box, and the new service is now listed. Click **OK** to return to the **Properties** dialog box.

Implementing Virtual Private Networks (VPNs)

Traditionally, when you are setting up a private network that spans multiple locations, you use one or more private wide area network (WAN) links to connect the locations (for example, T1 lines). While this provides secure high-speed communication between the loca-

tions, it is also relatively expensive. A VPN eliminates the need for dedicated WAN links by taking advantage of readily available connections to the public Internet.

A VPN is defined as a private network that uses virtual links through a public network rather than dedicated WAN links. These virtual connections use a technology called *tunneling* to encrypt private data and encapsulate it in packets to be transmitted over the public network.

Windows Server 2003 includes VPN functionality as part of RRAS. You can configure a Windows Server 2003 machine to act as a VPN server, which manages the VPN connections between clients or networks.



TEST DAY TIP

One advantage of using a VPN connection, rather than a dedicated leased line, is that the VPN connection is flexible. For example, if you move a location, all that is required to reconnect to the VPN is an Internet connection of any type.

Internet-based VPNs

One common use for a VPN server is to allow clients to remotely access the network. For example, you might have employees who work from home or who need network access from their laptops while on the road. Traditionally, this would require a pool of modems and a dial-up RRAS server, or a dedicated WAN link. With a VPN, since remote clients often have Internet connectivity, you can configure a VPN server to accept connections from these clients over the Internet. This provides them with a secure connection to the network without the need for modems or phone lines, and it often saves money, since a client can use a low-cost ISP with a local phone number rather than making a long-distance call.




NOTE

Microsoft refers to a VPN connection used for remote access as an *Internet-based VPN*. This is also known as a client-server VPN connection. The other type is a router-to-router connection. Although both types use the Internet for connectivity, Internet-based VPN refers to client-server connections.

How Internet-based VPNs Work

Figure 5.7 shows how a typical Internet-based VPN works. The remote client connects to the public Internet and uses VPN client software to initiate a connection with the VPN server. Communications for the VPN are encrypted and encapsulated into packets sent over the Internet.

Figure 5.7 Communications in an Internet-based VPN

Configuring Internet-based VPNs

RRAS supports the protocols needed for a VPN. You can configure these individually or use the RRAS Setup Wizard to configure a VPN server. Exercise 5.03 guides you through the process of configuring a VPN server using the Wizard.

EXERCISE 5.03

CONFIGURING A VPN SERVER USING THE WIZARD

If you have not yet configured RRAS on a server, you can use the Routing and Remote Access Server Setup Wizard to configure the server with the basic options for a VPN server.



NOTE


If you have previously configured the server to use RRAS, in order to perform this exercise you will need to first disable it. To do so, right-click the RRAS server name in the left console panel of the Routing and Remote Access MMC and select **Disable Routing and Remote Access**.

Follow these steps to configure the VPN server:

1. Select **Start | Programs | Administrative Tools | Routing and Remote Access** to start the Routing and Remote Access MMC snap-in.
2. Click the RRAS server name (usually the current machine) in the left column to highlight it.
3. From the menu, select **Action | Configure and Enable Routing and Remote Access**.
4. The Routing and Remote Access Server Setup Wizard displays a Welcome window. Click **Next** to continue.

5. The Configuration window appears (see Figure 5.1, earlier in the chapter). Select **Virtual Private Network (VPN) access and NAT** from the list and click **Next**.
6. The Wizard displays a final confirmation window, as shown in Figure 5.8. Click **Finish** to enable the RRAS and VPN features.

Figure 5.8 Completing the Routing and Remote Access Server Setup Wizard



7. A dialog box asks whether you wish to start the RRAS service at this time. Click **Yes**.

Windows Server 2003 next starts the RRAS service and can accept VPN connections. You are returned to the Routing and Remote Access MMC snap-in, where you can customize the settings for the VPN server.

Router-to-Router VPNs

While an Internet-based VPN provides easy remote access for individual clients, you can also configure a larger-scale VPN to connect two geographically separated LANs. A router-to-router VPN requires an Internet connection for each LAN, and it encapsulates traffic on the Internet to create a virtual WAN between the locations.

A router-to-router VPN can either use *demand-dial connections*, creating the VPN only when it is required for traffic between the networks, or *persistent connections* for an always-on

VPN. In either case, it can save money, since Internet connectivity is usually available at a lower cost than a dedicated WAN link between geographically separated sites. The longer the distance, the more money you are likely to save.

On Demand/Demand-Dial Connections

A demand-dial connection is often the most practical choice for small remote sites that only occasionally require VPN connectivity. RRAS supports one or more demand-dial connections. You can configure a connection using the Network Interfaces node in the RRAS MMC snap-in. Exercise 5.04 demonstrates how to add a new demand-dial interface.


EXERCISE 5.04

CONFIGURING A DEMAND-DIAL INTERFACE

You can add a new demand-dial interface on any RRAS computer that has RRAS configured. If you have not yet configured and enabled RRAS, see the instructions earlier in this chapter. Follow these steps to create a new demand-dial interface:


1. From the **Routing and Remote Access** MMC snap-in, right-click the **Network Interfaces** item in the left column and select **New Demand-dial Interface**.
2. The **Demand-Dial Interface Wizard** displays an introductory message. Click **Next** to continue.
3. You are prompted for a name for the new interface, as shown in Figure 5.9. Enter the name and click **Next**.

Figure 5.9 Enter a Name for the Demand-Dial Interface



4. The Connection Type window appears. Select **Connect using virtual private networking (VPN)** and click **Next**.
5. The VPN Type window is displayed. You can choose one of the VPN protocols (described in the “VPN Protocols” section later in this chapter). Select **Automatic selection** and click **Next**.
6. You are prompted for the host name or IP address of the remote router. Enter an address or name and click **Next**.
7. The Protocols and Security window is displayed, as shown in Figure 5.10. Enable the **Route IP packets on this interface** option and click **Next**.

Figure 5.10 Choose Protocols and Security Options



8. The **Static Routes for Remote Networks** window is displayed. Click **Add** to add a static route. Specify a destination address and subnet mask, and then click **OK**.
9. Click **Next** to continue.
10. The **Dial Out Credentials** window is displayed. Enter a username, domain name, and password to connect to the remote network, and then click **Next**.
11. The Wizard displays a completion message. Click **Finish** to complete the configuration of the demand-dial interface.

After you have completed this process, the new interface you created is listed in the **Network Interfaces** section of the **Routing and Remote Access** MMC snap-in. You can select this entry and open its **Properties** dialog box to change the configuration.

One-Way versus Two-Way Initiation

You can configure a demand-dial VPN with either one-way or two-way initiation:

- In one-way initiation, one VPN server is configured to accept demand-dial connections, and the other initiates the connection.
- In two-way initiation, both VPN servers are configured to accept connections. Whenever a client of one server requires access to the VPN, it initiates a connection to the other server.

Persistent Connections

Instead of using a demand-dial connection, a VPN server can use a persistent (always-on) connection to the Internet, such as an existing Digital Subscriber Line (DSL) connection. If the computer you are using as the VPN server is configured to use this type of Internet connection, it can be made available to VPN clients. To create a new persistent connection, select **Start | Control Panel | Network Connections | New Connection Wizard**.

Remote-Access Policies

You can secure a demand-dial connection in the same way that you secure a connection for a remote user. The calling router requires a user account on the VPN server. You can configure this user account's properties with the **Allow Access** option in the **Dial-in** properties section to explicitly allow access, or if access is controlled through a Remote Access Policy, the policy should grant the appropriate user remote access permissions. If you are using RADIUS authentication (explained in the "Using Internet Authentication Service (IAS)" section later in this chapter), the policy is configured on the RADIUS server rather than on the RRAS server.

Each remote-access policy is associated with a dial-in profile, which allows you to configure how the connection can be used. You can use the policy and profile settings to configure the authentication methods allowed, the hours in which dialing out is allowed, and other settings. These options are explained in detail in Chapter 7.

VPN Protocols

A VPN is created using a *tunneling protocol*. This is a standard communication protocol that creates a tunnel through the public network and transmits private data in encrypted form.

This is accomplished using *encapsulation*, a process that encrypts each VPN packet, combines it with a header to form a standard IP datagram, and sends it over the public network. Windows Server 2003 supports two standard tunneling protocols: the Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP).

PPTP

PPTP is the oldest and most common VPN protocol. PPTP is based on the Point-to-Point Protocol (PPP), which is typically used for dial-up connections. PPTP encapsulates PPP frames into IP packets, encrypts the data, and transmits them over the Internet.

PPTP in Windows Server 2003 is based on the existing PPP infrastructure and supports the same authentication methods as PPP, such as the Password Authentication Protocol (PAP) and Microsoft Challenge Handshake Authentication Protocol (MS-CHAP). When a higher-level authentication method is used, PPTP supports Microsoft Point-to-Point Encryption (MPPE), a strong method of encrypting VPN traffic before allowing it to traverse the public network.

L2TP

L2TP is a more recent tunneling protocol that offers additional features over PPTP. L2TP is a generic tunneling protocol that can encapsulate packets of many types for transmission over a network. Unlike PPTP, L2TP does not include encryption. Windows 2003 VPNs use the IP Security protocol (IPSec) to encrypt data sent over an L2TP tunnel. This provides end-to-end encryption and greater security than the MPPE encryption used with PPTP. Refer to Chapter 7 for more details on tunneling protocols.

VPN Security

A VPN combines encapsulation with encryption to create a connection between two systems. Depending on the VPN tunneling protocol you use, one of two encryption protocols is used to encrypt the data before it passes through the public network: MPPE or IPSec.

MPPE

MPPE is used with VPNs created by PPTP. MPPE provides encryption for the tunnel only; it does not provide end-to-end encryption from the client to the VPN server. MPPE requires that the client and server support either the MS-CHAP or Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) authentication method. These methods are described in detail in the “Authentication Methods” section later in this chapter.

IPSec


IPSec is an Internet standard for encrypted IP traffic. Since the L2TP tunneling protocol does not include encryption by itself, IPSec is used to encrypt the data before it is encapsu-

lated across the tunnel. Unlike MPPE, IPSec does provide end-to-end encryption. You can use IPSec over an established PPTP link to add end-to-end encryption.



TEST DAY TIP

IPSec also supports *tunnel mode*, a built-in ability to create a VPN tunnel without the use of L2TP. This mode works only with router-to-router VPNs. It is an advanced feature and is only necessary to support certain hardware that does not support the standard PPTP or L2TP tunneling protocols.



Using Internet Authentication Service (IAS)

While basic RRAS security is sufficient for small networks, a larger enterprise often needs a dedicated infrastructure for authentication. RADIUS is a standard for dedicated authentication servers. A RADIUS server provides centralized authentication and access control, and it can also provide detailed accounting for the use of its services. RADIUS services can be scaled to handle any enterprise's authentication needs and extended with multiple authentication servers.

Windows Server 2003 includes Microsoft Internet Authentication Service (IAS), an implementation of a RADIUS server. IAS supports authentication for Windows-based clients, as well as for third-party clients that adhere to the RADIUS standard. IAS stores its authentication information in Active Directory (AD), and you can manage it with Remote Access Policies.



NOTE

For more detailed information about configuring IAS for specific uses, such as wireless authentication, see Chapter 7.

Advantages of IAS

While IAS requires the use of an additional server component, it provides a number of advantages over the standard methods of RRAS authentication. These advantages include centralized authentication for users, auditing and accounting features, scalability, and seamless integration with the existing features of RRAS.

Centralized User Authentication and Authorization

In the RADIUS standard, remote users do not connect directly to the RADIUS server. Instead, they connect to a network access server (typically an RRAS server), which acts as a RADIUS client, connecting to the IAS server and authenticating the user. This provides for

centralized authentication. Any number of RRAS servers can connect to the same IAS server for authentication.

Centralized Auditing and Accounting

Along with authentication, IAS supports auditing features—tracking when the system is used, when errors occur, and so on—and can keep a centralized record of usage of the remote access or VPN servers. This record is stored in a log file, which you can import into a database or analyze to determine traffic patterns or potential problems.

RRAS Integration

IAS supports the same Remote Access Policy settings as RRAS. You can use these settings on a simple RRAS server in a small network, and later add an IAS server, move the policies to the IAS server, and configure one or more RRAS servers to authenticate using IAS. When using IAS for authentication, RRAS servers no longer have their own Remote Access Policies, since the IAS server manages a centralized policy.

Control via Remote-Access Policies

As with basic RRAS security, you can define remote-access policies to configure remote-access security with IAS. You can define a single set of remote-access policies on the IAS server, and they will be used by every RRAS server that uses IAS for authentication. This centralized authentication allows you to quickly define policies for the entire enterprise without the need to manage individual policies for each RRAS server.

Extensibility and Scalability

IAS provides an extensible architecture for authentication. While it provides only a small advantage over traditional Windows authentication methods when used on a small network, IAS excels in large enterprises because it provides centralized authentication. You can scale from a single IAS server to multiple IAS servers interacting with multiple RRAS servers in a global network. When you add a new RRAS server, you don't need to configure its security separately; simply configure it to use the existing IAS server for authentication.

IAS Management

To support IAS, you will need one or more IAS servers. You can install IAS on a domain controller or member server. The server can be used for other components, such as RRAS, but if the IAS server will be heavily used, you may wish to dedicate a server for this purpose. You can use a single server or configure a second server to act as a backup. RRAS servers that authenticate using IAS can contact the backup server if they are unable to reach the primary server.

The IAS component is included with all editions of Windows Server 2003 except the Web Edition. You can install IAS on a Windows Server 2003 computer using the

Add/Remove Programs option in **Control Panel**. Exercise 5.05 demonstrates how to add this component to a server.

EXERCISE 5.05

INSTALLING IAS

Follow these steps to install IAS on a computer running Windows Server 2003:

1. Select **Start | Control Panel | Add/Remove Programs**.
2. Select the **Add or Remove Windows Components** option.
3. Select **Networking Services** from the list and click **Details**.
4. Check the box next to **Internet Authentication Service** and click **OK**.
5. Click **Next** to complete the installation.

Activating IAS Authentication

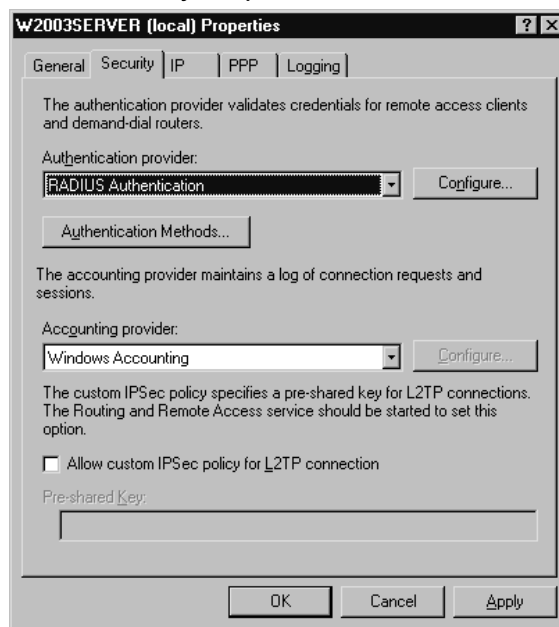
When you have a working IAS server on the network, you can configure the RRAS server to use IAS authentication. This will disable the normal Remote Access Policies in the Routing and Remote Access MMC snap-in and forward all authentication to the IAS server. You can then configure security settings for all RRAS servers centrally at the IAS server. Exercise 5.06 guides you through the process of enabling IAS authentication for an RRAS server.

EXERCISE 5.06

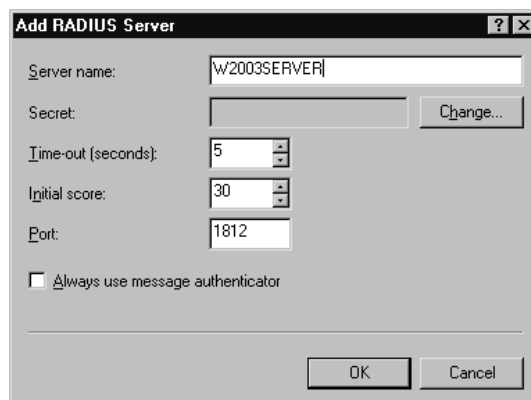
SELECTING IAS AUTHENTICATION

To select IAS authentication, you must have already configured and enabled RRAS services on the computer. Follow these steps to enable IAS authentication:

1. Select **Start | Administrative Tools | Routing and Remote Access**.
2. Click the RRAS server name in the left column to highlight it. Select **Action | Properties** from the menu, or right-click the RRAS server name and select **Properties** from the context menu.
3. The **Properties** dialog box is displayed. Click the **Security** tab. The **Security** properties are displayed, as shown in Figure 5.11.

Figure 5.11 Security Properties

4. In the **Authentication provider** drop-down list, select **RADIUS Authentication**.
5. Click the **Configure** button to display the RADIUS server options.
6. Click **Add** to add a RADIUS server to the list.
7. The **Add RADIUS Server** dialog box is displayed, as shown in Figure 5.12. Enter the name of the RADIUS server. You can optionally specify a shared secret using the **Change** button. Click **OK**.

Figure 5.12 Add a RADIUS Server

8. Click **OK** to exit the **Properties** dialog box.
9. A dialog box reminds you to restart RRAS to enable the new authentication method. Click **OK** to continue.
10. You are returned to the **Routing and Remote Access** MMC snap-in. Select the RRAS server in the left column and select **Action | All Tasks | Restart** from the menu, or right-click the server name and select **All Tasks | Restart** from the context menu.

RRAS is now restarted, and RADIUS authentication is enabled using the IAS server.



EXAM WARNING

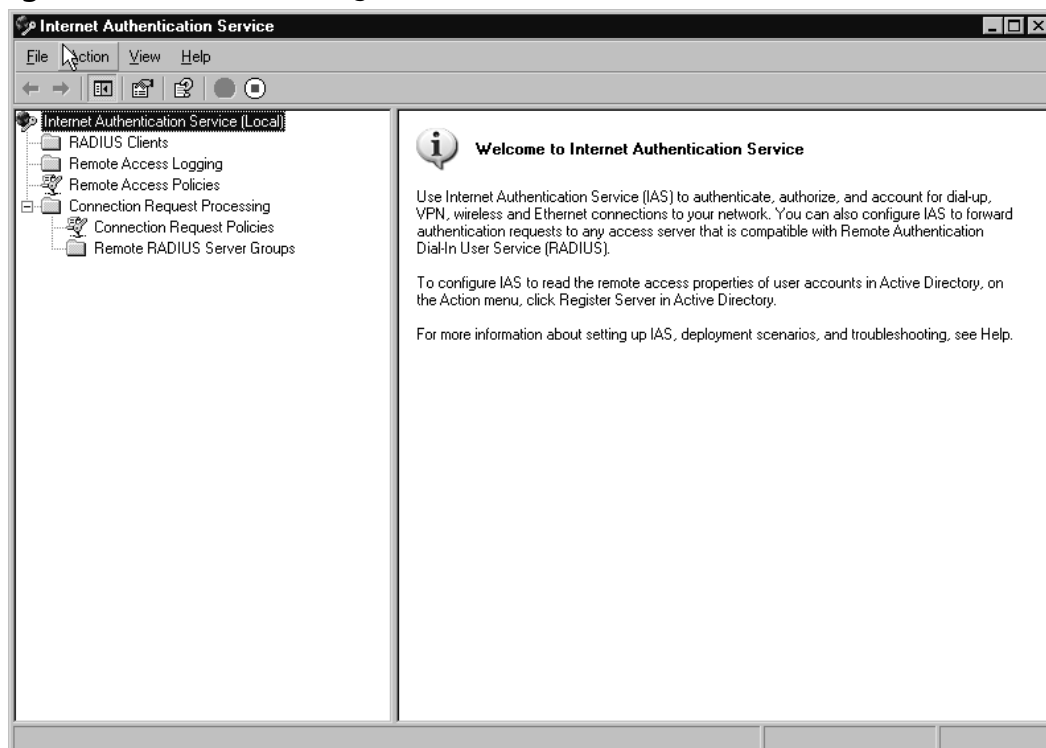
If you enter a shared secret (password) in the RADIUS Authentication settings of RRAS, it must be the same one you already specified in the properties of the IAS server. This password system provides a basic level of security between RADIUS clients and servers. Its primary purpose is to ensure that an unauthorized RADIUS server cannot be added to the network and used to provide incorrect authentication information.

Using the IAS MMC Snap-in

You can manage the configuration of an IAS server using its MMC snap-in. To launch the IAS management console, select **Start | Programs | Administrative Tools | Internet Authentication Service**. The IAS console is shown in Figure 5.13. The left column of the window displays several components of the IAS server that you can manage, including the following:

- **RADIUS Clients** Lists the clients (RRAS servers) currently configured and allows you to add new clients.
- **Remote Access Logging** Lists log files and allows you to configure additional logging options.
- **Remote Access Policies** Lists current policies and allows you to add policies. IAS policies are identical to those used on RRAS servers.
- **Connection Request Processing** Includes options for forwarding authentication requests to another IAS or RADIUS server for processing.

Figure 5.13 The IAS Management Console



IAS Monitoring

You can monitor the status of the IAS server using Windows Server 2003's standard monitoring facilities, including Event Viewer and System Monitor. IAS also supports Simple Network Management Protocol (SNMP) for centralized monitoring of IAS, along with other devices and services.

IAS also adds a number of objects to the System Monitor utility when you install it. You can use the counters within these objects to monitor the performance of the IAS server. To use **System Monitor**, select **Start | Administrative Tools | Performance**, click the **Add Counters (+)** button, and select one of the IAS objects to view a list of the available counters.

IAS SDK

Microsoft also makes an IAS Software Development Kit (SDK) available. You can use this to create customized behaviors for IAS, control the number of network sessions available to users, and create customized methods of authorization and authentication. The SDK also includes development tools for the Extensible Authentication Protocol (EAP) to allow you to create new types of authentication. EAP is described in the next section.

Authentication Methods

The Windows Server 2003 IAS server supports a number of different authentication methods. These range from basic, unencrypted authentication to highly secure methods. Windows Server 2003 also supports an infrastructure that allows external methods of authentication, such as smart cards. In the following sections, we will discuss authentication methods supported by IAS.

PPP-based Protocols

IAS supports several simple authentication methods based on the authentication used with PPP. These are the same basic methods supported by native RRAS authentication. The following are the basic authentication methods you can select:

- **Unencrypted Password (PAP)** This option uses PAP, a basic unencrypted authentication method. Since PAP transmits passwords as plaintext, it provides very little security.
- **Shiva Password Authentication Protocol (SPAP)** SPAP is Shiva's extended version of PAP and is slightly more secure. This protocol is included for use with legacy devices and systems that require it.
- **Encrypted authentication (CHAP)** CHAP is a standard protocol that uses encryption to prevent password snooping. In CHAP, the server sends an encrypted challenge to the client, and the client uses the password entered by the user to decrypt it and send a response.
- **Microsoft encrypted authentication (Microsoft-CHAP)** MS-CHAP is Microsoft's extension of CHAP, which improves security and integrates with Windows authentication. Version 1 of MS-CHAP is included to support older operating systems.
- **Microsoft encrypted authentication version 2 (MS-CHAP v2)** MS-CHAP version 2 is an improved version that increases security. Since version 2 is supported by all current versions of Windows, you should choose it over version 1, unless you are supporting older clients.

EAP

Another choice for Windows Server 2003 and IAS authentication is EAP. EAP is not strictly an authentication protocol; it is a structure that allows numerous plug-in authentication methods. EAP also allows clients and servers to negotiate the most secure authentication method they both can support.

The EAP Infrastructure

Authentication protocols that fit into EAP are called *EAP types*. Each of these types is handled by a plug-in module. When a client connects to the server and both support EAP, they negotiate an EAP type for authentication, depending on which types each of them supports. A server that responds to authentication requests is called an *authenticator*. The authenticator can make any number of requests for information from the client, depending on the authentication type.

Enabling EAP-based Authentication

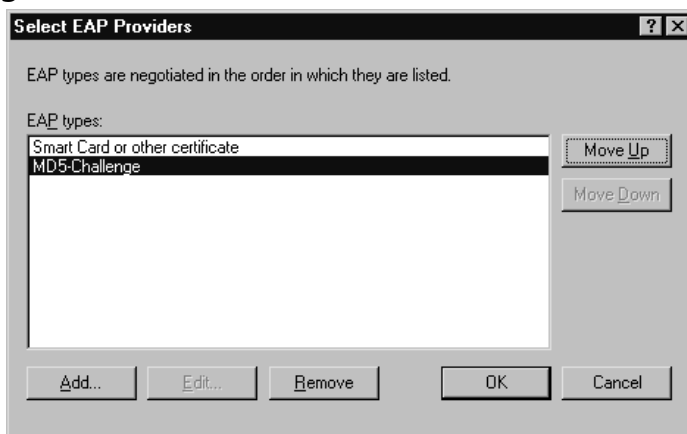
To enable EAP authentication on an IAS server, you create a Remote Access Policy that allows EAP authentication, or you modify an existing policy. Exercise 5.07 demonstrates how to modify a policy to allow the use of MD5 CHAP authentication through EAP.

EXERCISE 5.07

ENABLING EAP-BASED AUTHENTICATION

You can enable EAP authentication for any Remote Access Policy and specify the EAP types that can be used. Follow these steps to enable EAP authentication:

1. Select **Start | Administrative Tools | Internet Authentication Service**.
2. The IAS management console is displayed. Click to highlight **Remote Access Policies** in the left column.
3. In the right column, select **Connections to Microsoft Routing and Remote Access Server**.
4. Select **Action | Properties** from the menu, or right-click and select **Properties** from the context menu.
5. The **Properties** dialog box is displayed. Click the **Edit Profile** button.
6. The **Edit Dial-in Profile** dialog box is displayed. Select the **Authentication** tab.
7. The authentication methods supported by IAS are displayed, as shown in Figure 5.14. You can enable or disable the non-EAP authentication methods here. You can also change the order in which the selected EAP types are negotiated by moving them up or down in the list, using the **Move Up** and **Move Down** buttons.

Figure 5.14 Authentication Methods

8. Click the **EAP Methods** button. A list of the currently enabled EAP types is displayed.
9. Click **Add** and select **MD5-Challenge** from the list.
10. Click **OK**, then click **OK** in the **EAP types** list.
11. Click **OK** to exit the **Edit Profile** dialog box.
12. Click **OK** to exit the **Properties** dialog box.

EAP authentication is enabled as long as one or more EAP types appears in the list during this procedure. You can also remove available types from the list to disable EAP types or remove support for EAP altogether.

EAP-MD5 CHAP

EAP-MD5 CHAP is an implementation of the same challenge-response system as MS-CHAP within the EAP infrastructure. It supports the same level of security as MS-CHAP v2, but clients must support EAP in order to authenticate with this protocol. Clients that support MS-CHAP but not EAP will require the non-EAP version of this protocol.

EAP-TLS

Transport Level Security (TLS) is an authentication protocol that uses public-key encryption. All messages between the client and server are securely encrypted. The encryption is similar to that used with the Internet Secure Sockets Layer (SSL) protocol. This is the highest level of security provided by Windows Server 2003's authentication methods.



TEST DAY TIP

EAP-TLS also supports *smart cards*. These are hardware devices that implement public-key encryption. Smart cards answer challenges within the hardware and do not transmit the private key, so they provide higher security than simple password authentication. For more information about smart card authentication, see Chapter 7.

EAP-RADIUS

EAP-RADIUS is not a true authentication method. This option is an interface between EAP and RADIUS. When you select EAP-RADIUS, you specify an external RADIUS server, and all requests for authentication are forwarded to the RADIUS server for processing. This provides a way for clients that only support EAP to be authenticated using the RADIUS server.

Authorization Methods

IAS supports a variety of methods of authorization, to determine whether a connection is allowed and what tasks it can perform. Custom authorization methods are also supported. The following sections discuss different types of authorization in IAS.

Dialed Number Identification Service (DNIS)

DNIS is a phone company service that identifies the number being called and allows you to authorize the connection based on that number. It is usually used with 800 and 900 numbers, where there are several different numbers that go into the same public exchange (PBX) system. In dial-up modem pools where several phone numbers can reach the same group of modems, you can use DNIS authorization to ensure that users are calling a valid number.

Automatic Number Identification (ANI) and Calling Line Identification (CLI)

You are probably familiar with caller ID, which works on consumer phone lines to provide the number from which a call originated. ANI and CLI are the business-line equivalent services. IAS can authorize connections based on ANI or CLI to allow access to valid incoming numbers.

Guest Authorization

Windows Server 2003's IAS service can optionally allow guest access for unauthorized users using the Guest user account. Because this access is unauthenticated, its use is not recommended in most cases, and it is disabled by default.

Access Server Support

In the RADIUS standard, the RADIUS server works with one or more network access servers (NASs) that provide access to the network. In Windows terminology, this usually means RRAS servers. IAS also supports the following alternate types of access servers:

- **RADIUS access server support** IAS supports RADIUS standard access servers, whether they are Microsoft servers running IAS or those from other vendors. The standards for RADIUS access servers are defined in RFCs 2865 and 2866.
- **Wireless access points** IAS can also provide authentication for wireless access points using the various 802.11 protocols for wireless networking. For this to work, the access point hardware must support RADIUS authentication using an external server.
- **Authenticating switches** Some Ethernet switches support RADIUS authentication to authorize nodes attached to the switch. IAS includes the Ethernet port type, which allows you to manage authentication for these switches.

Outsourced Dialing

IAS supports *outsourced dialing* (sometimes called *wholesale dialing*), a standard for the use of ISP modem pools. In this system, you contract with an ISP to provide your employees remote network access using the ISP's existing modems. Users connect to a modem at the ISP, and a server at the ISP creates a VPN tunnel to connect them to the LAN. A RADIUS server at the ISP can forward records to your organization's IAS server, which allows you to manage access to the modems and obtain auditing and accounting information for their use.

Outsourced dialing has a number of advantages. The ISP already maintains pools of modems, and you may be able to obtain access to them at a lower price than the cost of configuring your own modems. The ISP may also have physical presence in areas you do not have a facility to provide for local calls, and it relieves you of the burden of managing modem pools.

Using Connection Manager

Connection Manager is a Windows application that enables a client to initiate a dial-up or VPN connection to a server running RRAS. To set up a connection, you need to know whether you are using dial-up, VPN, or another connection type; the phone number or VPN server to connect to; and other information.

Fortunately, if you frequently have clients or employees that need to create a connection to the RRAS server, you can distribute a customized version of Connection Manager that already contains most of the required information to connect to the server. Microsoft

EXAM
70-293

OBJECTIVE

2
2.3

distributes the Connection Manager Administration Kit (CMAK), which guides you through the process of customizing Connection Manager and creating a distribution package.



TEST DAY TIP

Along with employees who wish to remotely access a company network, CMAK is often used by ISPs to provide a simple way to set up connections for their customers.

Using CMAK

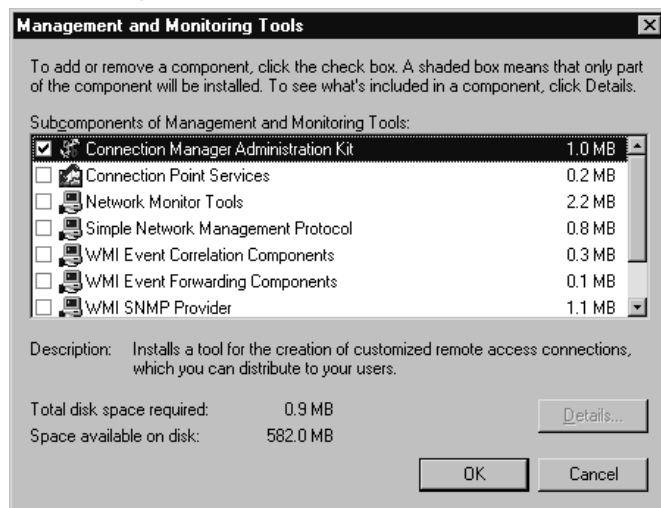
CMAK works as a Wizard that presents a series of questions about the connection you are using, and then creates a custom service profile that can be used with Connection Manager to easily initiate the connection.

Installing and Running CMAK

CMAK is included with Windows Server 2003. To install CMAK, follow these steps:

1. Select **Start | Control Panel | Add or Remove Programs**.
2. Select the **Add/Remove Windows Components** option.
3. Select **Management and Monitoring Tools** from the list and click **Details**.
4. Check the box next to **Connection Manager Administration Kit**, as shown in Figure 5.15.

Figure 5.15 Installing CMAK



5. Click **OK**, and then click **Next** to complete the installation. You will need the Windows Server 2003 CD-ROM.

After CMAK is installed, select **Start | Programs | Administrative Tools | Connection Manager Administration Kit** to launch the Wizard. Exercise 5.08 guides you through the process of using CMAK to create a simple service profile.

EXERCISE 5.08

USING THE CONNECTION MANAGER ADMINISTRATION KIT

The CMAK prompts you for several items of information. Follow these steps to use CMAK:

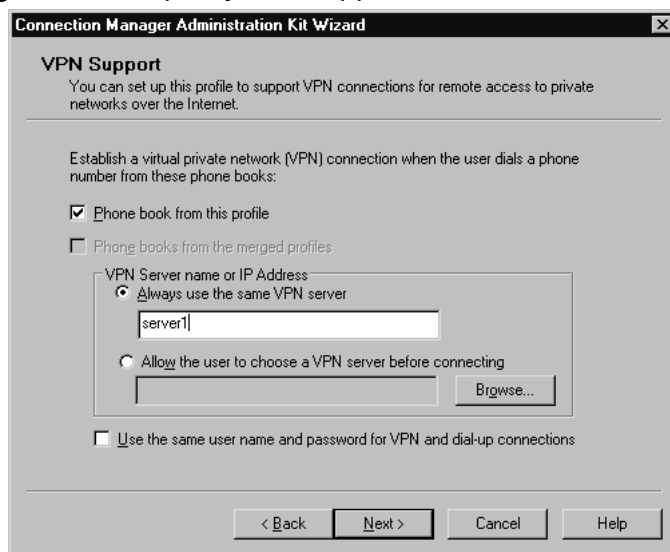
1. Select **Start | Programs | Administrative Tools | Connection Manager Administration Kit**.
2. An introductory window is displayed. Click **Next** to continue.
3. The next window asks whether you wish to create a new service profile or edit an existing one. Select the **New profile** option and click **Next**.
4. You are now prompted for a service name. Enter **Test Connection** in the **Service name** text box and **test** in the **File name** text box, as shown in Figure 5.16. Then click **Next**.

Figure 5.16 Specify a Service Name and Filename

The screenshot shows a dialog box titled "Connection Manager Administration Kit Wizard" with a close button (X) in the top right corner. The main title is "Service and File Names". Below the title is a descriptive text: "The service name identifies your profile to end users; the file name identifies your profile to administrators." There are two text input fields. The first is labeled "Service name:" and contains the text "Test Connection". The second is labeled "File name:" and contains the text "test". At the bottom of the dialog box, there are four buttons: "< Back", "Next >", "Cancel", and "Help".

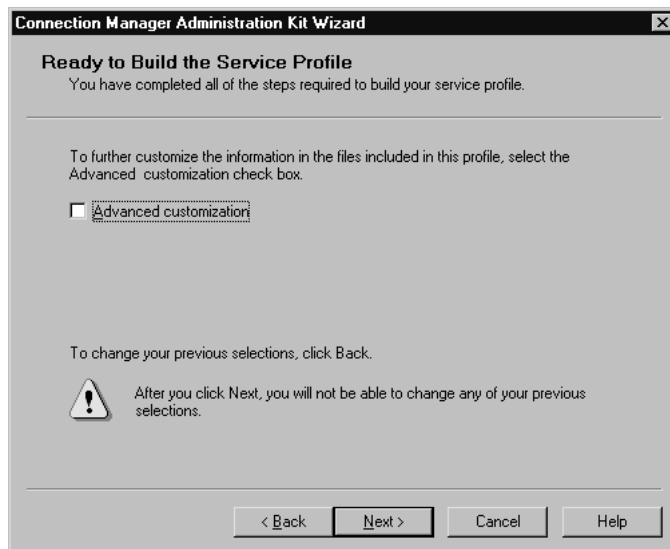
5. The next window asks whether you will be using a realm name. This allows you to add a standard prefix or suffix to usernames. Select **Do not add a realm name to the user name** and click **Next** to continue.
6. The **Merge Profiles** window is displayed. This allows you to merge phone numbers or other information from other profiles to the new profile. Click **Next** to continue.
7. The **VPN Support** window is displayed. This allows you to specify that a VPN connection will be created. Check the box next to **Phone book from this profile** and enter **server1** in the **VPN Server name or IP Address** text box, as shown in Figure 5.17. Then click **Next**.

Figure 5.17 Specify VPN Support



8. The **VPN Entries** window is displayed. Here, you can choose an existing VPN connection for the profile to support or create a new entry. Click **Next** to continue.
9. The **Phone Book** window is displayed. You can select a phone book file to provide access numbers to clients. Disable the **Automatically download phone book updates** option and click **Next**.
10. The **Dial-up Networking Entries** window is displayed. You can choose a current dial-up networking entry to use with the profile or create a new one. Click **Next** to continue.
11. The **Routing Table Update** window is displayed. Click **Next** to continue.

12. The **Automatic Proxy Configuration** window is displayed. Here, you can specify settings for a proxy server to be used with the connection. Click **Next** to continue.
13. The **Custom Actions** window is displayed. Custom actions are described later in this section. Click **Next** to continue.
14. The **Logon Bitmap** window is displayed. You can choose a default graphic or your own 330-by-140 pixel graphic to be displayed in the Connection Manager dialog box. Click **Next** to continue.
15. The **Phone Book Bitmap** window is displayed. You can choose a default graphic to be displayed in the phone book dialog box or specify a custom 114-by-309 pixel graphic. Click **Next** to continue.
16. The **Icons** window is displayed. You can choose custom icons for the connection or use the defaults. Click **Next**.
17. The **Notification Area Shortcut Menu** window is displayed. You can choose items to be included in a menu available from the icon in the notification area. This is useful to provide a default list of Internet applications, such as Web browsers or e-mail programs. Click **Next** to continue.
18. The **Help File** window is displayed. You can use a custom help file, as described later in this section. Click **Next** to continue.
19. The **Support Information** window is displayed. Enter a single line of text that will be displayed in the Connection Manager dialog box and click **Next** to continue.
20. You can choose whether to include the installation files for Connection Manager with your service profile. Select **Install Connection Manager** and click **Next** to continue.
21. In the next window, you can specify an optional text file to be displayed as a license agreement. Click **Next** to continue.
22. The **Additional Files** window is displayed. You can specify any files you wish to be included with the distribution. Click **Next** to continue.
23. The **Ready to Build the Service Profile** window is displayed, as shown in Figure 5.18. Click **Next** to begin building the service profile.

Figure 5.18 Ready to Build the Service Profile

24. A final window is displayed after your profile is created. Click **Finish** to exit the Wizard.

Service Profiles

When you complete the CMAK Wizard, your connection profile is stored as a self-extracting executable file. Any additional files you specified are also included in the distribution directory. CMAK creates a directory for your profile, typically under `C:\Program Files\CMAK\Profiles`. If you are distributing your customized version of Connection Manager to customers or employees, copy the files in this directory to a floppy disk or CD-ROM, or share the folder and provide them with the network path.

Custom Actions

CMAK supports *custom actions*, to run programs automatically during the Connection Manager process. This allows you to incorporate any custom software you wish into the Connection Manager. CMAK supports a variety of different actions that execute at different times:

- **Pre-init actions** Execute when Connection Manager starts.
- **Pre-connect actions, pre-dial actions, and pre-tunnel actions** Execute before starting a connection, depending on the type of connection in use.

- **Post-connect actions** Execute after a successful connection.
- **On cancel actions** Processed when the user cancels the connection.
- **On error actions** Used when an error occurs while connecting.

Custom Help

You can specify a custom help file for use with Connection Manager from the **Help File** window in the CMAK Wizard. You can use the default Connection Manager help file as a basis for your custom version. When you install CMAK, the source files for this help file are stored in the C:\Program Files\CMAK\Support\CMHelp folder. You can use any standard help file development tool, such as Microsoft's Help Workshop, to modify these files and compile the new help file.

VPN Support

CMAK supports VPN connections as well as dial-up connections. You can specify a VPN server, or a list of servers, and the protocols that will be enabled by default in Connection Manager. This makes it easy for clients with existing Internet connections to connect as VPN clients.

Connection Manager Security Issues

Although customizing Connection Manager with CMAK allows you to simplify the process of connecting to your network, it can also create several potential security issues. The following sections discuss some common security concerns when using CMAK and how you can address them.

Preventing Editing of Service Profile Files

You can edit service profiles using the CMAK Wizard, as explained earlier in this chapter. Only administrators can install this tool on other computers, and users must be members of the Power Users group to run an existing installation of CMAK. However, because the profiles created by CMAK are stored as simple text files, anyone who has access to the text file can modify any of its settings with a text editor.

To minimize the risk of users editing the text files, store them in a secure location. However, once you distribute the files to users, keep in mind that savvy users can edit the text files on their own computers. While this does not compromise your network security, realize that the constraints you created using CMAK might not always be followed.

Client Operating System, File System, and Configuration

CMAK can create Connection Manager profiles for a wide variety of Windows operating systems, which vary greatly in the levels of security they provide. Some features of

Connection Manager, such as user certificates, are not supported by older versions of Windows. For maximum security, require users to have a more recent operating system.

Preventing Users from Saving Passwords

When a computer is accessible by multiple users, there is always the risk of an unauthorized user using a connection. To minimize this risk, you can prevent users from using the **Remember Password** option to store the password for the connection on their computers. To disable this feature, set a value of **1** for the **HideRememberPassword** option in the connection profile. You can do this by selecting **Edit Advanced Options** from the CMAK Wizard's final screen or by editing the .cms file in a text editor.

Secure Distribution of Service Profiles

Your service profile might include private information, such as phone numbers, network server addresses and settings, and pre-shared keys. Depending on the level of detail this information includes, you might need to make sure that only authorized users can download or obtain a copy of your customized Connection Manager.

Summary of Exam Objectives

Internet connectivity is an important consideration in most networks today. The first consideration when planning Internet connectivity is whether to use a routed connection or a translated connection to the Internet. A routed connection places all machines in the network on the Internet, and each requires an IP address. A translated connection uses a separate private addressing scheme on the local network, and a server translates between public and private IP addresses to provide shared Internet access. In Windows networks, translated connections typically use NAT or ICS, a simplified version of NAT supported by Windows systems.

A VPN is an extension of a private network using a public network, such as the Internet, functioning as a conduit between two points. There are two basic types of VPNs: Internet-based VPNs, used by clients for remote access, and router-to-router VPNs, used to connect two segments of a WAN. VPNs use a tunneling protocol (such as PPTP or L2TP) to encapsulate data, in conjunction with an encryption protocol (such as MPPE or IPSec) to encrypt it before sending it over the public network. RADIUS is an Internet standard for a server that provides centralized authentication, authorization, and accounting services. Microsoft IAS is an implementation of a RADIUS server. RADIUS allows you to centralize the authentication and auditing features of one or more RRAS servers. RRAS servers connect to the IAS server with authentication requests. The IAS server supports Remote Access Policies, which replace the individual policies normally stored at each RRAS server.

Windows operating systems can use the Connection Manager utility for creating connections to dial-up networks or VPNs. Windows Server 2003 includes CMAK, which lets you use a Wizard to customize Connection Manager for your particular connections. CMAK allows you to specify dial-up or VPN server information, authentication settings, and other options to create the connection. You can also specify custom icons, graphics, and a help file to personalize Connection Manager.

Exam Objectives Fast Track

Connecting the LAN to the Internet

- ☑ LAN connections to the Internet can be either routed or translated.
- ☑ Windows Server 2003 supports NAT for translating addresses.
- ☑ ICS is a simplified, limited version of NAT.
- ☑ NAT is part of RRAS and can be installed using the Routing and Remote Access Server Setup Wizard or configured using the Routing and Remote Access management console.

Implementing Virtual Private Networks (VPNs)

- ☑ VPNs can be either Internet-based (providing remote access to clients) or router-to-router (connecting two networks that are in geographically separate locations).
- ☑ Router-to-router VPNs can use either persistent connections or demand-dial connections.
- ☑ Demand-dial VPNs can use either one-way or two-way initiation.
- ☑ VPN tunneling protocols include PPTP and L2TP.

Using Internet Authentication Service (IAS)

- ☑ RADIUS is the Internet standard for centralized authentication. IAS is the RADIUS server included with Windows Server 2003.
- ☑ IAS provides centralized authentication, accounting, and auditing.
- ☑ IAS integrates with RRAS and supports Remote Access Policies.
- ☑ IAS supports PPP-based authentication methods, such as MS-CHAP, as well as EAP.

Using Connection Manager

- ☑ Connection Manager is software you can use to make a connection, which automates much of the process for you. CMAK lets administrators use a Wizard to customize Connection Manager.
- ☑ CMAK stores the choices you enter in a text file called a connection profile and compiles the information into a customized executable version of Connection Manager.
- ☑ CMAK can be used to create connections for dial-up networks or for VPN clients.
- ☑ Because the customized Connection Manager can include specific access information for the network, using CMAK creates security concerns, and distribution of the connection profile should be restricted.

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 need to choose one VPN protocol or the other?

A: No, you can configure the VPN server to support both PPTP and L2TP, and clients can connect using the most secure protocol that is supported on their computers.

Q: What are the limitations of ICS as compared to NAT?

A: ICS supports a single LAN and a single Internet connection. It also lacks some of the configuration options of the full NAT service. For example, you cannot configure IP address assignment options. You also cannot use ICS on a network that has a DNS and/or DHCP server; NAT should be used in that case.

Q: Can a single RRAS server provide multiple functions, such as NAT and VPN access?

A: Yes, an RRAS server can support any of the features of RRAS simultaneously, although this will require you to customize the configuration.

Q: Can a single Windows Server 2003 computer act as both RRAS server and IAS server?

A: Yes, you can install IAS on a computer that is already running RRAS, and you can configure RRAS to use the local IAS server for authentication.

Q: What other options are included in a service profile for CMAK?

A: Along with the options the Wizard guides you through, profiles include a number of options for dealing with passwords, dialing options, VPN settings, username settings, and advanced options. Search Microsoft's TechNet site at www.microsoft.com/technet for a complete list of the configuration options CMAK supports.

Q: Are there alternatives to RRAS for forming VPN connections?

A: Yes, a number of hardware VPN devices are available. While they require additional expense, they provide a convenient "out-of-the-box" solution and may be a more robust solution than using a software VPN.

- Q:** Some routers support NAT. Is this the same translation feature supported by Windows Server 2003?
- A:** The Internet NAT standard defines a general process for address translation. The exact implementation varies between devices, but the functionality is the same.
- Q:** Can a client computer connect to two VPNs at the same time?
- A:** Yes, all this requires is a separate network connection entry for each VPN. You can connect to both using a single Internet connection.
- Q:** If I have ICS running for network translation, is there an easy way to upgrade to NAT?
- A:** No, you will need to configure NAT manually. Any custom service entries you have defined in ICS will need to be reconfigured in NAT.
- Q:** What is the difference between authentication and authorization?
- A:** Authentication refers to the methods RRAS or IAS use to determine a user's identity and verify that he or she is a legitimate user. Passwords, smart cards, and challenge-response systems provide authentication. Authorization is the process of determining what a client is allowed to do on the network after authentication.
- Q:** Is there any way to restrict connections to certain client operating systems?
- A:** A new Windows Server 2003 feature, Network Access Quarantine Control, allows you to create a script that must be run before a client is allowed access, and the script can check the client operating system or other factors. This feature is discussed in Chapter 7.

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.

Connecting the LAN to the Internet

1. You have five Windows XP clients on a network with a Windows Server 2003 server. The server has an always-on Internet connection with an ISP. What service can you install on the server to allow the clients to access the Internet, without requiring you to obtain additional IP addresses from your ISP?
 - A. PPTP
 - B. NAT
 - C. DHCP
 - D. DNS
2. You are configuring a simple network with two computers, both running Windows Server 2003. Both will be used as Web servers and must be accessible over the Internet. You have chosen to assign an Internet IP address to each machine, and you want to configure a single Internet connection for use by both machines. Which of the following is the best strategy?
 - A. Use a routed connection.
 - B. Use NAT.
 - C. Use ICS.
 - D. Two separate connections are required.
3. Your network includes a Windows Server 2003 computer and several workstations running Windows 2000 and Windows XP. You need to configure the server to provide shared Internet access to all machines on the network. The server will also act as a Web server. In addition, one of the workstations is providing an FTP service and requires its own Internet IP address. Which solution will address all of these requirements?
 - A. ICS
 - B. A hardware router
 - C. NAT
 - D. IAS

4. You have a DHCP server on the network that automatically assigns IP addresses to clients. You are configuring a NAT server to provide shared Internet access. You want clients to use internal addresses from the same pool, whether or not they are using the Internet. What is the most efficient way to do this?
 - A. Divide the address pool between the NAT server and the DHCP server.
 - B. Define identical address pools on the NAT server and the DHCP server.
 - C. Configure NAT to forward IP addressing requests to the DHCP server.
 - D. Remove the DHCP server from the network and use NAT exclusively.

Implementing Virtual Private Networks (VPNs)

5. You are planning a VPN to allow traveling employees to access the network from remote locations. Employees will be using a variety of ISPs to connect to the Internet. You want to ensure that the VPN offers end-to-end encryption between the VPN client and server for maximum security. Which VPN protocol should you use?
 - A. PPTP
 - B. L2TP only
 - C. L2TP and IPSec
 - D. PPP
6. You have configured a VPN server running RRAS under Windows Server 2003. A number of remote workstations are able to access the network by connecting to the Internet using local access methods and establishing a VPN connection. Which of the following terms describes this type of VPN?
 - A. Router-to-router
 - B. Point-to-point
 - C. Internet-based
 - D. One-way
7. You have configured a router-to-router VPN using two Windows Server 2003 computers as VPN servers, each with a local Internet connection. You have configured the VPN servers at each end of the VPN to use the PPTP protocol. Which of the following types of encryption will the VPN use in this configuration?
 - A. L2TP
 - B. MPPE
 - C. IPSec
 - D. EAP

8. You need to configure a VPN connection between the local network and a remote branch. The remote branch has access to a dial-up ISP and will be billed by the hour by the ISP for the time spent online. Which of the following is the best strategy to configure the VPN?
- A. Use a demand-dial connection.
 - B. Use a persistent connection.
 - C. Use dial-up access via RRAS.
 - D. Create a dedicated WAN link.

Using Internet Authentication Service (IAS)

9. You have three RRAS servers configured for VPN access for remote clients. The servers are currently using Windows authentication, and you wish to use IAS for centralized authentication. You have installed the IAS component on a Windows Server 2003 computer. What additional task is necessary to enable IAS authentication?
- A. Install IAS on all RRAS server computers.
 - B. Configure each RRAS server to use RADIUS authentication.
 - C. Install a RADIUS client.
 - D. Choose authentication protocols.
10. You have installed the IAS component on a Windows Server 2003 server. You are planning the authentication strategy for the IAS server and have configured the IAS server to use EAP for authentication. Which of the following protocols are supported by EAP? (Select all that apply.)
- A. MD5 CHAP
 - B. PAP
 - C. SPAP
 - D. EAP-TLS
11. You have an IAS server running Windows Server 2003. It supports a group of RRAS servers used to manage VPN connections for clients. You are configuring the authentication methods for the IAS server and want to allow the clients to use smart cards for secure and convenient authentication. Which of the following authentication protocols should you select?
- A. MS-CHAP
 - B. EAP-TLS

- C. MD5 CHAP
 - D. MS-CHAP v2
12. You have configured an RRAS server on one Windows Server 2003 computer and an IAS server on another, and configured the RRAS server to use the IAS server for authentication. In RADIUS terminology, which computer(s) are referred to as network access servers?
- A. The IAS server
 - B. The RRAS servers
 - C. The clients of the RRAS server
 - D. Both the IAS and RRAS servers
13. During a security audit, you are monitoring network traffic and notice that plaintext versions of passwords are passing through the network. You are using an IAS server to handle authentication. Which protocol do you need to disable at the IAS server to prevent this security risk?
- A. MS-CHAP
 - B. PAP
 - C. EAP-TLS
 - D. CHAP
14. You have an IAS server running Windows Server 2003. You need to enable and configure EAP to support clients that use EAP authentication. In the IAS MMC snap-in, where do you find the options for configuring EAP?
- A. Properties
 - B. Remote Access Policies
 - C. Protocols
 - D. Connection Request Processing
15. You wish to create client software for VPN clients to connect to the network so that clients do not need to manually specify the VPN server, tunneling protocol, and other settings. Which program allows you to customize the client software?
- A. Connection Manager
 - B. Connection Manager Administration Kit
 - C. RRAS MMC snap-in
 - D. IAS MMC snap-in

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**

2. **A**

3. **C**

4. **C**

5. **C**

6. **C**

7. **B**

8. **A**

9. **B**

10. **A, D**

11. **B**

12. **B**

13. **B**

14. **B**

15. **B**

MCSE 70-293

Planning, Implementing, and Maintaining a Name Resolution Strategy

Exam Objectives in this Chapter:

- 2.7 Plan a host name resolution strategy.
 - 2.7.1 Plan a DNS namespace design.
 - 2.7.2 Plan zone replication requirements.
 - 2.7.3 Plan a forwarding configuration.
 - 2.7.5 Examine the interoperability of DNS with third-party DNS solutions.
 - 2.7.4 Plan for DNS security.
- 2.8 Plan a NetBIOS name resolution strategy.
 - 2.8.2 Plan NetBIOS name resolution by using the Lmhosts file.
 - 2.8.1 Plan a WINS replication strategy.
- 2.5.2 Diagnose and resolve issues related to name resolution cache information.
- 2.9 Troubleshoot host name resolution.
 - 2.9.1 Diagnose and resolve issues related to DNS services.
 - 2.9.2 Diagnose and resolve issues related to client computer configuration.


Introduction

Computers “think” in ones and zeros, and the computers and routers on your Windows Server 2003 network communicate with one another using numbers in the form of IP addresses and MAC addresses. People, on the other hand, prefer to think in terms of names, which they use to represent computers and resources. This means that there must be a way to resolve the names used by the people to the numbers used by the computers, and that is where name resolution comes in. Without this mechanism, users will be unable to connect to resources using the “friendly” names they’re used to employing. Thus, it is an important part of the network administrator’s job to design and implement an effective strategy for name resolution on the network.

In this chapter, you’ll learn how to plan for the best way of resolving host and NetBIOS names on your network. We’ll first present an overview of host naming, including how host names are resolved using the hosts file and using the Domain Name System (DNS). We’ll discuss issues involved in designing a DNS namespace, such as choosing the parent domain name, the conventions and limitations that govern host names, the relationship of DNS and Active Directory (AD), and how to support multiple namespaces.

Then we move onto planning DNS server deployment. You’ll find out how to consider factors such as the number of servers, server roles, server capacity, and server placement. We’ll also show you how to plan for zone replication between your DNS servers, and we’ll address planning for forwarding and how DNS interacts with the Dynamic Host Configuration Protocol (DHCP) on a Windows Server 2003 network. We’ll discuss Windows Server 2003 DNS server interoperability with Berkeley Internet Name Domain (BIND) and other non-Windows DNS implementations. You’ll learn about zone transfers between Windows Server 2003 DNS servers and BIND servers, and we’ll discuss supporting AD with BIND. You’ll learn about split DNS configurations and how interoperability relates to other services such as Windows Internet Name Service (WINS) and DHCP. Next, we’ll address DNS security issues, including common DNS threats such as footprinting, redirection, and DNS denial-of-service (DoS) attacks. You’ll learn how to best secure your DNS deployment by using a split namespace and packet filtering. We’ll discuss how to determine the best DNS security level for your network. Next, we’ll look at DNS performance issues. We’ll show you how to monitor DNS server performance and how to analyze DNS server tests.

In the next section, we’ll address NetBIOS name resolution and provide an overview of how NetBIOS names are resolved using LMHOSTS files and NetBIOS name servers such as WINS servers. You’ll find out what’s new for WINS in Windows Server 2003, and we’ll show you how to plan WINS server deployment and WINS replication. We’ll walk you through the process of configuring WINS replication partnerships, including push-only, pull-only, and push/pull configurations. We’ll also discuss common WINS issues, including configuration, performance, and security issues. We’ll show you how to plan for WINS database backup and how to troubleshoot name resolution problems related to both host names and NetBIOS names.



Planning for Host Name Resolution

One of the most common sources of trouble on any Windows network—whether it’s a Windows NT, Windows 2000, or Windows Server 2003 network—is faulty name resolution. Computers cannot resolve the computer names to the proper IP addresses, or they cannot find an IP address associated with a computer name at all. When name resolution (the process of finding the IP addresses associated with computer names and services running on those computers) is not working perfectly, a multitude of problems can arise, including (but not limited to) the following:

- Users might not be able to log on to the network.
- Users might not be able to connect to applications and services residing on remoter computers.
- Domain controllers might not be able to communicate with each other.

In fact, problems with name resolution are so common that a typical first step in troubleshooting problems on a Windows network is to ensure that name resolution is working flawlessly. A common mantra that reflects this situation is the following: “The problem is irrelevant. The answer is DNS.” Although this is a gross oversimplification of the problems that can arise on a Windows network, it does contain a germ of truth.

It is critically important that an appropriate name resolution strategy be planned, implemented, and maintained on every Windows network. Starting with Windows 2000, correct host name resolution is a necessary condition for the proper operation of the network. This contrasts with Windows NT 4 and earlier networks in which correct NetBIOS name resolution is a necessary condition for the proper operation of the network. NetBIOS name resolution can still play an important and central role in Windows 2000 and 2003 networks; However, its importance is subordinate to that of host name resolution, and in some situations reliance on NetBIOS name resolution can be completely eliminated with careful planning.

Planning for host name resolution on a Windows Server 2003 network means developing and implementing a fault-tolerant and secure strategy, whereby host computers on the network are always able to resolve computer names to IP addresses and locate services running on the network in a timely manner. For example, to log on to a Windows Server 2003 network, client computers must be able to locate domain controllers that are able to process logon requests. On a Windows Server 2003 network, the primary mechanism for locating the domain controllers is host name resolution through DNS.

Understanding Host Naming

We have mentioned two different kinds of name resolution: host name resolution and NetBIOS name resolution. In order to understand host naming, you might find it useful to understand the differences between NetBIOS and host names. In the following sections, we’ll discuss the characteristics of each.

NetBIOS over TCP/IP

NetBIOS was originally developed to run on small broadcast-based networks. An early and commonly implemented network/transport protocol that relies on NetBIOS is NetBEUI, which is designed to run on a segment. NetBIOS itself was not designed to run on multi-segment networks, and it was not initially designed to run on TCP/IP networks. For NetBIOS applications to work properly, they must be able to locate computers by their NetBIOS computer names. An example of the use of a NetBIOS application is the use of the Universal Naming Convention (UNC) path to gain access to a share on a remote computer. The UNC path has the form `\\computername\sharename`.

On a single-segment network running NetBEUI, a computer trying to connect to a file share on a remote computer sends a broadcast request to find the Media Access Control (MAC) address (a unique 12-digit hexadecimal number on Ethernet networks) associated with the network adapter of the target computer. After it receives a reply to its request for the MAC address of the computer that owns the NetBIOS name, the requesting computer can establish a session with the target computer. NetBIOS names either belong exclusively to the device, such as a NetBIOS computer name, or they are group names that are not exclusive, such as domain names. In either case, each NetBIOS name must be unique on the network.

On a TCP/IP network, IP addresses, rather than NetBIOS names, are used to connect to destination hosts. The process of IP address resolution is similar to NetBIOS resolution in that it is broadcast-based. When a computer tries to establish a connection with a destination host on the same network segment, it sends out a broadcast request for the MAC address of the computer configured with the IP address of the destination host. When the destination host is on a separate network segment, it sends a request for the MAC address of the default gateway. When the computer learns the MAC address of the destination host (or the default gateway, if the host is on a remote subnet), it can begin communicating with it.

Obviously, for NetBIOS applications running on TCP/IP networks, some method must be implemented so that these applications can use computer names and resolve them to the appropriate IP and MAC addresses. This is accomplished through the use of a specific NetBIOS interface called NetBIOS over TCP/IP, also known as NetBT or NBT, implemented in the Windows TCP/IP protocol stack. This interface allows NetBIOS applications to translate NetBIOS names to IP addresses, which are then subsequently used to resolve to the appropriate MAC address (the MAC address of the destination host, if on the same local subnet, or the default gateway, if on a remote subnet).

Host Names

NetBIOS names are required only when using NetBIOS applications that provide access to services running on remote or local computers. In contrast, WinSock applications, which are specifically written to run on a TCP/IP stack, use the WinSock interface in the Windows protocol stack. These applications include Web browsers and servers, FTP servers

and clients, Internet e-mail clients and servers, and so on. However, since these applications are specifically designed to run on TCP/IP networks, they rely on IP addresses and not names to establish communications with a remote computer. Unlike the case with NetBIOS applications, it is not necessary to use a name to establish communications. When using a WinSock application, you need to use only the IP address of the destination host to establish communications. Host names are used in place of IP addresses to make it easier for the human operators of computers. It is much easier for most of us to remember a name than it is to remember a number.

In contrast to NetBIOS names, there is no necessary relationship between host names and the IP addresses of the computers they represent. In fact, multiple host names can be assigned to the same IP address, and a single host name can be assigned to multiple IP addresses. This last technique is used, for example, to provide a type of simple load balancing (round-robin DNS resolution) among multiple Web servers that are all hosting the same Web site. Also, unlike NetBIOS names, host names are not a necessary part of the configuration of the computer. A Web server, for example, does not need to be configured with the host name used to reach it.

For host names to resolve to the appropriate target computer IP address, the client computer needs to have some means of being able to resolve the host name to the remote IP address. There are two primary methods for resolving host names to IP addresses: using a hosts file or using a DNS server. (On a Windows network, the situation is a little more complicated, because methods of NetBIOS name resolution can be used when host name resolution fails.)

Understanding the Hosts File

The hosts file is a text file that is found on the local computer. On Windows-based computers, the path to the hosts file is `%systemroot%\system32\drivers\etc\hosts`, where `%systemroot%` is a variable used to identify the folder where the operating system is installed, such as `C:\Winnt` or `C:\Windows`.



NOTE

Even though the hosts file is a text file, it does not have a `.txt` extension. Therefore, you must ensure that the file is not saved with this extension appended to it. For example, if you open the file in Notepad and then save it as a text file the `.txt` extension will be the default. Then Windows will not recognize the hosts file. This can be particularly problematic on Windows machines that are set to hide common file extensions by default, because the file will appear to not have the extension when you view it in the file list in Windows Explorer. This is one of the first things you should check if you have a hosts file that doesn't seem to be working.

The hosts file contains a list of host names and the IP addresses associated with them. By default, the hosts file contains only one active entry for the host name localhost, which points to the loopback IP address of 127.0.0.1.

The structure of the file is simple. To add a host name to IP address mapping, simply insert a new line containing the IP address of the destination computer followed by the host name. You can enter either a simple name or a fully qualified domain name (FQDN) that contains dots, such as `www.syngress.com`.

Although the hosts file has largely been superceded by DNS as a method of name resolution, it still has a number of valid uses. For example, you can use it to substitute a shorter, simpler name for a longer, more complex name that is stored on a remote DNS server. You can also use it for testing purposes. Another purpose of the hosts file is to use it to deliberately block Web sites, such as those that serve banner ads on Web sites, by mapping the FQDN of the Web sites to the loopback address (127.0.0.1) of your computer. This technique is useful if you want to speed up browsing on Web sites and do not wish to be subjected to large numbers of banners ads. (See <http://pgl.yoyo.org/adserver/> for more information about this use.)

The use of a hosts file can also speed up the process of host name resolution. When you try to connect to a remote computer using a host name, Windows operating systems prior to Windows 2000 will first consult their DNS cache stored in memory, then consult the hosts file, and then consult the DNS server. Beginning with Windows 2000, the hosts file is parsed whenever modifications are made to it and the contents are stored in the DNS resolver cache, eliminating the second step. Prior to Windows 2000, the hosts file is parsed every time host name resolution is required and the result is not found in the DNS cache.

To verify that the contents of the hosts file are stored in the DNS cache, open a command prompt and enter the command **ipconfig /displaydns**. This will display the contents of the DNS cache. Make a modification to the hosts file and save it. Then run the same command again. You will see the entry you made in the hosts file listed in the DNS cache. Figure 6.1 shows the output of the **ipconfig /displaydns** command after the hosts file had been saved with the addition of a host record named `test_record`, pointing to 192.168.100.1.

Figure 6.1 Output of the `ipconfig /displaydns` Command Showing the Contents of the DNS Cache

```

C:\WINDOWS\system32\cmd.exe
C:\Documents and Settings\Administrator>ipconfig /displaydns

Windows IP Configuration

1.0.0.127.in-addr.arpa
-----
Record Name . . . . . : 1.0.0.127.in-addr.arpa.
Record Type . . . . . : 12
Time To Live . . . . . : 604347
Data Length . . . . . : 4
Section . . . . . : Answer
PTR Record . . . . . : localhost

test_record
-----
Record Name . . . . . : test_record
Record Type . . . . . : 1
Time To Live . . . . . : 604347
Data Length . . . . . : 4
Section . . . . . : Answer
A (Host) Record . . . . : 192.168.100.1

1.100.168.192.in-addr.arpa
-----
Record Name . . . . . : 1.100.168.192.in-addr.arpa.
Record Type . . . . . : 12
Time To Live . . . . . : 604347
Data Length . . . . . : 4
Section . . . . . : Answer
PTR Record . . . . . : test_record
    
```

Understanding DNS

In the early days of the Internet (prior to 1984), there was no such thing as DNS. A single individual was responsible for updating a `hosts.txt` file whenever computers were added to the network. This `hosts` file was downloaded to other computers in order to maintain an up-to-date list of host names. Obviously, this solution was not effective when large numbers of computers were added to the Internet. The solution that replaced the use of the `hosts` file was DNS.

A DNS server is a computer that contains a database of host names and IP addresses. A computer configured as a DNS client can use the DNS server to query the database for the purpose of resolving names to IP addresses. An important characteristic of DNS is that the DNS server itself runs on a remote computer and will resolve names to IP addresses, as long as the DNS client has access to it.

There are many DNS servers in use on the Internet. Collectively, these DNS servers comprise a distributed, hierarchical database containing *resource records* (RRs) that allow DNS clients to resolve the host names to IP addresses in the case of forward lookup zones, and IP addresses to host names in the case of reverse lookup zones. DNS is also responsible for supplying mail routing and other information for various Internet applications. Because the billions of RRs that compose the DNS database are distributed, and because DNS uses an efficient protocol for name resolution (UDP), its performance is exceptional and is for the most part unaffected by the very large number of host names to IP addresses that it must resolve on a daily basis.

Understanding Common Resource Records (RRs)

Knowing the nature and purpose of common DNS resource records (RRs) is important to an understanding of DNS in general. RRs are defined in RFC 1034. However, since the publication of this RFC, a number of new RR types have been added. RRs have the following components:

- **Owner Name** The domain name where the RR is found.
- **Type** A 16-bit value that identifies the type of RR such as an A, a PTR, an NS, an MX, or an SOA record.
- **Time To Live (TTL)** The amount of time that an RR will be cached on a server. This is an optional field for many RRs.
- **Class** A 16-bit value that identifies the class of the resource, such as IN for Internet. Windows Server 2003 DNS supports only the IN class. This is a mandatory field.
- **RDATA** A required field that contains information describing the resource. The length and format of this information vary according to the type and class of the RR.

Common RR types include the following:

- **A** Address record used to map names to IP addresses. When a DNS client queries for an address record, it will receive an IP address in the reply. Here is an example:

```
host1.syngress.com.    IN    A    192.168.100.5
```

- **PTR** Pointer record used to map IP addresses to names in reverse lookup zones. When a DNS client queries for a PTR record, it will receive a FQDN as the reply. Here is an example:

```
5.100168.192.in-addr.arpa.    PTR    host1.syngress.com.
```

- **MX** Message Exchanger record used to identify name(s) and priority of server(s) responsible for handling Simple Mail Transfer Protocol (SMTP) mail for a domain. Note in the following example that a specific host is identified for handling mail for the domain. A corresponding address record must be associated with the host name.

```
syngress.com.    MX    10    host1.syngress.com.
```

- **NS** Name Server record used to identify name servers that are responsible for identifying DNS servers for DNS resolution for a domain. Note in the following example that a specific host is identified for authorita

Continued

tive name servers for the domain. A corresponding glue address record must be associated with the host name.

```
syngress.com.      IN NS   ns1.syngress.com
```

- **CNAME** Canonical name used to map an alternate or aliased name to a primary or canonical domain name. The canonical name must exist, and there can be only one CNAME per alias. Here is an example:

```
aliasname.syngress.com CNAME www.syngress.com
```

- **SRV** Service locator record that all allows multiple servers hosting TCP/IP-based services to be located by means of a DNS query. This is used extensively to support AD. For example, SRV RRs allow clients to locate domain controllers that can process logon requests. Here is an example:

```
_ldap._tcp._msdcs      SRV 0 0   389 dc1.syngress.com
                        SRV 10 0  389 dc2.syngress.com
```

- **SOA** Start of Authority record used to indicate the name of origin for the zone, the name of the server that is the primary authority for the zone, and other properties (such as e-mail of the responsible administrator, the version number of the zone data file, and other fields). This record is always the first record to appear in the DNS data. In the following example, note the @ on the left side to designate the owner for the RR. This symbol is a shorthand designator to indicate the origin (domain name). It can be used with any record, but it is most often used with the SOA record.

```
@      IN      SOA      ns1.syngress.com.  dnsadmin.syngress.com. (
                                1              ; serial number
                                3600           ; refresh   [1h]
                                600            ; retry    [10m]
                                86400         ; expire   [1d]
                                3600 )        ; min TTL  [1h]
```

The fields in the SOA record merit some special attention:

- The *serial number* indicates the version number of the zone file. When this number is incremented on a primary DNS server, a secondary DNS server that is polling the primary DNS server will learn that it needs to update its zone file through a zone transfer.

Continued

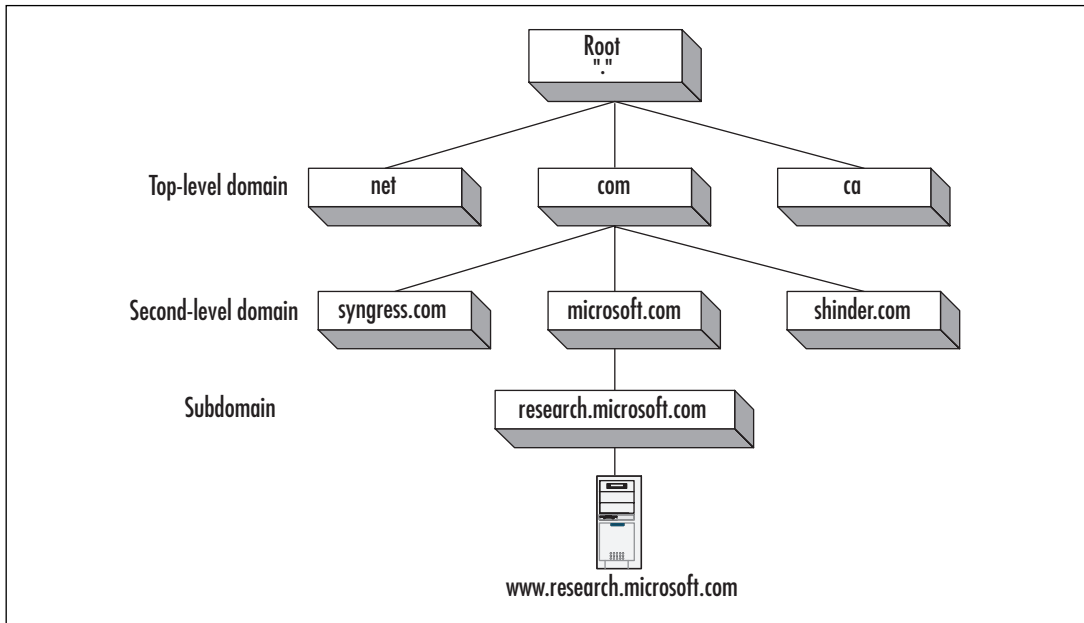
- The *refresh interval* indicates the length of time a secondary DNS server will wait before polling the primary DNS server to determine if it needs to copy the DNS data from the primary DNS server. This setting can be used to control the frequency of zone transfer traffic to secondary DNS servers.
- The *retry interval* indicates how long a secondary DNS server will wait after a failed zone transfer before trying again.
- The *expire interval* indicates how long a secondary DNS server will keep its records after failing to contact the primary DNS server. This prevents a secondary DNS server from retaining out-of-date data.
- The *min TTL* is the length of time a DNS resolver will cache records that it has queried on this server. The min TTL is a global value that is applied to all records, unless a specific TTL is specified in a particular RR.

Windows Server 2003 DNS also supports new RR types for IPv6, such as the AAAA RR for 128-bit IPv6 address. Here is an example of an AAAA RR:

```
host1.syngress.com.      IN      AAAA
                        4321:123:12:322:3:4:567:34de
```

For more information about DNS extensions to support IPv6, see RFC 1886 at www.rfc-editor.org/rfc/rfc1886.txt. For more information about IPv6, see Chapter 3 of this book.

The hierarchical tree on which DNS is based is called the *domain namespace*. At the top of the domain namespace is the root, or the dot (.), domain. Below the root domain are the various subdomains, beginning with the top-level domains, such as .com, .net, edu, and the various domain names that indicate country codes, such as .ca, .us, .de, and so on. Below the top-level domains are subdomains, referred to as the second-level domains, such as microsoft.com, syngress.com, and so on. These second-level domains can have further subdomains, such as authors.syngress.com or research.microsoft.com. Figure 6.2 shows an example of the hierarchical domain namespace.

Figure 6.2 Hierarchical DNS Namespace

In the example in Figure 6.2, the domain `research.microsoft.com` contains an RR for a host called `www`. When the name of the host is concatenated with the complete domain name from right to left (`www.research.microsoft.com`), the result is the FQDN for the host. The FQDN indicates the full path from the host to the root domain when read from left to right.

A true FQDN includes a period at the rightmost end of the domain name to indicate termination at the root zone. Thus, `www.research.microsoft.com.` (with the period at the end) is the true FQDN. This point is a source of some confusion among administrators, because they normally do not include this rightmost period when they, for example, type a destination in a Web browser. However, the output of diagnostic utilities such as `NSLookup` will be affected if you do not include the period.

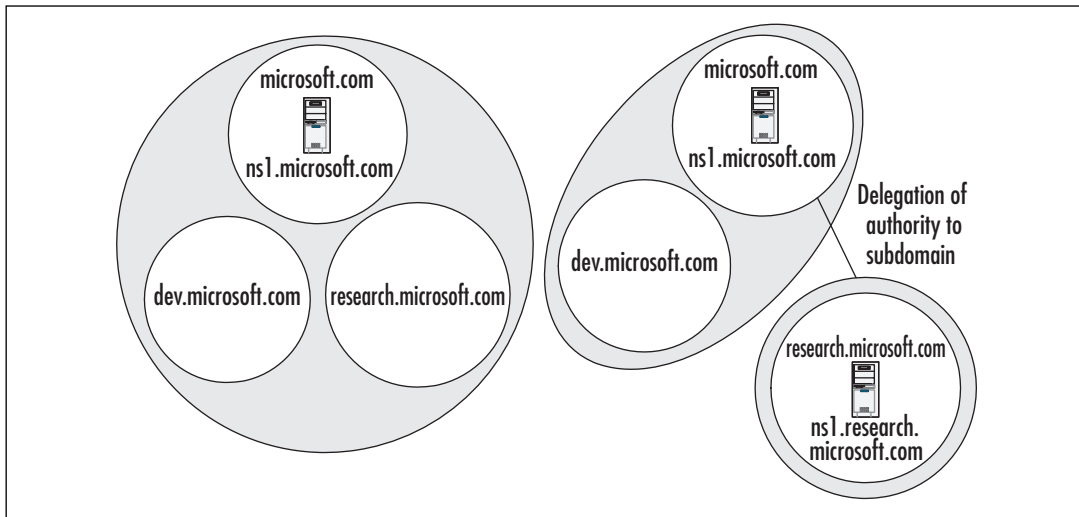
Domains versus Zones

It is critically important to understand the difference between domains and zones. As you can see in Figure 6.2, the domain namespace is partitioned into various subdomains. However, the hierarchy is also partitioned into various zones. A zone comprises the total set of RRs contained in an authoritative name server for a domain *and its subdomains*, starting from a particular point in the DNS hierarchy.

Consider, for example, the domain `microsoft.com`. This domain and its subdomain, `research.microsoft.com`, might consist of a single zone administered by authoritative name servers in the `microsoft.com` domain; that is, all the RRs for the parent and its subdomain

are contained in the same authoritative name servers. However, it is possible to delegate authority for the subdomain, `research.microsoft.com`, to another set of name servers through the use of NS and A RRs. When this delegation of administrative authority takes place, the subdomain is administered by a separate zone. The authoritative name servers for the parent domain, `microsoft.com`, do not contain records for the subdomain, with the exception of those records that are necessary to delegate authority for the subdomain to other name servers. Figure 6.3 illustrates two possible zone configurations for the `microsoft.com` domain.

Figure 6.3 Zones versus Domains



The left side of Figure 6.3 represents a *zone of authority*, which includes both the parent and the subdomains for `microsoft.com`. The authoritative name servers in the parent domain contain all the RRs for the parent and the subdomain. In the example in Figure 6.3, a name server called `ns1.microsoft.com` holds all of the RRs for the three domains: `microsoft.com`, `dev.microsoft.com`, and `research.microsoft.com`.

The right side of Figure 6.3 shows a delegation of authority from the parent domain to the `research.microsoft.com` subdomain. In this case, the name server for the parent domain does not control the records for the `research.microsoft.com` subdomain, but it does control the records for the parent and the `dev.microsoft.com` domain. A server called `ns1.research.microsoft.com` holds all the RRs for the `research.microsoft.com` domain.

Creating different zones of authority can be an efficient way of optimizing zones that contain a great many RRs. However, creating zone delegations can involve a security trade-off in that different administrators might be responsible for the servers that are authoritative for the child domains.

Configuring Delegations to Child Domains from Parent Domains

Knowing how to delegate authority to a child domain is important in implementing and maintaining a DNS infrastructure. To delegate authority from a parent domain to a child domain, the DNS servers that are authoritative for the parent domain must have NS records that identify the names of the DNS servers that are authoritative for the child domain, as well as the A records that point to the IP addresses of the DNS servers in the child domain. In our fictional example for the microsoft.com domain, the primary zone file for the microsoft.com domain would contain a set of records (NS and A records) to delegate authority for the research.microsoft.com zone:

```
research          NS          ns1.research.microsoft.com.
ns1.research      A           192.168.100.21
```

Note the lack of a trailing period after “research” on the left side of the NS and A records. The lack of a period indicates that this name is *unqualified*; that is, it is not an FQDN, and the domain name microsoft.com is implicitly appended to the left side of the name.

There are a number of ways to delegate authority to a child domain. You can enter the records manually into the zone file or you can use the New Delegation Wizard, found on the context menu of the zone that is invoked when you right-click the zone in the DNS Microsoft Management Console (MMC). A third way to delegate authority is to create a *stub zone* for the child domains on the DNS servers that are authoritative for the parent domain. If you do this, you do not need to include records to delegate authority in the zone file of the parent domain. Stub zones are a new feature of DNS in Windows Server 2003. We will discuss their use in more detail later in this chapter.

In a standard DNS environment, authoritative servers are either *primary* or *secondary servers*. (Secondary servers are sometimes referred to as *slave servers*.) The primary server has an updatable version of the flat text file that contains the RRs for the domains for which it is authoritative. The primary server is the only server on which updates to the RRs can be made. The secondary server has a read-only copy of the zone file, which is updated by a process known as *zone transfer*.

The zone transfer process is usually initiated when the secondary server polls the primary server according to a predefined interval. The secondary server reads the SOA RR and compares the version number in the record with the version number in its SOA. If the version number is higher on the primary server, it will initiate the zone transfer process and copy the zone file over TCP port 53. It is possible to configure a primary DNS server to contact the secondary DNS servers on its list when there are changes to the zone file. It is also possible to use an *incremental zone transfer* (IXFR) to copy only the changes to the zone file, rather than the entire file, but this depends on whether the DNS servers support the

IXFR protocol. If the secondary server is capable of IXFR transfers, it will request that the primary use IXFR to transfer the zone information; otherwise, it will request a standard zone transfer.

In a Windows 2000 and 2003 environment, it is also possible to store the zone information in AD rather than in flat text files. This configuration is known as an *Active Directory-integrated* zone. Updates can be made to any Active Directory-integrated zone; this is, Active Directory-integrated zones are primary DNS servers. Synchronization of Active Directory-integrated zones occurs through AD replication, rather than through the standard DNS mechanism of zone transfer. We will discuss these and other DNS server roles later in the chapter.

DNS Name Resolution Process


Distributing DNS RRs among many different zones and domains has an effect on the name resolution process that needs to occur for a DNS client to find a host name-to-IP address mapping. Let's take the example of a client trying to connect to `www.research.microsoft.com`. The DNS client is configured to use another DNS server to perform *recursion* on its behalf. (Performing recursion simply means that the DNS server will issue *iterative queries* to other DNS servers and accept referrals from these servers until it receives a positive or a negative response, and then forward that response to the DNS client.) The DNS client issues a *recursive query* to the DNS server; the DNS server subsequently issues a series of iterative queries to resolve the name. Figure 6.4 shows the process that occurs in order to resolve `www.research.microsoft.com` to the IP address.



EXAM WARNING

Don't be confused about the difference between iterative and recursive queries. Iterative queries occur when a DNS resolver asks a DNS server to perform the work of finding the answer for it. An analogy will be helpful to illustrate the concept of iterative and recursive queries. You ask your class instructor a question about AAAA records. Your instructor says he doesn't have the answer but will find out. After doing some research and asking other people for advice and direction, he finds the answer and relays it to you. In the meantime, you wait for either an informative answer to your question or a negative response. Your instructor has performed a series of iterative queries to find the answer. You have issued only a single recursive query and simply waited for a positive or negative response.

Figure 6.4 DNS Server Issuing Iterative Queries to Resolve an IP Address on Behalf of a DNS Client



The DNS client requests that dns1.shinder.net use recursion to return an answer to its query for the IP address of www.research.microsoft.com. (By default, both the DNS client and the DNS server service are configured to support this arrangement.) The DNS server first checks to see whether it can answer authoritatively from locally configured zone information. If it doesn't have the zone information, the DNS server then checks its cached information to see if it has previously answered the same query. If it doesn't have this information in cache, it then begins the process of recursion to find the answer for the DNS client.

The process of recursion begins with the contacting of the root DNS servers, which are authoritative for the top-level domain on the Internet. To find these authoritative servers, the DNS server will consult its *root hints file*, which is a list of RRs that provides information about the name servers that are authoritative for the top-level domain on the Internet. Windows 2000 and Windows Server 2003 servers will automatically install this file when you install the DNS service on your server, in most circumstances. You can also get the most current version of this file from <ftp://rs.internic.net/domain/named.root>.

Note that the root hints file is present on the DNS server *only* if the DNS server has *not* itself been configured with a root, or ., zone. If this zone is present on your DNS server, it means that this server is the highest level of authority for the root domain, and the server will not be able to perform DNS queries on the Internet. If you use the Dcpromo utility to install and configure the DNS server as a prerequisite for installing a domain controller, that utility will automatically configure the DNS server with the . zone. If you wish to use the root hints file on this server to perform recursion on the Internet, you will need to first delete the . zone from the DNS server.



NOTE

The root hints file is found in the `%systemroot%\system32\dns\cache.dns` file. By default, this file is prepopulated with the root hints for Internet servers that are responsible for resolving top-level domain names and delegating authority to second-level domains. This file can be modified directly or from the **Root Hints** tab of the DNS server property pages. On servers that are configured with a root, or ., it is recommended that this file be removed completely. In a Windows environment where you have deployed a private root, DNS servers will learn of the servers hosting the root zone and automatically update this file, as long as the TCP/IP properties are configured with the IP addresses of the root servers. You can also modify the file to reflect your DNS infrastructure.

In this example, the root DNS server is not authoritative for the .com domain, but it does contain NS records for the servers that are authoritative for this domain. It sends this information back to dns1. Then dns1 contacts a server that is authoritative for the .com domain. Again, because authority for the microsoft.com domain has been delegated to other servers, it sends the name server referral information for the microsoft.com domain to dns1. Then dns1 contacts a name server that is authoritative for the microsoft.com domain. If this server had also been authoritative for the research.microsoft.com domain, it would respond with the IP address of the requested host. However, because authority for this sub-domain has been delegated to other name servers, it sends name server referral information back to dns1, which is finally able to contact an authoritative server and receive a positive reply to its query for the IP address of www.research.microsoft.com. Once it finds this information, dns1 sends the positive reply containing the IP address information to the DNS client, which is then able to connect to the Web site.

This recursion process assumes that no information about the FQDN for www.research.microsoft.com is cached on either the DNS client or dns1. However, over a period of time, dns1 would cache information about the domain namespace and would learn the IP addresses of authoritative name servers for domains and hosts on the Internet, thereby eliminating steps and speeding up the process of name resolution. But even without cached information, DNS host name resolution is very efficient, because it will normally

use small UDP packets (512 bytes), unless the response is too large to be contained in a single UDP packet; in which case, TCP will be used.

In our example, three kinds of common responses to DNS queries are used:

- **An authoritative answer** This means that a response is sent from a server that is authoritative for the record of domain.
- **A referral answer** This means that an answer was sent back to the DNS requester that contained information not originally requested to provide hints to find the answer. For example, if the request is for an A RR, the DNS server might return a CNAME or an NS record in response to the query to help the requester find the answer.
- **A positive answer** This means that a positive response to the query is sent to the requester.

New & Noteworthy...

Using Extension Mechanisms for DNS (EDNS0) to Change the Default Size of UDP Packets Used by DNS

The original RFC for DNS (RFC 1035) limits the size of UDP packets to 512 bytes. However, Windows Server 2003 implements a more recent standard for UDP packet size (RFC 2671) that allows the administrator to configure a larger allowable UDP packet size for responses to DNS queries. When EDNS0 is configured on the DNS requester to allow UDP packets that are larger than the default size, the DNS requester sends this information to the DNS server in a query that contains an OPT RR that advertises the maximum size of the UDP packet to use in the response. When the DNS server receives this information, it will truncate the packet at the maximum allowable size specified in the OPT RR. If this information is not present, the DNS server assumes that the DNS requester does not support packets larger than 512 bytes.

Care must be taken when configuring support for EDNS0 to ensure that the UDP packet does not exceed the maximum transmission unit (MTU) packet size of any device, such as a router, that the request and response must traverse. To change the UDP packet size and EDNS0 cache settings, you must modify the Registry. For more information about EDNS0, see RFC 2671 at www.rfc-editor.org/rfc/rfc2671.txt.

A fourth possible response is a negative answer. This means that the authoritative server does not have a record for the queried name, or that it does have a record for the queried name, that is a different RR type than specified in the query.

Regardless of the answer that is returned, the results are cached so that subsequent DNS queries can be answered with nonauthoritative responses from name servers that contain the cached information. With the exception of a negative answer, the results are cached

according to the value specified for the minimum TTL in the authoritative zone's SOA RR; that is, the authoritative name server controls the TTL of the RR for cached records on DNS requesters. In the case of a negative response, this information is also cached for a period of five minutes by default to prevent unnecessary consumption of resources if the name is queried again. The period for caching negative responses is relatively short to allow the query to be resolved if the RR becomes available in the future. Negative caching is a DNS standard that is documented in RFC 2308.

It is possible to set up *caching-only* DNS servers. These are DNS servers that contain no zone information and function only to provide support for the recursion process for DNS clients. We will discuss the various DNS server roles later in this chapter.

Forward versus Reverse Lookup Zones

In most of the preceding discussion, we have focused on *forward lookup zones*. These are DNS data files that provide answers to *forward* queries that ask for the IP address of a particular FQDN. However, *reverse lookup zones* are also widely used to provide answers to *reverse* queries that ask for the FQDN of a particular IP address. For example, if you wanted to find the FQDN associated with a particular IP address, you would perform a reverse lookup against a reverse lookup zone.

To handle reverse lookups, a special root domain called `in-addr.arpa` was created. Subdomains within the `in-addr.arpa` domain are created using the reverse ordering of the octets that form an IP address. For example, the reverse lookup domain for the `192.168.100.0/24` network would be `100.168.192.in-addr.arpa`. The reason that the IP addresses are inverted is that IP addresses, when read from left to right, get specific; when the IP address starts with the more general information first. FQDNs, in contrast, get more general when read from left to right; the FQDN starts with a specific host name. In order for reverse lookup zones to work properly, they use a special RR called a PTR record, which provides the mapping of the IP address in the zone to the FQDN.

Reverse lookup zones are used by certain applications, such as NSLookup (an important diagnostic tool that should be part of every DNS administrator's arsenal). If a reverse lookup zone is not configured on the server to which NSLookup is pointing, you will get an error message when you invoke the `nslookup` command.

Security Considerations for the Presence of a Reverse Lookup Zone

Being able to make NSLookup work against your DNS servers is not the most important reason why you should configure reverse lookup zones. Applications on your internal network, such as DNS clients that are trying to register PTR records in a reverse lookup zone, can “leak” information about your internal network out to the Internet if they cannot find a reverse lookup zone on the intranet. To prevent this information from leaking from your network, you should configure reverse lookup zones for the addresses in use on your network.

For more information about security and reverse lookup zones, see <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q259922>. Note, however, that the information regarding the name of the blackhole servers in this article is out of date. The Internet Assigned Numbers Authority (IANA) has set up two blackhole servers, blackhole-1.iana.org and blackhole-2.iana.org, to handle the bogus addresses from private networks that leak onto the Internet. For more information on this topic, see Kent Crispin’s FAQ at <http://archives.neohapsis.com/archives/incidents/2002-09/0059.html>.

EXERCISE 6.01

INSTALLING WINDOWS SERVER 2003 DNS SERVICE AND CONFIGURING FORWARD AND REVERSE LOOKUP ZONES

The exercises in this chapter require that you install Windows Server 2003. You can download a 180-day evaluation copy of Windows Server 2003, Enterprise Edition, from www.microsoft.com/windowsserver2003/evaluation/trial/evalkit.msp. If you wish to preserve your current operating system, you can install Windows Server 2003 in a VMware virtual machine, which allows you to emulate a PC on which to install Windows Server 2003. You can download a 30-day evaluation copy of VMware Workstation 4.0 from www.vmware.com/vmwarestore/newstore/wkst_eval_login.jsp.

This exercise assumes that a single Windows Server 2003 server is installed as a stand-alone server and is not a member of any domain.

Before you install the DNS service, you might wish to ensure that the domain name in the FQDN for the computer name matches the domain name of the DNS forward lookup zone you plan to install. It is not a requirement that the domain name of the FQDN and the DNS forward lookup zone match. However, if they do match, you will find that Windows Server 2003 adds the appropriate records to the forward lookup zone for the DNS server. To change the FQDN for the computer, follow these steps:

1. On the Windows Server 2003 desktop, right-click the **My Computer** icon and select **Properties** from the context menu.
2. Select the **Computer Name** tab, and then click the **Change** button.
3. In the **Computer Name Changes** property pages, click the **More** button.
4. In the **DNS Suffix and NetBIOS Computer Name** property page, change the primary DNS suffix to **tacteam.local** (or a name of your own choosing) and click **OK**. Reboot the computer when prompted.

Another prerequisite for installing DNS is that your TCP/IP properties should be configured with a static IP address and the primary DNS settings should be configured to point to the address of the computer on which you are installing DNS. To configure TCP/IP properties, follow these steps:

1. On the Windows Server 2003 desktop, right-click the **My Network Places** icon and select **Properties** from the context menu.
2. In the **Network Connections** folder, right-click the **Local Area Connection** icon and select **Properties** from the context menu.
3. Highlight TCP/IP, and then select Properties.
4. In the **TCP/IP** properties page, configure a static IP address, and then configure the primary DNS server settings to point to the IP address of the server. (For the examples in this chapter, we are using addresses on the **192.168.100.0/24** network.)

After you have configured your computer with the appropriate FQDN and IP address, you can install the DNS service. There are a couple of ways you can do this. You can install the DNS service through the **Manage Your Server** page that appears when you first log on to your Windows Server 2003 computer, or you can install the service through **Control Panel | Add/Remove Programs | Windows Components**. In this exercise, we will install the service through **Control Panel**. To install the DNS service, follow these steps:

1. Select **Start | Control Panel | Add or Remove Programs**.
2. Select **Add/Remove Windows Components**.
3. In the **Windows Component Wizard** dialog box, scroll down the list of Windows components, highlight **Networking Services**, and then click **Details**.

4. In the **Networking Services** dialog box, click **Domain Name System (DNS)** to place a check mark in its box, and then click **OK**.
5. If prompted, insert the Windows Server 2003 source CD to provide the installation files for the DNS service, or enter the name of a network path to the installation files.

The DNS service is now installed on your Windows Server 2003 computer. By default, the DNS server is installed with the root hints file and will resolve queries to the Internet. If you have an Internet connection, you can verify this by using the browser on the Windows Server 2003 server and connecting to a Web site. (Alternatively, you can verify this by performing the test labeled **Perform a recursive query to other DNS servers**, which you can find in the DNS console on the **Monitoring** tab of the properties of the DNS server.)

Next, we cover the steps to add a forward lookup zone. We begin by creating a standard primary forward lookup zone:

1. Navigate to the DNS console by selecting **Start | Programs | Administrative Tools | DNS**. (You can also invoke the DNS console through the **Manage Your Server** page that is displayed when logging on to the Windows Server 2003 computer.)
2. In the DNS console, right-click **Forward Lookup Zones** and click **New Zone** in the context menu.
3. The **New Zone Wizard** appears. Click **Next**. Ensure **Primary Zone** is selected as the zone type and click **Next**.
4. Type in **tacteam.local** as the zone name, and then click **Next**. (You can also type in a domain name of your own choosing. For ease of configuration later, it should match the domain name portion of the FQDN of the computer name.)
5. Select the option to **Create a new file with this name**. (A filename has already been created based on the domain name.) Click **Next**.
6. On the subsequent page, click **Next** again to accept the default setting not to allow dynamic updates, and then click **Finish**.

We now need to verify the records in the new zone. To do this, perform these steps:

1. In the DNS console, expand **Forward Lookup Zones**, and then click the zone you just created.
2. Examine the contents of the zone on the right side of the window. You should see three records: an SOA, an NS, and a Host (A) record. If you are missing any of these records, the reason is that the domain you chose to create did not match the domain in the FQDN for the com-

puter name, or the TCP/IP configuration was not pointing to the configured IP address for the primary DNS.

We now can create a reverse lookup zone. The reverse lookup zone is used to resolve IP addresses to names. In addition, if we want to use NSLookup to query the DNS server, we need a reverse lookup zone containing a PTR RR that points to the authoritative DNS server in the zone. The domain name will be based on the IP subnet and the suffix, in-addr.arpa. In these exercises, we are using the subnet 192.168.100.0/24, so the reverse lookup domain will be 100.168.192.in-addr.arpa.

1. In the DNS console, right-click **Reverse Lookup Zones** and click **New Zone** in the context menu.
2. Follow the previous steps for creating a forward lookup zone. However, you will need to type the network ID of your network when prompted. (The **New Zone Wizard** will create the appropriate domain name based on your network ID, so do not change the order of the octets in your address. If you are following the setup for these exercises, you should type **192.168.100** as the network ID in the Wizard.)

After you have created the reverse lookup zone, examine the records that are created in it. You should see only two records: an SOA record and an NS record. Open a command prompt and invoke the **nslookup** command. You should see an error message, such as the following:

```
*** Can't find server name for address 192.168.100.21: Non-existent
    domain
Default Server:  UnKnown
Address:  192.168.100.21
```

To correct this situation, we need to add a PTR RR for the DNS server. To do so, follow these steps:

1. Right-click the reverse lookup zone you just created and select **New Pointer (PTR)** from the context menu.
2. In the **New Resource Record** dialog box, enter the host ID for the DNS server (the last number in the IP address), click **Browse**, and navigate to the A record for your DNS server in the forward lookup zone you created previously.
3. Finish creating the record. You should now have a PTR record in addition to the NS and SOA records. To verify the record is correct, invoke the **nslookup** command from a command prompt. You should see the name of the DNS server (instead of “Unknown”) in the output.

Now that you have installed a DNS server and have created forward and reverse lookup zones, you will be able to explore and examine DNS server settings. You should use the **New Delegation Wizard** to create a delegation of authority to a subdomain of the domain you just created. To create a delegation of authority from a parent domain, right-click the forward lookup zone for the parent domain and select **New Delegation**. Follow the steps presented by the Wizard.

It's obviously better if a DNS server that is authoritative for the subdomain actually exists, but if this is not the case, you can still create the records used to delegate authority. If you are able, you should install a second Windows Server 2003 server to further explore the features of DNS, such as zone transfers, stub domains, and so on. This server can be installed on a virtual machine using VMware; you can run multiple virtual machines, all of which can communicate with one another on the network.

EXAM
70-293
OBJECTIVE
2.7.1

Designing a DNS Namespace

Designing a DNS namespace is a critically important function for any business that relies on both the public and the private identities provided by the DNS namespace(s) for interaction with its customers and for the smooth and secure operation of its network. You should take some of the following considerations into account:

- **Uniqueness** Domain names on the Internet must be unique. To guarantee uniqueness of the public domain namespace, the public domain must be registered with the Internet Corporation for Assigned Names and Numbers (ICANN) through one of many authorized registrars. Although it is not a requirement that your internal domain namespace be unique, it is prudent to ensure its uniqueness.
- **Integration and interaction of public and private DNS namespaces** It is possible to use the same or different DNS namespace(s) for the public and private networks. Each of these alternatives provides different challenges. To separate the public and private zones requires both planning and administrative effort. One method of separating the public and private namespaces is to base the DNS namespace for AD on the internal network on a delegated subdomain of the public domain. Another method involves choosing a different domain suffix, such as .local instead of .com, for the private namespace that is the root of AD.
- **Security** Designing a DNS namespace should take into account the security requirements and configuration of your network. For example, it is extremely inadvisable to allow any RRs that are specific to your internal network to be publicly available through DNS queries. You should set up separate name servers to respond to queries for the IP addresses of the organization's Internet hosts, such

as Web and mail servers. Deploying a private root zone can also help to enhance the security of your DNS infrastructure. Additionally, you need to consider fire-wall placement and access rules when designing the DNS namespace. Does the security organization's security policy allow or restrict access from the Internet to internal DNS servers? In addition to considering who can query DNS servers, it is important to consider who can update RR records in the authoritative zones and how those records are updated. For example, you might not wish to allow dynamic updates in the top-level domain, but you might want to allow updates in the child domain. You would design your namespace accordingly.

- **Administration** The design of the DNS namespace will affect administration. For example, using the same domain namespace for both the private and the public networks will require, at a minimum, a split DNS configuration, where two name servers (one that is authoritative for the public RRs and one that is authoritative for the private RRs) will need to be implemented and maintained. In this scenario, special configurations might need to be implemented to allow users on the corporate network to connect to the organization's public Web servers.

Choosing the Parent Domain Name

When choosing the parent domain name to support your organization's business and infrastructure, consider whether to use or acquire an Internet domain name that is registered to your organization. If the name you choose is for use on your internal network only, you can use any name you want. However, although it is not a requirement that domain names used on your internal network be unique, it is a good idea to ensure that they are.

The best way to ensure the use of a unique domain name for the internal network is to base the domain name on one your company has registered for use on the Internet. If your organization has not registered a domain name or its currently registered name is not acceptable for use on the internal network, you should register a new domain name with an ICANN-accredited registrar.



NOTE

You can find a complete list of ICANN-accredited domain name registrars at www.icann.org/registrars/accredited-list.html.

Depending on the nature of your organization, you will want to register a domain name that has a top-level domain (TLD) name like .com, .net, .edu, or .org. You can find a complete list of top-level domains supported by ICANN, along with a description of their appropriate uses, at www.icann.org/tlds/. Sometimes, organizations will register their domain names in as many top-level domains as possible to prevent others from taking advantage of any brand recognition that the chosen domain name might possess. For

example, in addition to registering a name that has the form `mydomain.com`, you might also wish to register `mydomain.net`, `mydomain.biz`, `mydomain.org`, and so on. Furthermore, organizations that have a prominent presence on the Internet may also register common misspellings of the domain name to ensure connectivity for users who mistype the name in their browsers or e-mail clients. You should try to find a domain name that you can register with as many common top-level domains as possible. For example, if another company has already registered `mydomain.com`, but not `mydomain.net`, you might wish to expand your search and find a new domain for which you can register a `.com` extension. Many users will try a `.com` extension before trying a `.net` or other extension to reach your organization's Web servers.

Before you can register a name, you need to determine if it is unique. Most domain name registrars provide a service for determining whether a name is available for registration. However, you can also use the Whois application on the InterNic.net Web site to determine if a name has been registered and who owns the name. You can find the Whois application at www.internic.net/whois.html.

When you register a domain name, you must provide the registrar with the IP addresses and host names of one or more DNS servers that will be authoritative for your zone. This DNS server can be located on your network or on the ISP's network. In addition, many registrars offer a service whereby you can host your zone files on their DNS servers and manage these files directly (usually through a Web-based application).

Host Naming Conventions and Limitations

Regardless of the choice you make for the domain namespace of your internal and external networks, you should abide by host naming conventions and limitations. According to RFC 1123, "Requirements for Internet Hosts—Application and Support," which defines naming standards for host names, the following US-ASCII-based characters are allowed:

- Uppercase letters (A through Z)
- Lowercase letters (a through z)
- Numbers (0 through 9)
- The hyphen (-)

Note that, according to RFC 1053, DNS resolution is supposed to be case-insensitive. For this reason, the Microsoft DNS service will "downcase" any uppercase characters that it encounters to lowercase (it is an *optional* requirement that case be preserved for use with DNS; to ensure maximum compatibility Microsoft does not implement the optional requirement for case preservation). In other words, all uppercase characters will be treated as lowercase characters.

The RFC 1123 standard is a relatively old one (created in October 1989) and places limitations on non-English organizations that might wish to use an extended or non-Roman-based character set for their names. Windows 2000 and Windows Server 2003 pro-

vide support for the more recent RFC 2181, which states that any binary string can be included in a DNS name. To allow for the use of more characters than are available with US-ASCII, Windows 2000 and Windows Server 2003 DNS servers provide support by default for UTF-8, which is a Unicode transformation format. Furthermore, Windows 2000 and higher client operating systems, such as Windows XP, are UTF-8 aware.

UTF-8 is a superset of extended ASCII and additionally provides support for UCS-2, which is a Unicode character set that allows for the use of the majority of the world's writing systems. UTF-8 is backward-compatible with US-ASCII in that the binary representations of characters are identical between the two formats. However, because characters in some writing systems require more than 8 bits to represent a character, it is not possible to use character length as a means of calculating the maximum allowable length for a DNS name, which according to RFC 2181 is 63 octets per label and 255 octets per name. Because the last byte is used for the terminating dot of an FQDN, the maximum length of the name is 254 octets (bytes).

It is important to remember that not all DNS servers are UTF-8-aware. It is also possible to turn off UTF-8 support on individual Microsoft DNS servers by configuring the name-checking format in the DNS server property pages. Therefore, care must be taken in environments where not all name servers support UTF-8. In particular, when zone information is being transferred between UTF-8 and non-UTF-8 name servers, the zone can fail to reload on servers that do not support UTF-8 if the zone contains UTF-8 information.



NOTE

Even though Microsoft DNS provides support for UTF-8, it is generally a good idea, if possible, to limit host and DNS names to the US-ASCII character set supported by standard DNS to ensure maximum compatibility.

The Underscore Character

While it is legitimate to use the underscore character in NetBIOS names, the inclusion of this character in a host name is problematic in environments that use older DNS standards in which its use is prohibited. (The underscore character is allowed in domain names, however, so its use is legitimate in SRV records.) Support for UTF-8 guarantees that the underscore character can be used safely in Microsoft environments. In fact, the underscore is a reserved character that is used extensively in Microsoft DNS to identify SRV records as per RFC 2782. However, third-party standard DNS servers such as older UNIX BIND DNS servers, might not recognize host records that use the underscore. Consequently, host names, especially those used by Internet-facing servers, should not use the underscore character as a best practice. If you are upgrading a Windows NT 4 environment to Windows

Server 2003, you might wish to consider changing the NetBIOS and host names of computers whose names include the underscore character before performing the upgrade.

DNS and Active Directory (AD)

AD was introduced with Windows 2000 and is improved and enhanced in Windows Server 2003. AD is an X.500-based directory service (similar to Novell Directory Services), which stores information about users, computers, printers, and other objects that compose your network. AD also provides a consistent naming convention for users and other objects, making it easy to locate and gain access to these objects.

In addition to providing centralized control of resources and a means to either centralize or decentralize resource management, AD provides a means of logically organizing objects into administrative units. The core administrative unit of AD is the *domain*. A domain is a collection of objects that are grouped together into a single administrative unit in a common database. These objects share common security policies (for example, minimum password length). Furthermore, the domain itself is a unit of replication within AD among all the domain controllers that are members of a particular domain. There is a very close relationship between DNS and AD: the AD domain name is also the DNS domain name, which is stored in a DNS zone.

Domains are grouped into a logical hierarchy referred to as a *domain tree*. This logical hierarchy mirrors the hierarchy of the DNS namespace. When a new domain is added to the domain tree, it becomes a *child domain* of the *parent domain* to which it is added, as is the case with the DNS namespace. Furthermore, the DNS name for the new child domain is contiguous with the parent domain; that is, both the parent and child domain are part of the same DNS namespace.

Let's take the example of a root Windows Server 2003 domain that is named shinder.net. We add a child domain named corp to the parent domain. The resulting unique FQDN for the child domain is corp.shinder.net.

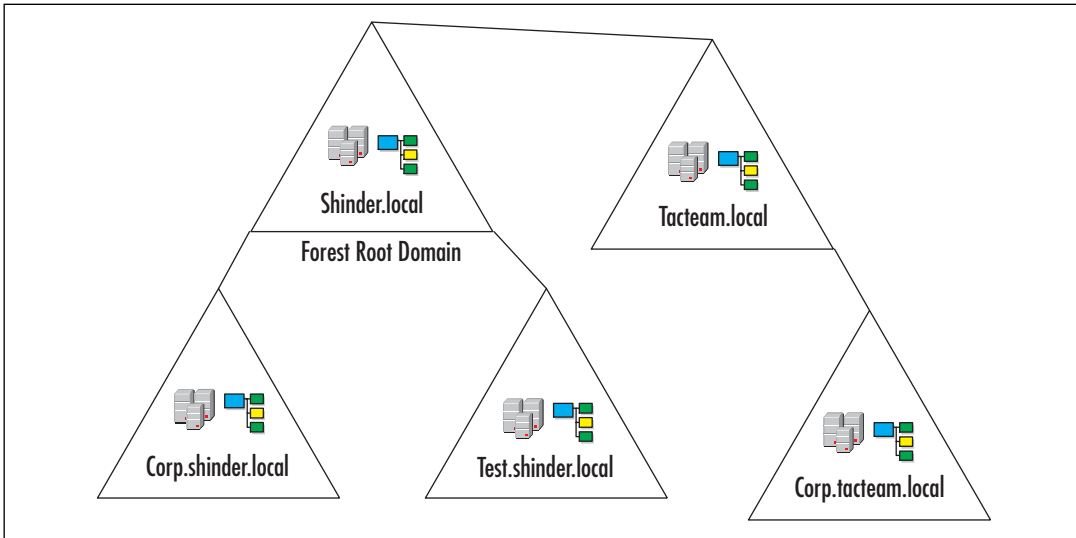
AD can comprise more than one domain tree. The resulting group of domain trees is called a *forest*. The domain trees do not share a contiguous DNS namespace. However, they do share trust relationships and a common AD schema, that is replicated to domain controllers throughout the forest. If there is only one domain tree in the forest, the subdomains in the tree are child domains of and contiguous with the *forest root domain*, which is the first domain controller installed into AD. The forest root domain and its child domains form another administrative and security boundary.

If there is more than one domain tree in the forest, the forest has a disjointed DNS namespace. That is, the namespace for the entire forest is *not* contiguous. Disjointed namespaces may require special DNS configurations in order to ensure proper name resolution throughout the forest. For example, if you have a private root zone, you need to ensure that you add delegations for your top-level domains to the root zone so that DNS requesters can find the servers that are authoritative for the appropriate domain. Alternatively, you might need to configure secondary servers and conditional forwarders to ensure name resolution.

Figure 6.5 shows the relationship between an AD forest and domain trees. Note the similarity of the AD domain names with the DNS namespace.

As noted earlier, AD has a close relationship with DNS. AD is, in fact, dependent on DNS, which is fundamental to its operation. In a Windows 2000 or Windows Server 2003

Figure 6.5 An Active Directory Forest with Two Domain Trees



network, hosts must be capable of resolving names to IP addresses using DNS.

As a prerequisite to installing AD, you must first have a DNS infrastructure in place on your network and your TCP/IP stack must be configured to use an appropriate DNS server. The DNS server must be authoritative for the domain name of your AD and must be able to support a special kind of RR known as an SRV record, which provides information about well-known network services and replaces the legacy WKS record. By default, Windows 2000 and Windows Server 2003 DNS servers provide support for these records. Other DNS servers, such as those that implement the most recent version of BIND (BIND 9 as of this writing), might support these records as well, but this needs to be confirmed beforehand if you are using something other than Microsoft DNS.

The DNS server should also be capable of supporting the following:

- Dynamic DNS (DDNS) updates** DDNS is a protocol that allows servers and DNS clients to update DNS records in the master zone file. Although it is not a requirement that the DNS server support DDNS, it is highly recommended that it do so. Support for DDNS eliminates a considerable amount of administrative work that must be performed in the form of manually adding DNS records to support AD and the network infrastructure in general. Windows 2000 and Windows Server 2003 DNS servers support DDNS, as does BIND 9.

- **Incremental zone transfers (IXFR)** When a zone file on a master DNS server is updated on a secondary DNS server, the entire file is transferred over TCP port 53 using the AFXR protocol. To eliminate unnecessary traffic associated with zone transfers, the IXFR protocol allows for the transfer of specific updated records, rather than the entire file, between master and secondary servers. The Microsoft DNS service supports IXFR, as do BIND versions 8 and 9.

If an appropriate DNS server is not available when you install your first Windows Server 2003 domain controller, the Dcpromo.exe application will prompt you to install and configure the DNS service on the computer you are promoting to a domain controller.

If you choose to install DNS through the Dcpromo.exe application, you should note that a . (root) zone will also be installed at the same time. If this zone is present on the DNS server, you will not be able to use the DNS server to resolve queries for hosts in zones for which the server is not authoritative. That is, you will not be able to use this DNS server to resolve queries on the Internet. You can correct this situation by deleting the root zone and either configuring the DNS server as a forwarder or adding the root hints file.

AD is capable of storing DNS zone information in the form of Active Directory-integrated zones. We will discuss this feature in more detail later in this chapter.

NOTE



When you install a domain controller, a file called netlogon.dns is created in the %systemroot%\system32\config folder. This file contains the SRV and other RRs required to support AD DNS resolution. You can use this file to assist in populating the zone file of a DNS server that does not support dynamic updates.

Supporting Multiple Namespaces

When you plan to use DNS for name resolution on your intranet and also plan to have a presence on the Internet, you need to consider how to support one or multiple name spaces. Assuming that you have a publicly registered Internet domain name and wish to base the internal domain name on this one, you have three choices for the selection of your internal domain name:

- **Same domain name for external and internal use** In this scenario, if your publicly registered domain is mydomain.com for use on the Internet, you use mydomain.com as your internal domain name for your intranet. This configuration requires that you manage separate DNS servers for your internal network and the external network that are both authoritative for the same domain name. This configuration is sometime referred to as a *split DNS*. However, the internal DNS servers will contain RRs that are specific to your internal network and possibly contain RRs for your publicly available Web and mail servers. The DNS servers

that are authoritative for the internal network should not be available to external clients. Depending on your security requirements and network configuration, you might find it necessary to maintain a copy of your Internet-facing servers such as your Web server on your intranet for use by your internal clients. The external DNS server that is authoritative for the domain will contain RRs for your publicly available Internet-facing servers only (such as the Web and mail servers) and will not contain RRs for your internal network. This model increases the administrative effort for managing DNS records and security, so it is not a recommended solution. However, a key advantage is that your organization's users do not need to remember different domain names for your organization's externally available servers.

- **Different namespace for internal use** In this scenario, you would use either a completely different name for the internal name of the intranet or use a domain namespace based on the registered domain name but with a different top-level domain suffix, for example, mydomain.local. Microsoft recommends using a namespace based on a registered domain name in the (unlikely but possible) event that two organizations that are using the same AD name merge. If the domain name is registered, it must be unique by definition. A key advantage of this approach is that it provides you with a unique and separate namespace for use on your internal network. With this configuration, the administrative effort required to manage the domain namespace is minimized, compared to using the same domain name for internal and external use. Also, security is enhanced and easier to manage for the following reasons:
 - The internal namespace is not exposed in the form of NS and A records used to delegate authority to the child domain in the parent domain.
 - The internal domain namespace is not reachable by clients on the Internet.
 - It is not necessary to transfer zone information between the publicly available DNS servers to internal DNS servers that might function as primary masters or secondary servers for the parent domain zone file.

A disadvantage of this option is that it requires that you manage two separate DNS namespaces, increasing administrative complexity. For example, using an unrelated internal domain name might require you to register this name with ICANN. Furthermore, using an unrelated internal domain name might cause confusion among users in your company.

- **Delegated subdomain for internal use** In this scenario, your internal domain namespace begins at a subdomain of the publicly registered domain namespace. For example, if your domain name is mydomain.com, you would use something like internal.mydomain.com for your internal namespace on your intranet. To support this configuration, you need internal DNS servers that are authoritative for the subdomain and are available only to your internal network (that is, the

child domain namespace is not accessible to external users). Your internal clients, however, would be able to gain access to both the internal and external DNS servers. This approach has a number of advantages:

- Administrative effort to maintain the DNS namespace is minimized.
- Both your internal and Internet-facing servers share the same contiguous namespace, making it easier for users to connect to these resources.
- Any DNS records used for AD are isolated in the child domain and its subdomains. The delegated child domain becomes the forest root domain for AD.

Disjointed Namespaces

Many companies have needed to deploy a disjointed namespace; that is, they design their DNS infrastructure to support two or more noncontiguous namespaces. For example, because of the high level of trust required for Domain Admins in a forest, many companies have deployed multiple forests to meet strict security requirements. In other cases, because of mergers and acquisitions, companies have needed to create Windows NT-style trusts between individual domains in the separate forests to enable resource access.

In Windows Server 2003, it is now possible to create one-way or two-way, cross-forest transitive Kerberos trusts. A two-way transitive trust simplifies resource management because it automatically enables trusts between all domains in the separate forests. This feature, along with complex business needs to deploy disjointed namespaces for separate business units, will make disjointed namespaces more common. Implementing a stable DNS infrastructure to support DNS resolution for a disjointed namespace creates challenges for the DNS administrator. For example, the DNS administrators in the separate forests might need to host secondary zones for the primary zones in the remote forests. The Windows Server 2003 DNS service includes two new features that make it easier to support disjointed namespaces:

- **Conditional forwarding** Makes it possible to configure a DNS server to automatically contact predefined DNS servers based on the domain name in the query request. Thus, when a DNS server encounters a query request for name resolution for resources in a separate namespace, it can forward this query to a particular, predefined set of DNS servers.
- **Stub zone** A concept borrowed from implementations of BIND. The stub zone is a special kind of secondary zone and consists of only a subset of records from the primary zone of the child domain: the SOA, NS, and A records that identify the DNS servers that are authoritative for the child domain. The NS and A records (sometimes known as *glue records*) are updated on the DNS server hosting the stub zone based on the refresh interval specified in the SOA record. A DNS server hosting a stub zone can respond to recursive queries and contact the DNS servers that are authoritative for the child domain, or it can respond to iterative

queries and provide referrals to the DNS servers that are authoritative for the child domain.

When a DNS server hosts a stub zone for another domain, the server can contact the authoritative servers for the domain directly when it receives a request to resolve a name query, helping to reduce DNS name query traffic and the load on the primary DNS server. Stub zones are useful in situations where authority is delegated to DNS servers in a child domain from a parent domain, such as when you are deploying your own internal root (discussed in the next section) and need to support a disjointed namespace. Stub zones remove the need to manually maintain glue records for the child domain in the parent domain. If a DNS administrator changes the NS or glue records in the child domain, this information will be updated in the stub zone, making it unnecessary for the DNS administrator in the parent domain to manually update records used to delegate authority.

These automatic updates serve to prevent a specific and common problem in a DNS infrastructure, which is known as *lame delegation*. A lame delegation occurs when the NS and glue address records used to delegate authority from a parent to a child domain are incorrect and prevent DNS servers from contacting DNS servers that are authoritative for a child domain.

NOTE



Because a stub zone is a kind of secondary zone, it is important to ensure that the zone transfer security is configured appropriately in the authoritative subdomain so that the stub zone can be replicated to the parent domain that is hosting the stub zone. By default, when you set up a primary zone, the zone transfer security allows zone transfers only to secondaries listed on the Name Servers tab. You will need to change these settings to allow zone transfers to occur to specific IP addresses, including those for the DNS servers that are configured to host the stub zone and are not listed in the Name Servers tab.

EXAM WARNING



You should know how to support disjointed namespaces and how to prevent problems arising from improperly configuring delegations of authority to other domain servers, because these are important issues in a Windows DNS infrastructure. You should be prepared for exam questions that require a thorough understanding of the challenges and solutions involved in supporting a disjointed namespace, such as the use of stub zones and conditional forwarding (which can be used as an alternative to stub zones). Additionally, you should know how to manually delegate authority from a parent to a child domain.

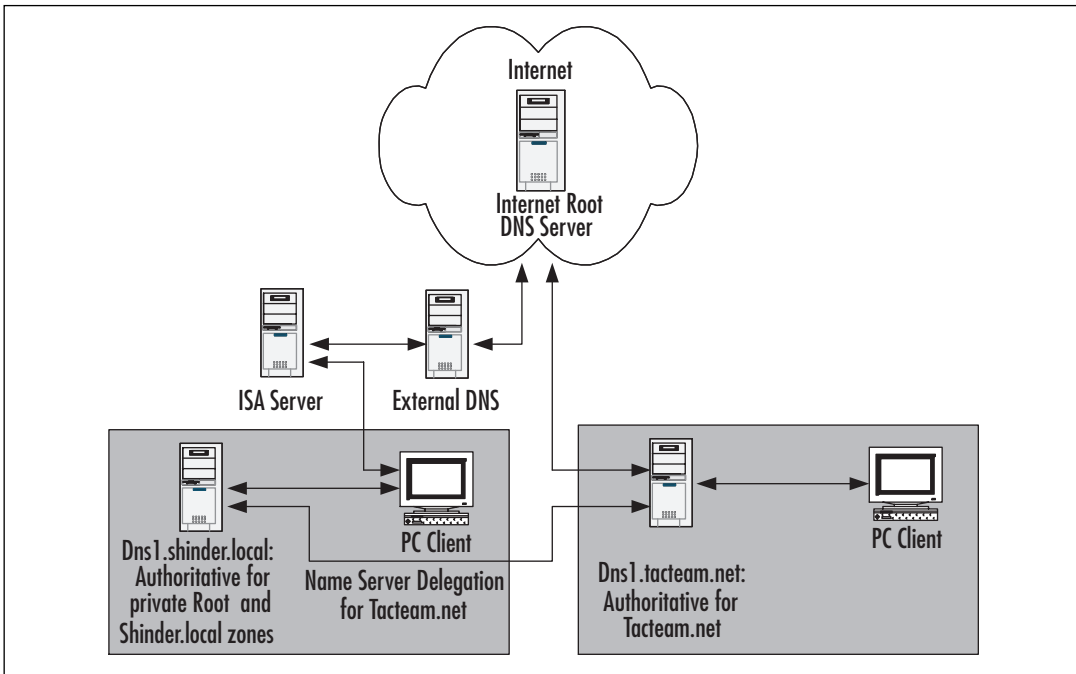
Deploying an Internal DNS Root Zone

In considering your DNS infrastructure, you should determine whether it is necessary or desirable to deploy an internal DNS root zone (the . zone). When you deploy a private root zone, you create a configuration whereby your DNS servers are authoritative for the entire DNS namespace. The private root zone contains only delegations to your internal top-level domains. Consequently, these DNS servers will not perform DNS name resolution on the Internet. If you wish your DNS servers to perform name resolution outside your organization (for example, to servers belonging to a partner or merged organization), you can add delegations from your root zone and top-level domains in the form of NS and glue A records to external DNS servers that are authoritative for other domains. In this situation, it might be advantageous to deploy a stub zone on `dns1.shinder.local` so that the NS and glue A records for DNS servers in the `tacteam.net` domain are automatically updated.

A primary advantage of this approach is enhanced security. Your DNS clients and servers that are authoritative for your DNS zones never send DNS information on the Internet. Furthermore, for large and complex networks that span WAN links, deploying a private root zone helps to simplify your DNS infrastructure.

If Internet name resolution is a requirement on your network, you might not be able to deploy a root zone. However, if your client computers are capable of using proxy servers such as ISA Server 2000, client computers can access Internet resources through the proxy server, which will perform name resolution on their behalf. The proxy server and computers that cannot use the proxy client software need to be configured to use separate, internal DNS forwarders or other DNS servers for Internet name resolution.

Figure 6.6 shows a possible deployment of an internal private root zone in combination with a proxy server to allow connectivity to external Web sites for client PCs. The figure also shows a delegation to a disjointed namespace (`tacteam.net`) to allow an internal DNS server to resolve host names on the `tacteam.net` network. Note that `dns1.shinder.local` does not perform Internet name resolution for client PCs. The ISA Server contacts a DNS server capable of performing name resolution on the Internet. However, `dns1.shinder.local`, by virtue of a name server delegation, performs recursive DNS resolution for hosts in the `tacteam.net` network.

Figure 6.6 Deployment of a Private Root Zone

In the example in Figure 6.6, a considerable amount of DNS name resolution traffic can cross a WAN link between the shinder.local and the tacteam.net networks. To reduce this traffic, you can host a secondary zone for tacteam.net on dns1.shinder.local and host a secondary zone for shinder.local on dns1.tacteam.net. In fact, in order for dns1.tacteam.net to perform name resolution for hosts on the shinder.local network, you must either host a secondary zone for shinder.local, or use some other configuration, such as conditional forwarding to make it possible for this name resolution to occur.

General Guidelines for Internal Domain Namespaces

In deciding which approach is best for your organization, take into account a number of complex factors, such as the presence of firewalls and proxy servers, client software, and the number and location of DNS servers under your control. Regardless of the approach you take, you should follow some common-sense guidelines:

- Keep it simple. Don't create a DNS infrastructure with too many subdomains (limit the number to five or fewer subdomains). As a corollary to this, try to limit the number of authoritative zones to a minimum number; don't create separate zones of authority for individual subdomains, unless it is necessary.
- Use your own company or product names, not those of another company.

- Register the domain names used by your company and base internal names on registered names.
- Avoid acronyms and geographical names that might not be easily understood.
- Don't base names on things that are likely to change, such as business units or divisions that can disappear or be renamed during the next company reorganization.
- Don't repeat names that occur on the Internet. For example, don't create a top-level domain name that already exists on the Internet, such as .ca, .biz, and so on. This will cause problems for external name resolution.
- Consider security and ease of administration—these goals might be mutually exclusive and require trade-offs.
- Use host names that are unique across your entire DNS infrastructure (keep in mind that DNS is not case-sensitive).
- Develop a convention for naming internal computers that is consistent, informative, and easily understood and remembered.
- If possible, use US-ASCII characters only for host and domain names and consider changing any NetBIOS computer names to ensure conformity with the US-ASCII character set.
- If you're using AD, make sure that the primary DNS suffix on your computers matches the AD domain name.

Planning DNS Server Deployment

Once you have determined your requirements for your DNS namespace and host names and have determined the number of subdomains, you must plan for the deployment of the DNS infrastructure on DNS servers. The goal of this planning is to ensure maximum availability, fault tolerance, currency of updated DNS records, and security, while at the same time minimizing the amount of traffic associated with DNS query and zone transfer traffic. The size and placement of zone files in your DNS topology will have a direct bearing on these considerations. Your network topology also has a direct bearing on these considerations. For example, the presence of WAN links connecting remote subnets and the available bandwidth on those links will affect the deployment of your DNS infrastructure.

Planning the Number of DNS Servers

On a simple network consisting of a single zone and relatively few hosts, you should try to deploy a minimum of two DNS servers. With two DNS servers, you ensure fault tolerance in the event that one DNS server fails or is temporarily removed from the network for maintenance.

On larger, more complex networks, you should deploy at least two DNS servers for each zone of authority you administer on your network. To reduce administrative complexity, keep the number of DNS servers you deploy to a minimum, while at the same time ensuring a high level of availability, fast query response times, and currency of records.

To reduce administrative complexity and to ensure fast query response times and fault tolerance, you can configure servers in a variety of roles. For example, you can configure *conditional forwarders* and other types of *caching-only* servers and use these in combination with DNS servers that are authoritative for particular domains. We will discuss forwarders and other DNS server roles later in this chapter.

To determine the number of DNS servers you need, you should keep the following guidelines in mind:

- A Windows Server 2003 DNS server on a 700 MHz Pentium III or higher computer with at least 256MB RAM can handle a large number of queries, more than 10,000 per second. If you experience slow response times, you can add additional DNS servers in the form of secondary servers or Active Directory-integrated zones.
- A DNS server can host many different zones—as many as 20,000 small zones that contain only a few RR in addition to the SOA, NS, and glue address records. If there is excessive traffic related to recursive queries on the network as a result of delegation to other zones, DNS servers can be configured as secondary servers to remote primary servers.
- If you have high-speed, reliable WAN links, you can use centrally located DNS servers to resolve queries for clients located in remote subnets.
- If WAN links are not reliable, you can set up a secondary DNS server on the remote network to ensure availability of zone information.
- Because DHCP servers and clients can automatically update DNS zone records using DDNS, zone replication traffic can become an issue on large networks even though Windows Server 2003 DNS supports incremental zone updates. If zone replication traffic across WAN links is a consideration, you can set up caching-only forwarders on the remote subnets to eliminate this traffic.
- DNS servers can have multiple roles. For example, a DNS server hosting a primary zone for a particular domain can be configured as a conditional forwarder for other domains. Configuring a server as a conditional forwarder allows it to build up a cache of frequent queries for host name resolution, helping to reduce DNS-related traffic for particular domains.



NOTE

When determining how many DNS servers you need, consider the importance of fault tolerance. You should never have only a single DNS server; a minimum of two is recommended, so that you will have a backup in case the primary server goes down.

Planning for DNS Server Capacity

Your DNS deployment plan will also depend on the capacity of your DNS servers to respond to queries in a timely manner and their ability to load zone files into memory. The Windows Server 2003 Resource Kit provides the following typical recommendation for a Windows Server 2003 DNS server:

- Pentium II computer running at 400 MHz
- 256MB RAM
- 4GB hard drive
- Network adapter

This should be considered a minimum configuration for a DNS server. Adding RAM or using a faster processor will increase performance, especially if the DNS server must respond to many queries or load large zone files. Adding RAM can be particularly helpful for improving DNS performance. On startup, an authoritative DNS server loads its zone files into RAM. A typical RR consumes approximately 100 bytes of RAM, although the precise value is determined by the kind of RR; for example, an SRV RR consumes more RAM than an A RR. The DNS service itself uses 4MB of RAM without loading any zones. You can use these figures to determine the amount of RAM you need to support your zone files.

You should also keep in mind that a DNS server caches query results in RAM and can return *nonauthoritative* responses to query requests from its cache. (When a DNS server performs a recursive query on behalf of a DNS client, it stores the result in cache. The next time a DNS client makes a query request for the same record, the DNS server responds with a nonauthoritative answer from its cache.) The more RAM available for caching responses, the better the performance for returning nonauthoritative answers to DNS clients on the network.

The performance of the DNS server is also influenced by the number and types of DNS queries to which it must respond. Also, a multihomed DNS server (a DNS server with more than one network interface) that is listening on more than one IP address for DNS queries consumes additional resources. If the DNS server is also a primary server, the number of secondary servers that are polling for updates of the primary zone also have an effect on performance.

Another factor that has an effect on performance is whether the DNS server is processing dynamic updates to zone files and whether the computer is also configured as a domain controller and processing secure updates to the zone files.



NOTE

In some baseline tests that Microsoft performed on a single-processor Pentium III 733 MHz computer with 256MB of RAM and a 4GB hard drive, the DNS service was able to handle 9500 queries per second and 1300 dynamic updates per second with an average CPU utilization of 75 percent. The test machine had all unnecessary services removed and was not a domain controller.

To gain a more precise understanding of the resources required for your DNS server, you can gather information from the DNS-related Performance Monitor counters that are installed with the DNS service. We will discuss the topic of monitoring DNS performance in more detail later in the chapter.

Planning DNS Server Placement

Considering where to place DNS servers, you should try to eliminate single points of failure to ensure the availability of DNS and AD services. This means that for every zone in your control, you should have at least two authoritative servers for fault tolerance. All DNS clients should be configured with the IP addresses of primary and at least one alternate DNS server to contact for name resolution. The following guidelines might assist in determining placement of your DNS servers:

- On segmented LAN environments, you should have at least two authoritative servers. These servers should be installed on different subnets.
- On a WAN, you should try to ensure that an authoritative DNS server is installed at each geographic location.
- If you are hosting an authoritative DNS for your Internet-facing hosts such as your Web and mail servers, consider hosting an offsite secondary DNS server at your ISP or on your domain name registrar's network.
- Consider which services will be unavailable if the router fails on your network segment. For example, if you have a small branch office that lacks a domain controller, users will not be able to use the services provided by AD if the router fails. In this case, there might not be any advantage to deploying a secondary server that is authoritative for your AD zones.
- Consider zone replication traffic across slow WAN links. If zone replication traffic consumes too much bandwidth, consider using forwarding servers in the remote location.

Planning DNS Server Roles

In order to properly plan, implement, and maintain a DNS infrastructure for your network, you should have an understanding of the various DNS server roles that you can install and configure.

- **Authoritative name servers** These are servers that contain the complete zone information for a domain and possibly its subdomains. Any domain will be served by one or more authoritative name servers. For purposes of fault tolerance and load balancing, there should be at least two authoritative name servers for each zone. In a Windows 2000 and Windows Server 2003 environment, it is possible to configure three types of authoritative name servers:
 - A *primary master server* is the authoritative name server that holds the updatable RRs. Any changes made to the zone file information must be made on this server. Unless you are using Active Directory-integrated zones, there is only one primary master DNS server for each zone of authority. A stand-alone server, member server, or Windows 2000 or Windows Server 2003 domain controller can be configured as a primary server.
 - *Secondary servers*, sometimes known as *slave servers*, hold a read-only copy of zone information that is transferred from the primary master server during a process known as *zone transfer* to ensure that RRs are synchronized between the secondary servers and the primary server. A zone transfer occurs in one of two ways. One way is for the secondary servers to poll the primary master server according to the refresh interval in the SOA RR and compare the version number in the SOA RR in the primary's zone file with its own. If the number is larger, it will initiate the zone transfer process. Alternatively, the primary master server can notify the secondary servers on its list whenever updates are made to the zone file. A secondary server can also be configured to do zone transfers to other secondary servers. This configuration is used primarily in situations where the polling of the primary DNS server by a large number of secondary servers puts an unacceptable load on it. The trade-off lies in currency of records, since updates from the primary DNS server must travel through more than one secondary server before all the records are synchronized among DNS servers.
 - The *Active-Directory-integrated* configuration is specific to Windows 2000 and Windows Server 2003. Instead of zone information being stored in flat text files as is the case with the primary and secondary DNS servers, zone information is stored in AD. Rather than relying on the mechanism of zone transfers, AD replication is responsible for ensuring that zone information is synchronized among all the participating DNS servers. Another key advantage of using Active Directory-integrated zones is that any DNS server that stores the zone information can update RRs; that is, more than one DNS server can

update the zone information. Secondary zones cannot be stored in AD. Active Directory-integrated zones provide enhanced security for DNS updates and zone replication traffic in several ways: all DNS servers hosting Active Directory-integrated zones must be registered in AD, AD replication traffic is encrypted, and you can use access control lists (ACLs) to restrict the hosts that are allowed to update RRs using DDNS (*secure dynamic updates*).

- **Stealth servers** When you register the name servers that are authoritative for your Internet domain namespace, you must supply at least one or two name servers that are authoritative for the zone so that authority can be delegated from the parent domain (.com, .net, and so on) to your servers. It is possible, however, for these servers to be secondary, or slave servers to a primary master server that is not listed in the registered NS records for the zone listed by the registrar as being authoritative for your domain. Usually, the primary master server is located behind a firewall, and access to the primary server itself and zone transfers to the secondary servers are tightly controlled by access rules on the firewall.
- **Caching name servers** A caching name server performs queries on behalf of DNS, but the server itself is not authoritative for any zones. When you first set up a Windows DNS server with the root hints file, it is a caching name server that can resolve queries for Internet hosts using information it possesses about the name servers that are authoritative for the root zone. After time, the caching name server builds up a list of commonly queried names in its cache, which is subsequently used to answer queries on behalf of clients.
- **Forwarding servers** A forwarding server is a kind of caching name server that sends queries to a predetermined list of name servers, known as *forwarders*, which can perform recursive queries on its behalf. The forwarding server will send its query to each forwarder in its list until it receives a positive or negative response. After it exhausts the name servers in its list, it can be configured to send requests to servers on the Internet using its root hints file. Alternatively, a forwarder can be configured to stop at this point, by disabling recursion, and send a negative response back to the original DNS requester if the forwarder cannot resolve the query. If recursion is disabled on the forwarding server, it is referred to as a *forward-only server*. There are a number of uses for forwarding servers and forwarders. They are often used when you want to tightly control which DNS servers (the forwarders) are able to send and receive DNS traffic through your firewall. Another common use of forwarders is to handle DNS queries performed across relatively slow WAN links on a corporate network. In the remote network, a name server is configured to forward queries to a more powerful caching name server that has a larger cache and is better able to resolve DNS queries as result of having access to more bandwidth, rather than send its queries directly to the Internet. A new feature of Windows Server 2003 DNS allows the configuration of *conditional forwarding*. Conditional forwarding allows the DNS administrator to

configure the forwarding server to contact specific name servers based on the domain name specified in the query. To configure a conditional forwarder, you specify the domain name and the IP addresses of the servers that are responsible for resolving host names in these domains. Conditional forwarders provide intelligent name resolution and are typically used to reduce the amount of traffic related to recursion on your network.

- **Nonrecursive servers** A nonrecursive server is one on which you have disabled recursion so that it is not able to perform recursive queries on behalf of DNS requesters. Disabling recursion on a name server also prevents it from using forwarders to resolve queries. Usually, recursion is disabled on authoritative name servers that provide name resolution for DNS requesters on the Internet, performing queries to locate your Internet hosts such as your Web and mail servers. By disabling recursion on these name servers, you ensure that the servers will respond positively only to queries for RRs in zones for which they are authoritative, and hence tighten the security of these servers. DNS clients on the Internet will not be able to configure their TCP/IP settings to point to your DNS servers for name resolution.

These name server roles are only logically separate from one another. It is possible to combine roles on a single name server. For example, a DNS server can be configured to be a primary master for one domain zone file and as a secondary for other domain zone files. However, it is often advantageous to separate these roles and place them on separate servers. By doing so, you are better able to design your DNS infrastructure to take into account the contingencies of your network infrastructure, such as the speed of your WAN links, the presence of firewalls, the need for security, and so on.

Domain Controller versus Member Server

In an AD environment, you have the choice to install and configure DNS on your domain controllers or on member servers. If you install DNS on your domain controllers, you can configure Active Directory-integrated zones.

Active Directory-integrated zones provide the following advantages over standard DNS zones:

- There is not a single point of failure for the primary zone. In a standard DNS environment, if the primary master DNS server fails and is not brought online within a particular amount of time (specified in the SOA record), the secondary servers will remove the RRs from their zone, and name resolution will fail for the entire domain.
- In large environments where DHCP servers and clients are updating RRs, this load can be distributed among domain controllers that store zone information in AD.

- Active Directory-integrated zones provide enhanced security for zone replication in that DNS servers must be registered in AD and AD replication traffic is encrypted.
- You can use secure dynamic updates with Active Directory-integrated zones to tighten security further.
- Synchronization of zone information occurs automatically through AD replication. No further configuration is necessary to facilitate transfer of zone information among participating servers.
- AD replication is more efficient than the standard zone transfer mechanisms. For example, AD replication propagates only the last changes. Even though an incremental zone transfer copies only the changes to the RRs, it propagates all the incremental changes to the RRs that have occurred since the last update. If you are not using IXFR, the entire zone file is copied whenever an update is made.
- AD replication will compress replication traffic in certain circumstances, further reducing the bandwidth needed for DNS-related traffic.

Using the Application Directory Partition for Active Directory-Integrated Zones

Windows Server 2003 enhances the design and functionality of AD through the *application directory partition*, which is a new feature of Windows Server 2003. In Windows 2000, Active Directory-integrated zones are contained in the domain partition and are replicated to all domain controllers, regardless of whether the DNS service is installed on those computers. In contrast, Windows Server 2003 installs an application directory partition on only those domain controllers that have the DNS service installed. The application directory partition allows you to confine DNS-related replication to a subset of computers that have the partition installed. By using application directory partitions, you can reduce the size of the Global Catalog and the amount of replication traffic between domain controllers. This is a significant advantage when you have a large infrastructure in which DNS or another application is making a large number of frequent updates to AD, which would otherwise flood your network with replication traffic and negatively affect domain controller performance.

When you are installing the first Windows Server 2003 AD domain controller, two application directory partitions are created by default: *ForestDNSZones*, a forest-wide partition, and *DomainDNSZones*, a domain-wide partition for each domain in the forest.

Active Directory-integrated zones can be used in combination with secondary servers. For example, you can use secondary zones on servers that are not configured as domain

controllers. This is advantageous in situations where you do not want AD traffic replicated across a WAN link, but you do want to have an authoritative DNS server available at a remote location. You cannot simultaneously load a standard text-based primary zone file and an Active Directory-integrated zone for the same domain on the same domain controller. However, you can combine primary, secondary, and Active Directory-integrated zones on the same domain controller. On a stand-alone or member server, primary and secondary zones can be combined on the same server. Furthermore, if you have multiple IP addresses bound to the server, you can emulate a secondary server on the same computer where the primary is located. This configuration is useful in very small environments where you have only one server.

EXAM
70-293
OBJECTIVE
2.7.2

Planning for Zone Replication

In planning your DNS infrastructure, you need to decide on the number and placement of your DNS servers. In particular, you must decide which servers will host zone files for your domains. Distributing zone files across your network has a number of advantages. For example, distributed zone files reduce the network traffic caused by DNS queries, increase availability and fault tolerance, provide load balancing, and result in shorter query response times. However, distributing zone files requires that you replicate zone information among your DNS servers, increasing traffic associated with zone transfers or AD replication (if you have enabled Active Directory-integrated zones). Zone files also increase the storage space requirements on DNS servers. Furthermore, replicating zone information increases the administrative effort required to maintain the DNS infrastructure.

In planning for zone replication, you must decide which mechanism you will use for zone replication: either standard DNS zone transfers or AD replication. This decision will depend on a number of factors including the storage location (file-based or AD), the type of zone information (primary, secondary, or stub), and whether you need enhanced security.

If you are using stand-alone, member servers, or other implementations of DNS such as BIND, you must use standard DNS mechanisms for zone transfers. Depending on the version of DNS or BIND you are using, you can use either full (AXFR) or incremental (IXFR) zone transfers to propagate zone information. Incremental zone transfers reduce traffic by propagating only the incremental changes since the last update.



NOTE

You cannot use IXFR on Windows NT 4 DNS servers or on versions of BIND earlier than BIND 8.2.1.

Microsoft and other DNS servers optimize traffic associated with standard zone transfers by compressing the zone transfer information and including multiple RRs in individual TCP packets. This mechanism is referred to as *fast zone transfers* (it should not be confused with IXFR). Versions of BIND earlier than 4.9.4 do not support fast zone transfers. Support

for fast zone transfers to BIND secondaries is enabled by default on Microsoft DNS servers, but it can be disabled.

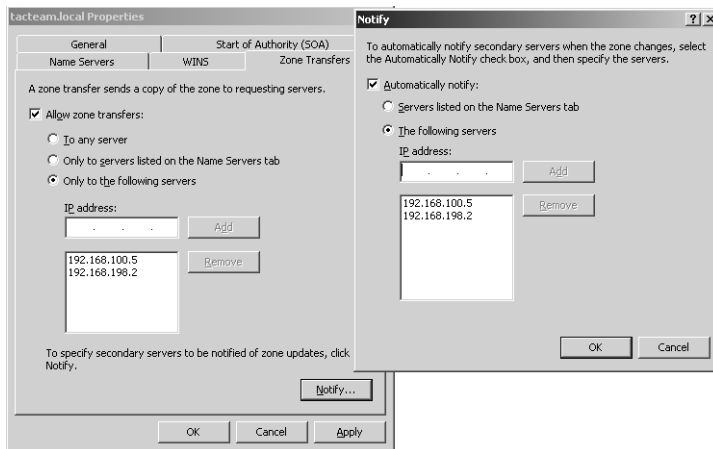
A zone transfer is initiated when the secondary servers determine that the version number in their SOA RR is lower than the version number in the primary's SOA RR, indicating an update to the primary zone. The secondary servers will compare the SOA version number in the following situations:

- When they are notified of a change by the primary server
- When the refresh interval specified in the SOA has elapsed
- When the DNS service on the secondary server is started
- When a zone transfer is manually initiated by the administrator

When the secondary server determines it needs to update its zone file, it will make a request for an incremental zone transfer (IXFR) or a full zone transfer (AXFR).

The notify list should contain only the IP addresses of secondary servers. It is not necessary to use this list to notify other domain controllers that have a copy of the Active Directory-integrated zone. Active Directory-integrated zones poll approximately every 15 minutes for updates. In fact, adding domain controllers to the notify list can actually degrade performance. Figure 6.7 shows the property pages for configuring a secondary zone transfer notify list.

Figure 6.7 Configuring a Notify List for Zone Transfers



NOTE

You should carefully consider the implications of the configuration settings for zone transfers. Configuring IP addresses in the notify list will increase the frequency and amount of zone transfer traffic on your network. If it is important that secondary servers be as up-to-date as possible, you should include their IP

addresses in the notify list. Increasing the refresh interval in the SOA RR will decrease the frequency of polling by secondary DNS servers and consequently decrease the frequency of zone transfers. If decreasing the amount of zone transfer traffic is a more important consideration than currency up to date DNS data on the secondary DNS servers, you should leave them off the notify list and increase the refresh interval. It might be desirable to do this, for example, if the secondary DNS is separated from the primary DNS by a slow WAN link.

Active Directory–integrated Zone Replication Scope

If you are using AD, you can use Active Directory–integrated zones that rely on AD to propagate zone information among domain controllers. Active Directory–integrated zones can further assist in reducing replication traffic because they replicate only the last change to RRs, rather than the incremental changes, and can compress replication traffic.

Furthermore, if all your domain controllers are running Windows Server 2003, you can further reduce this replication traffic by defining a *scope* for the replication of DNS–related information in AD. This is accomplished by leveraging a new feature of Windows Server 2003, the application directory partition (discussed previously). The broader the scope of replication, the more replication traffic that is generated.

In a Windows Server 2003 environment, you must specify an Active Directory–integrated scope. The choices for the replication scope are described in Table 6.1.

Table 6.1 Active Directory–integrated Zone Replication Scope Options

DNS Zone Replication Scope	Description and Usage
All DNS servers in the AD forest	This is the broadest scope for DNS zone replication and produces the most replication traffic. Zone data is replicated to all Windows Server 2003 domain controllers on which the DNS service is installed in the entire forest. You can use this option only when all your domain controllers are running Windows Server 2003.
All DNS servers in a specified AD domain	This is the default zone replication setting for DNS installed on Windows Server 2003 domain controllers. Zone information is replicated to all the Windows Server 2003 domain controllers on which the DNS service is installed in the domain. This option is desirable when you want to limit or restrict replication of zone information to only the domain controllers in your AD domain. Zone information is <i>not</i> replicated to Windows 2000 domain controllers.

Continued


Table 6.1 Active Directory-integrated Zone Replication Scope Options

DNS Zone Replication Scope	Description and Usage
All domain controllers in the AD domain	This option replicates DNS zone information to all domain controllers in the AD domain, regardless of whether or not the DNS service is installed on them. This option is desirable in mixed environment where Windows 2000 domain controllers are used.
All domain controllers specified in the replication scope of a DNS application directory partition	This option allows the customization of your zone replication environment. To use this option, your Windows Server 2003 domain controllers running DNS must be enlisted in the application directory partition. You can use the Dnscmd command-line utility to enlist DNS servers. The syntax for the command is dnscmd [DNS_server_name] /EnlistDirectoryPartition [FQDN of partition] . All fields are required.

A significant advantage of using the application directory partition to store zone data is that the data is not replicated throughout the AD forest in the Global Catalog. This would be the case if AD zone data were stored in the domain partition, as it is in Windows 2000. When using intersite replication (replication between different sites), the application directory partition is replicated according to the same schedule as the domain partition.

To change the replication scope, you can use the DNS console, which presents the choices indicated in Figure 6.8. There are four choices, corresponding to the descriptions in Table 6.1. The choices are to replicate zone data to all DNS server in the AD forest, to all DNS servers in the AD domain, to all domain controllers in the AD domain, and to all domain controllers specified in the scope of [a specified] application directory partition. The last choice to customize the zone replication environment is grayed out and unavailable because the server has not been enlisted in other partitions.

Figure 6.8 Changing Replication Scope for Windows Server 2003 Active Directory-integrated Zones

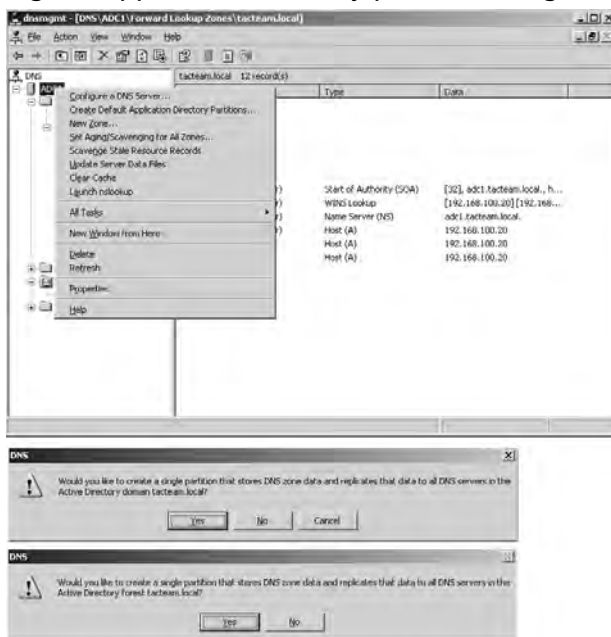


By default, when you first create an Active Directory-integrated zone and an application directory partition has not been created, you have the option of creating the partition using the DNS console utility. You can also use the Ntads utility to create or delete application directory partitions and the Dnscmd utility to create the default application directory partitions. If the default partitions have already been created, you will get an error message indicating that the partition already exists. When you use the DNS console utility to create the application directory partition, you are presented with two exclusive choices:

- To create a single application directory partition that stores DNS zone data and replicates that data to all DNS servers in the domain. If you respond **No** to this choice, you will be presented with the second choice.
- To create a single application directory partition that stores DNS zone data and replicates that data to all DNS servers in the forest. This creates the broadest scope for replication of DNS zone data.

Figure 6.9 shows the choices for creating an application directory partition using the DNS console. The two dialog boxes below the DNS console window appear when you use the DNS console to create the default application directory partitions.

Figure 6.9 Creating the application directory partition using the DNS console





NOTE

In order for the application directory partition to exist, the *domain naming master* Flexible Single Master of Operations (FSMO) role must be running on a Windows Server 2003 domain controller. In situations where you have upgraded a Windows 2000 domain controller to Windows Server 2003 and wish to change the replication scope from the domain to the application directory partition, you must first ensure that a Windows Server 2003 domain controller is the domain naming master. Otherwise, you will get an error message when you try to change the replication scope.

Security for Zone Replication

It is also important to ensure that zone replication traffic is secure, especially in situations where standard zone transfers are occurring over the Internet. To secure zone replication, you can configure Microsoft DNS to transfer zone information to only those servers that are found in the zone's name server list. However, you can further tighten security by specifying individual IP addresses that are allowed to receive zone transfers.

In situations where you are transferring zone transfer information over the Internet or you are concerned that this traffic can be intercepted, you should also consider using Virtual Private Network (VPN) tunnels or Internet Protocol Security (IPSec) to encrypt this traffic. Recent versions of BIND can use transactions signatures (TSIG) to secure zone transfers, but Microsoft does not support secure zone transfers to secondary zones. Hence the need for VPN tunnels and IPSec.


Using Active Directory-integrated zones also increases the security of your replication data by ensuring that all DNS servers are registered in AD and by using the security mechanisms inherent in AD replication. The security for zone transfers arises from the security of AD when you use Active Directory-integrated zones. Where possible, you should use Active Directory-integrated zones exclusively to improve performance and security of zone replication traffic.

General Guidelines for Planning for Zone Replication

You should keep the following guidelines in mind when planning for the distribution of zone files in your infrastructure:

- Limiting the number of zones of authority in your DNS infrastructure simplifies administration. For each subdomain that has a separate zone of authority, you must ensure that the delegation of authority is correct for the subdomain and plan for the appropriate zone replication for each of these subdomains.
- Distributing zone files increases the traffic associated with zone transfers or AD replication.

- Distributing zone files reduces the amount of traffic associated with name resolution queries.
- Distributing zone files provides a means for supporting a disjointed namespace.
- Distributing zone files increases availability and fault tolerance. It also reduces query response times.
- If you are using Active Directory-integrated zones and all your DNS servers are installed on Windows Server 2003 domain controllers, you can use an application directory partition to reduce the replication traffic associated with the transfer of zone information.
- You can minimize the bandwidth consumed by standard zone transfers by modifying the schedule for transfers to secondary zones.
- You should configure a primary server to notify only secondary servers. However, you should note that configuring the notify list to transfer zone information with the IP addresses of servers hosting the Active Directory-integrated zone can actually degrade performance.
- If you are using standard DNS zone transfers, you should try to implement incremental zone transfers and fast zone transfers where possible.
- A DNS server that is hosting an Active Directory-integrated zone or a standard primary zone can also host a standard secondary zone for another domain.
- A stub zone is a synchronized copy of a subset of an authoritative zone's RRs: the SOA, NS, and glue address records that identify authoritative name servers for a particular domain.
- A stub zone can reduce cross-domain referral and other DNS traffic.
- Security of zone data should be a consideration in your design and implementation. Active Directory-integrated zones provide more security than standard zone types. If you are using standard zone types, security can be enhanced by restricting the hosts that are allowed to receive zone transfers and by encrypting zone transfer traffic using VPN tunnels or IPsec using the strongest level of encryption possible.



Planning for Forwarding

Distributing zone files throughout your infrastructure provides one means of ensuring efficient DNS name resolution. However, it is not always desirable or possible to distribute zone files to facilitate efficient DNS name resolution.

Consider a situation in which a large company has a small branch office connected by a slow WAN link. If the branch office were to host a copy of the zone files, the zone replication traffic could overwhelm the slow WAN link. In a situation like this, it is advanta-

geous to configure a DNS server in the branch office that forwards DNS queries to specific servers in the main office. This increases the amount of name resolution traffic that crosses the WAN link, but it eliminates the more significant zone replication traffic.

It might also be advantageous in this situation to configure the DNS server in the branch office to forward all queries for Internet name resolution to a forwarder in the main office that is better able to resolve Internet queries. This forwarder can resolve queries directly by contacting authoritative DNS servers on the Internet using its more ample bandwidth and capacity, or it may be able to resolve queries from its larger cache.

A forwarder is simply a DNS server that receives queries that are *forwarded* to it by other DNS servers that are not capable of resolving the DNS query. Whenever a DNS server receives a query, it will try to answer the query from the data stored in its zone files or cache. Unless it has been configured otherwise (that is, as a nonrecursive server or a root-level server), if the DNS server cannot answer the query from its data it will either contact authoritative root servers or forward the query to a forwarder.

A forwarding server configured to use recursion if the configured sets of forwarders are unable to resolve name queries. This configuration might be desirable in situations where you want the DNS server to continue to attempt to resolve queries in the event that the forwarders are unable to do so. If the forwarders are unable to answer queries, the forwarding server will continue to use standard DNS methods to resolve the queries starting with the root-level servers. However, this configuration might not always be desirable or possible. In this case, you can configure the forwarding server to *not* use recursion if the queries to the forwarders fail. If resolution fails using the configured set of forwarders, the name resolution process stops and a negative response is sent to the DNS client.

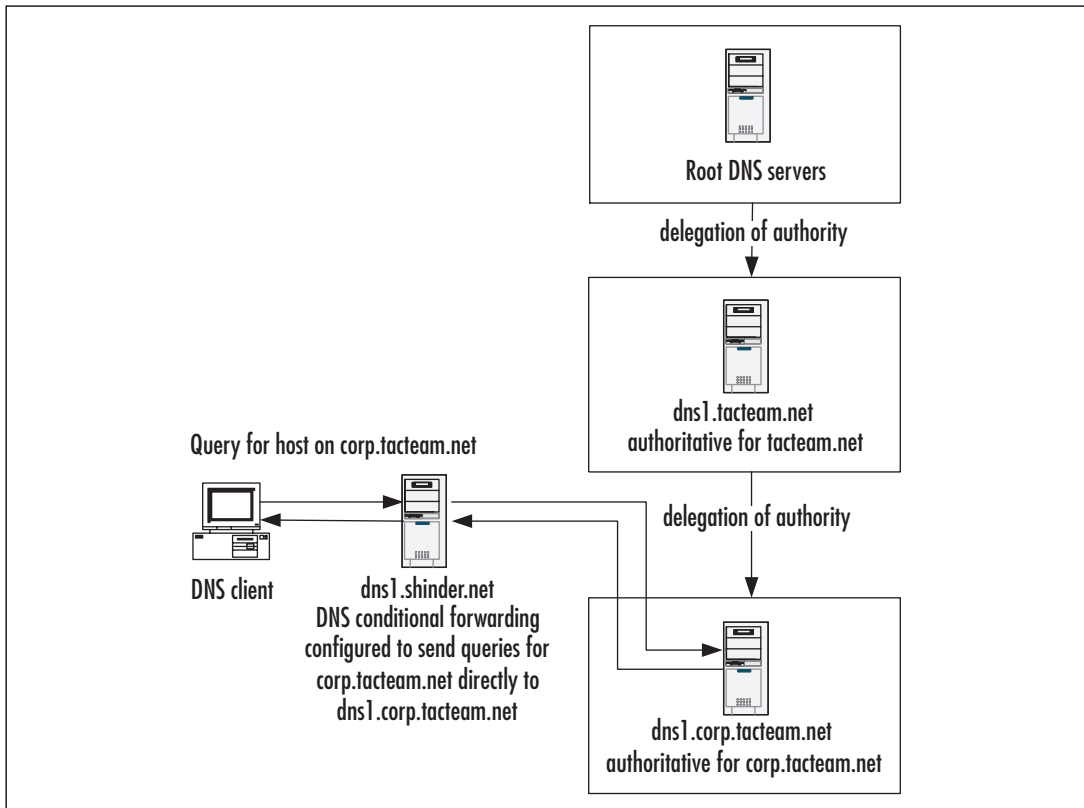
Servers that are configured to not use recursion are called *forward-only servers*. You configure a forward-only server by checking the box labeled **Do not use recursion for this domain** in the **Forwarders** property page (see Figure 6.11 in the next section).

Using forwarders can help reduce the amount of DNS traffic related to recursion in addition to reducing the traffic related to zone replication. Their use can also help to enhance security by minimizing the number of DNS servers that need to communicate with one another across firewalls. Other advantages can be realized by using conditional forwarding, a new feature of Windows Server 2003 DNS.

Conditional Forwarding

Conditional forwarding adds intelligence to the forwarding of DNS queries. In previous versions of Microsoft DNS, you could configure a forwarding server to forward queries for all domains it could not resolve to only a single set of forwarders. In this setup, the list of forwarders was responsible for resolving names for the entire domain namespace on behalf of the forwarding server. With conditional forwarding, it is possible for the DNS administrator to configure a forwarding server to contact different sets of forwarders based on the domain name in the query. Figure 6.10 shows a possible design configuration for conditional forwarding.

Figure 6.10 Conditional Forwarding Configured to Send Queries Directly to an Authoritative Server



In Figure 6.10, dns1.shinder.net has been configured to send any query requests for hosts in the corp.tacteam.net domain directly to dns1.corp.tacteam.net, which is authoritative for the zone. If conditional forwarding had not been configured, dns1.shinder.net would need to send a set of iterative queries to the root servers and dns1.tacteam.net in order to find the server that is authoritative for corp.tacteam.net. This configuration helps to eliminate network traffic related to DNS name resolution and reduces DNS query response time. Also, since dns1.shinder.net is a direct point of contact with dns1.corp.tacteam.net, over time it would acquire a significant number of cached RRs for hosts in the corp.tacteam.net domain.

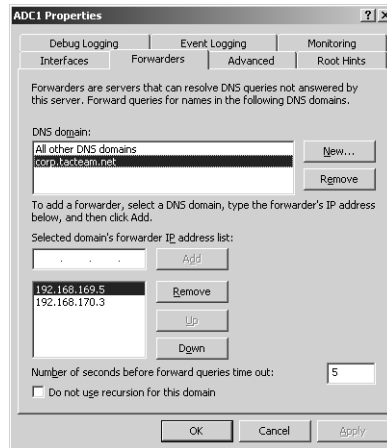
You can also imagine in this configuration that corp.tacteam.net is the forest root domain for AD. In this situation, it is both possible and highly desirable to limit DNS access through the firewall that protects the internal network for corp.tacteam.net to specific forwarding DNS servers.

As you can also infer from this scenario, using conditional forwarders eliminates the need to use secondary zones to support a disjointed namespace. It is not necessary to host a

secondary zone for the corp.tacteam.net zone on the shinder.net network. (You could also use stub zones to eliminate the need for the secondary zones, but conditional forwarding is a preferable solution.) To increase the fault tolerance of this solution, you should specify more than one forwarder in the list of servers for the forwarding server to contact to perform name resolution for the remote domain.

Figure 6.11 shows conditional forwarding configured for the corp.tacteam.net domain. Note that you can disable recursion on a per-domain basis.

Figure 6.11 Conditional Forwarding for the corp.tacteam.net Domain



General Guidelines for Using Forwarders

The following guidelines might assist you in planning to use forwarders as part of a DNS infrastructure:

- Forwarders can eliminate the need to host secondary zone files across slow WAN links that might otherwise saturate bandwidth during zone replication.
- Conditional forwarders can directly query authoritative name servers based on the domain name in the query.
- Conditional forwarders can assist in providing support for a disjointed namespace and are a preferred solution over using stub zones for the same purpose.
- Fault tolerance can be enhanced by specifying multiple forwarders and by enabling recursion if queries to forwarders fail.
- Using forwarders can enhance security by minimizing the number of DNS servers that need to communicate with each other across firewalls.



EXAM WARNING

Conditional forwarding is an important new and useful feature of Windows Server 2003 DNS. You should be familiar with configuring conditional forwarding and understand the reasons that conditional forwarding is a preferable solution in a given environment.

DNS/DHCP Interaction

As is the case with Windows 2000, Windows Server 2003 supports the DDNS standard (RFC 2136) to dynamically update both forward and reverse lookup zones with A and PTR RRs, respectively. (A forward lookup zone resolves host names to IP addresses; a reverse lookup zone resolves IP addresses to host names.) DDNS reduces much of the administrative burden in managing a zone files in a DNS infrastructure. In particular, DDNS makes it possible for AD domain controllers to create and update the SRV RRs that are fundamental to the proper operation of AD. DDNS is also used in combination with DHCP to ensure that DHCP clients will have the appropriate records registered for them in DNS and the DNS records are updated whenever IP addresses change or DHCP leases expire.

Both clients and DHCP servers are capable of updating the zone records. However, only clients that are running Windows 2000, Windows XP, or Windows Server 2003 operating systems are capable of directly updating DNS zones. This is the default configuration for these clients and can be disabled on the **DNS** tab of the **Advanced** property page for TCP/IP. Usually, DHCP clients will update their own A records in the forward lookup zone, but the DHCP servers will update the PTR record in the reverse lookup zone (the computer “owns” the host name, but the DHCP server “owns” the IP address). Clients with manually configured IP addresses will always try to register both an A and a PTR record. Other level clients, such as Windows 9x and Windows NT 4, must rely on DHCP servers to update both A and PTR RRs on their behalf.

When a client or a DHCP server attempts to update an RR, it will first query the DNS server that it is configured with to find the DNS server that is authoritative for the domain name it is trying to register. Once it determines this information, the DNS client will send an update request to the server that is authoritative for the zone. If the update request meets the prerequisites for updating the record, the record is updated. If the prerequisites are not met, the update fails. The client is notified of either the success or failure of the update. In the case of failure, the DNS client will attempt to register the record again in a 5-, 10-, and then a repeated 50-minute interval.

DHCP clients that are capable of dynamically updating DNS records use the DHCP client option 81 to provide the FQDN as specified by the full computer name in the properties of the **My Computer** object, and instructions for the DHCP server to handle DDNS registration. (This is configured on the **DNS** tab of the **Advanced** property page

for TCP/IP of the client computer.) The client's FQDN is used to register the name with the appropriate DNS server that is authoritative for the zone. Other level clients will be registered with DNS servers that are authoritative for the domain name configured for the DHCP scope.

DHCP Client and Netlogon Service

The ability for a Windows XP and Windows 2000 Professional client to update a DNS record requires the DHCP client service to be running. The DHCP client service, rather than the DNS client service, is responsible for sending dynamic update requests to the primary DNS. The reason for this is to ensure that updates to the zone file occur whenever there is a change in the IP address associated with a computer as a result of DHCP. This is true regardless of whether or not the client is configured to acquire its TCP/IP configuration from a DHCP service or has a static TCP/IP configuration. When a client creates a DNS registration, it will use a default value of 20 minutes for the TTL on record, which overrides the min TTL value in the SOA record. Using the DHCP client service, DNS clients will send an update request opcode every 24 hours for their A and PTR records. If there is no change to the name and IP address mapping, this update request is considered a refresh and does not result in a change to the version number of the DNS zone file.

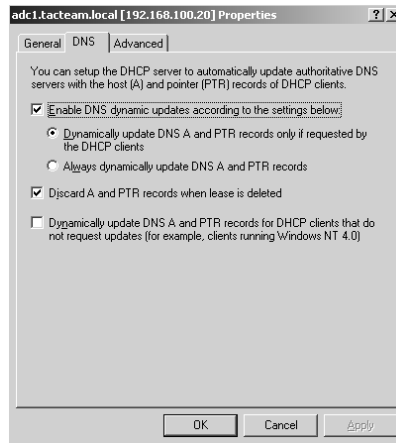
The situation for servers and domain controllers is a little different, owing to the importance of having accurate DNS data for these computers. These computers send an update request every hour. If the computer is a domain controller, the Netlogon server is responsible for sending the update every hour. A, PTR, CNAME, and SRV records. (In the case of Windows 2003 domain controllers the update interval is every 15 minutes.)

For more information about this topic, see the Microsoft Knowledge Base article "How to Enable/Disable Windows 2000 Dynamic DNS Registrations" at <http://support.microsoft.com/default.aspx?scid=kb;en-us;246804>.

A DHCP server will do the following, depending on its configuration:

- Update the A and PTR records, if requested by the client.
- Always update the A and PTR records, regardless of the client request.

A DHCP server will attempt to update A and PTR records if requested by the client. Figure 6.12 shows the default configuration for the DHCP server on the properties page for the DHCP server in the DHCP console. A similar property page exists for the DHCP scope.

Figure 6.12 Default DHCP Configuration for Dynamic DNS Updates

To configure the DHCP server to update DNS records, regardless of the client request, you can select the radio button labeled **Always dynamically update DNS A and PTR records**. If you wish to configure DHCP to perform DNS updates on behalf of legacy clients, you can select the check box labeled **Dynamically update DNS A and PTR records for DHCP clients that do not request updates (for example, clients running Windows NT 4.0)**. By default, the DHCP server is configured to remove both the A and the PTR records from the DNS zone. You can change this behavior by clearing the box labeled **Discard A and PTR records when lease is deleted**. When you clear this box, the DHCP will attempt to remove the PTR record when the lease expires.

Security Considerations for DDNS and DHCP

Implementing DDNS creates some security risks in that unauthorized computers and users might be able to update DNS records. In the case of public Web servers, the consequences of the unauthorized registration of a rogue Web server IP address to replace a valid one can be very significant indeed. For this reason, it is not a good idea to enable DDNS on any zones that are used to resolve names for your Internet-facing servers.

To mitigate the risk of unauthorized updates, you can require the use of *secure dynamic updates*. However, the option to use secure dynamic updates is available only if you are using Active Directory-integrated zones. (On a standard primary zone, you have two choices for security: secure and non-secure.) When you enable this option, you are able to control which computers, users, or groups are able to modify RRs in the zone. For this reason alone, you should consider the using DDNS only if you are using Active Directory-integrated zones.

If you have enabled secure updates, there is a potential for problems caused by the ownership of records. When a DNS client or a DHCP server updates a zone file with an RR, it becomes the owner of that record. Normally, this does not create a problem.

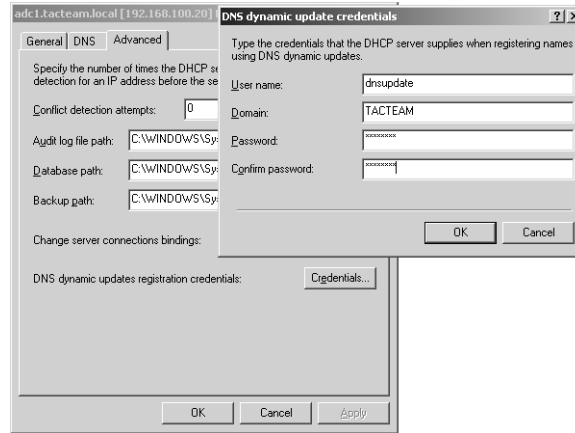
However, in some circumstances, the ownership of an RR can prevent a valid update to it. Consider the case of a client that is upgraded to Windows XP. After the upgrade, it attempts to update the RR in the zone. The attempt will fail because the record is owned by the DHCP server that originally created the record on the client's behalf. Or, consider the case where a different DHCP server other than the original one, tries to register an update on the client's behalf. Again, the attempt will fail. To resolve this problem, you can use a special security group called `DnsUpdateProxy`.

DnsUpdateProxy Group

Any objects that are created by members of the `DnsUpdateProxy` group have no security and are ownerless. Consequently, the first authenticated computer that updates the record is able to take ownership of the object. Therefore, if you enable secure dynamic updates only, you should place all DHCP servers in this group before they start registering names.

The `DnsUpdateProxy` group can create a security risk, however, if the DHCP server is installed on a domain controller. If the DHCP server that is a member of the `DnsUpdateProxy` group is installed on a domain controller, all the SRV, the A records for domain controller on which DHCP is installed and other critical records created by the domain controller for AD functionality will be ownerless, allowing the first authenticated user who tries to update them to become the owner. For this reason, you should not install a DHCP server on a domain controller if you are using the `DnsUpdateProxy` group.

If, for whatever reason you do need to install DHCP on a domain controller, or if DHCP is updating A records for clients in forward lookup zones, you should configure your DHCP server(s) to use DNS dynamic update credentials. To do this, you configure a security principal (a user account in this case) for use by all your DHCP servers when they update a DNS zone. You then configure your DHCP servers to use this account for dynamic updates. (This is a new feature of Windows Server 2003 and is not available on Windows 2000.) This obviates the problems arising from ownerless records created by DHCP servers in the `DnsUpdateProxy` group. In particular, enabling this configuration prevents a DHCP server from using the elevated permissions it inherits by virtue of its being installed on a domain controller. Figure 6.13 shows the **Advanced** tab on the DHCP server property page where you configure credentials for dynamic updates.

Figure 6.13 Configuring Credentials for DHCP Updates to Dynamic Zones

Head of the Class...

Generic Security Service TSIG (GSS-TSIG) and Dynamic Updates

Microsoft uses a dialect of transaction signatures (TSIG) as the underlying mechanism for secure dynamic updates, as specified in RFC 2485. This dialect, Generic Security Service TSIG (GSS-TSIG), is not spoken by other implementations of DNS. A version of BIND 9.x is supposed to provide this support in the future, but as of this writing, BIND 9.2 (the most current version) does not provide this support. This lack of interoperability can cause issues if you are trying to integrate BIND into your Windows environment. For example if you want a BIND server to handle all your dynamic updates, which makes the zone become a much more complex administrative challenge, as well as if you want a BIND DNS client to be able to update records using secure dynamic update.

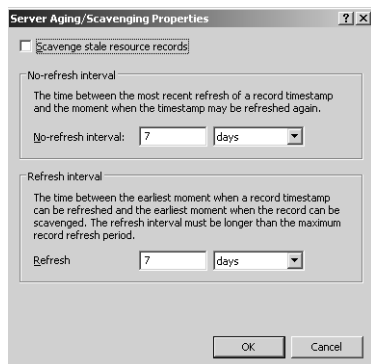
In BIND 9, TSIG is used primarily for secure server-to-server communications (for example, zone transfer, notify, and recursive query messages). However, TSIG can be used in a BIND environment for secure dynamic updates.

Aging and Scavenging of DNS Records

When you enable zones for dynamic updates, it is possible that the zone data files will acquire a large number of superfluous and outdated records that might have a negative effect on DNS performance. For example, if you retire a user's workstation and disconnect it from the network, the RRs for that computer might remain in the DNS data. To help ensure the integrity and currency of DNS data, you can enable aging and scavenging of outdated DNS records. (By default, the aging and scavenging option is not enabled.)

Aging and scavenging can be set on a per-zone or per-DNS server basis. Per-zone settings override per-DNS server settings. Figure 6.14 shows the server-wide aging and scavenging property page.

Figure 6.14 Aging and Scavenging Settings for a DNS Server



The **No-refresh interval** setting is the amount of time that must elapse before a DNS client or DHCP server can refresh a timestamp for a record. When a DNS client creates a record, it is assigned a timestamp. The DNS client attempts to refresh this record every 24 hours. Unless the record is changed (for example, the client receives a new IP address), the timestamp cannot be refreshed for a default period of seven days. After the seven days have elapsed, the DNS client can refresh the timestamp, which starts the timer on the no-refresh interval for the record. If the record is not refreshed in the seven-day period, it can be scavenged. When the record is scavenged, however, depends on another setting, the **Scavenging period**. This setting is enabled and configured on the **Advanced** tab of the property pages for the DNS server. To enable scavenging, you must enable this setting, as well as the settings for **No-refresh interval** and **Refresh interval**.



EXAM WARNING

DDNS and its interaction with DHCP are important concepts. You should be thoroughly familiar with the implementation of DDNS and DHCP to support dynamic updates to DNS zones. Your understanding of these concepts should also be informed by a thorough understanding of the security implications for enabling DDNS.

EXAM
70-293
OBJECTIVE
2.7.5

Windows Server 2003 DNS Interoperability

In addition to its interoperability with DHCP, the Windows Server 2003 DNS is designed to interoperate with other implementations of DNS such as BIND, and with other

Windows Server 2003 services such as WINS. In this section, we examine the interoperability of Windows Server 2003 with other DNS servers and Windows Server 2003 services.

BIND and Other DNS Server Implementations

One of the design goals of Windows 2000 and Windows Server 2003 is to ensure that they conform as much as possible with TCP/IP and other standards, as defined by various organizations and governing bodies. This, in turn, helps to ensure that Windows can interoperate with a wide variety of heterogeneous systems.

With some exceptions, such as the addition of functionality required for the interoperability of DNS and WINS, Windows Server 2003 DNS is a completely standards-based implementation of DNS. As such, it will interoperate with other standards-based implementations of DNS, such as BIND. In fact, in many cases, it is not necessary to forsake a current implementation of DNS for Windows Server 2003 DNS as long as the implementation of DNS supports current DNS standards. That said, management of your DNS infrastructure is easier if all your DNS servers are Windows Server 2003 servers.

The degree of interoperability will depend on the version of BIND with which the Windows Server 2003 DNS server interacts. Like other standards, the standards for DNS are evolving, and earlier implementations of DNS such as the DNS in Windows NT 4 or earlier versions of BIND will not interoperate completely with Windows Server 2003 DNS. In some cases, the presence of downlevel and legacy implementations of DNS can create problems in the DNS infrastructure.

BIND stands for Berkeley Internet Name Domain and was developed by a group of graduate students at University of California at Berkeley in the mid-1980s for use on UNIX operating systems. BIND is now the responsibility of the Internet Software Consortium (ISC). The ISC's first release was BIND 4.9.3. BIND 8 was released in 1997. BIND 9.2 is the most current version as of this writing. BIND 8 is still widely used. The latest version is 8.4.1, and it should be implemented because it fixes a number of security holes and bugs with earlier versions. Version 4 of BIND has been officially deprecated by ISC, and its use is not recommended. However, if BIND 4 cannot be upgraded to BIND 8 or 9, you should upgrade to BIND 4.9.11.

Table 6.2 shows a comparison of features support by various implementations of DNS.

Table 6.2 Windows DNS and BIND Compatibility Comparison

Feature	Windows Server 2003	Windows 2000	Windows NT 4	BIND 9.2	BIND 8.4.1	BIND 4.9.3
RFC 2782–SRV RRs	Yes	Yes	Yes, with Service Pack 4 or higher installed	Yes	Yes (minimum version is BIND 8.1.2)	No

Continued

Table 6.2 Windows DNS and BIND Compatibility Comparison

Feature	Windows Server 2003	Windows 2000	Windows NT 4	BIND 9.2	BIND 8.4.1	BIND 4.9.3
Fast zone transfer	Yes	Yes	Yes	Yes	Yes	No (but is supported in versions of BIND later than 4.9.4)
Incremental zone transfer	Yes	Yes	No	Yes	Yes (but not supported in versions of BIND earlier than 8.1.2)	No
Dynamic updates	Yes	Yes	Yes	Yes	Yes	No
Stub zones	Yes	No	No	Yes	Yes	Experimental
Conditional forwarding	Yes	No	No	Yes	No	No
DNSSEC	Limited support to allow loading of DNSSEC RRs in secondary zones	No	No	Yes	Yes	No
ACLs on RRs	Yes, if using AD-integrated zones with secure updates only	Yes, if using AD-integrated zones with secure updates only	No	No	No	No
GSS-TSIG for secure dynamic updates	Yes	Yes	No	No (Support for only simple secure updates, as per RFC 3007)	No (support for only simple secure updates, as per RFC 3007)	No
TSIG for securing zone transfers and notify messages	No	No	No	Yes	Yes	No

Continued

Table 6.2 Windows DNS and BIND Compatibility Comparison

Feature	Windows Server 2003	Windows 2000	Windows NT 4	BIND 9.2	BIND 8.4.1	BIND 4.9.3
Kerberos for secure zone transfers	Yes, when using AD-integrated zones	Yes, when using AD-integrated zones	No	No	No	No
WINS and WINS-R records	Yes	Yes	Yes	No	No	No
UTF-8 character encoding	Yes	Yes	No	No	No	No
Aging and scavenging of RRs	Yes	Yes	No	No	No	No

Zone Transfers with BIND

BIND supports standard primary and secondary DNS zones. Thus, BIND servers can be used as both primary DNS servers that transfer zone files to Microsoft DNS secondary servers and vice versa. A BIND server can also be configured as a secondary server to an Active Directory-integrated zone. However, an Active Directory-integrated cannot be a secondary zone, so it is not possible for a BIND server to host a primary zone that transfers zone information to a secondary zone configured in AD. Also note that if you want to secure zone transfers between BIND and Microsoft DNS servers, you will not be able to use the TSIG mechanisms available to recent implementations of BIND.




NOTE

To secure transfers of DNS zones, you must either implement zones that are exclusively Active Directory-integrated or use some other mechanism, such as VPN tunnels and IPsec if you are using standard DNS zones.

Versions of BIND earlier than BIND 4.9.4 do not support the fast transfer method for zone replication. When the fast transfer method for zone replication is enabled, multiple zone RRs are compressed in the TCP/IP packet. Fast zone transfers are enabled by default in Windows DNS. You should disable fast zone transfers only if your secondary DNS servers are running versions of BIND earlier than version 4.9.4. The configuration for fast zone transfers can be enabled or disabled only on a server-wide basis. You cannot enable or disable it on a per-zone basis. Disabling fast zone transfers does not affect zone replication between Windows DNS servers. Figure 6.15 shows the default configuration that enables

fast zone transfers. To disable fast zone transfers for BIND secondary servers, navigate to the **Advanced** tab of the property pages for the DNS server and clear the check box for **BIND secondaries**.

Figure 6.15 Enabling Fast Zone Transfers for BIND Secondaries



Windows DNS zone files can contain RRs that can cause problems for BIND secondaries. These records include those that use an underscore in the host or domain name and the WINS and WINS-R records. On some versions of BIND, notably BIND 8.0, the presence of these records can cause the zone to fail to load.

Although the underscore is a valid character in a NetBIOS name, it is not a valid character for DNS host names, according to RFCs 851, 952, and 1123. (The underscore is a valid character for domain names, and the more recent RFC 2181 specifies that any binary string can be used to represent a host name, but not all DNS servers conform to the standards specified in RFC 2181.) BIND version 8, in particular, will have problems if it encounters underscores in the host or domain names when it loads the data for the secondary zone. This is a result of a feature in BIND 8 known as name checking, which restricts the character set used for host and domain names. If underscores are present in host names, you have two choices: rename the computers so that their names do not have underscores, or disable name checking on the BIND 8 server by changing the default **check-name** setting on the BIND 8 server from **Fail** to **Warn** or **Ignore**.

If a BIND 8 server is hosting a primary or secondary zone for AD SRV records, the only choice is to disable name checking, because these records contain underscores in the domain names, and these cannot be changed. (BIND 9 does not restrict the character set for domain names, so this is not an issue if you are running BIND 9.)

The proprietary WINS forward and reverse lookup records also create problems for BIND secondaries. In this case, the issue is caused by the deed WINS record is not part of the DNS standard and not recognized by other DNS servers. Non-Microsoft DNS servers will see the WINS forward and reverse lookup records as bad records, causing either data

errors or the failure of the zone to load. If you are using BIND secondaries for a zone hosting WINS records, you have two choices: configure the WINS records not to replicate or configure a separate referral zone for WINS records. It is preferable to configure a separate referral zone for WINS records, because clients who contact secondary DNS servers might get different answers from those clients who contact the primary DNS server. We will discuss WINS and DNS interaction in more detail later in this chapter.

Supporting AD with BIND

As we mentioned earlier, you can support AD using BIND servers rather than Windows Server 2003 DNS. The minimum requirement for a DNS server to support AD is that it be able to host SRV records in its data. DDNS is only an optional requirement for a DNS server. Thus, a Windows NT 4 DNS with Service Pack 4 or later could be used to support AD records.

To host AD records, the minimum version of BIND that must be used is version 8.2.2 patch 7. If you use BIND 8, you must configure the **check-name** setting to **Ignore** so that it will load a zone containing underscores in domain names. This setting is not necessary on BIND 9 servers because they do not restrict character sets used for domain names.

Both BIND 9 and BIND 8.2.2 are capable of supporting dynamic updates. To allow domain controllers to dynamically register their DNS data, you can configure the **allow-update** setting in the `named.conf` configuration file on the BIND servers. However, it is not possible to configure ACLs on individual RRs (as it is when you are using Active Directory-integrated zones configured for secure updates only).

BIND might be uncomfortable, for security and other reasons, with allowing dynamic updates in the master zone file that hosts the DNS records currently in use. The **allow-update** setting allows you to specify the IP addresses of the servers that can dynamically update records in the zone. However, IP addresses can be spoofed, so this isn't a very strong level of security.

NOTE



Secure dynamic updates can be configured for secure zones hosted on BIND servers by using DNSSEC, as per RFC 3007. However, because many of the standards that govern secure updates and related issues are in the immature stages of being developed as officially accepted standards, Microsoft chose not to implement the same standards as BIND. (There is currently no single IETF standard for secure dynamic updates that addresses interoperability of the various mechanisms for secure updates.) Thus, Windows clients and DNS servers are not able to use DNSSEC mechanisms to provide secure dynamic updates for zones hosted on BIND servers.

One way to mitigate the risk of using BIND servers for dynamic updates is to create subdomains to host the AD DNS data. For example, if the domain name is mycompany.com, you can create a separate zone called ad.mycompany.com. To create this zone, you must issue a zone statement specifying the zone name and the location of the files in the named.conf file on the BIND server. However, Microsoft Active Directory-integrated zones still provide a much higher level of security. For this reason, it is preferable to use Active Directory-integrated zones. BIND administrators can delegate authority to a subdomain hosted in Active Directory-integrated zones and configure BIND servers as secondaries to this zone to enhance fault tolerance and availability.

Split DNS Configuration

Many organizations want to use the same name on their internal network as they do on their publicly available external network. For example, suppose that a company's name is mycompany.com and its Web server and e-mail servers located in the DMZ use this domain name in their FQDN. The company also wants to use this name for its AD domain on the internal network. This situation creates a number of challenges. Foremost among these is security of internal DNS records. It is not desirable to expose internal host names and IP addresses to external clients, even if these hosts cannot be reached by external clients because of restrictions on the firewall. Also, it is not a recommended DNS best practice to include any record in a zone file for a host that is unreachable.

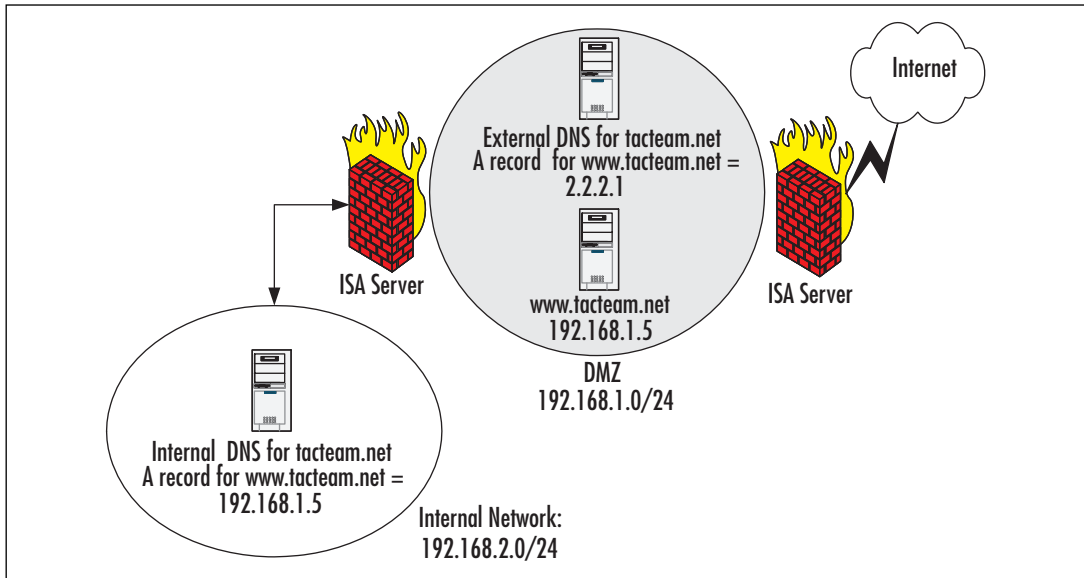
At a minimum, a properly secured DNS configuration requires that the DNS records for the internal namespace be accessible to internal clients only and not accessible to external clients. Furthermore, internal clients should be able to resolve queries for external hosts on the Internet so that e-mail servers are able to send mail to external hosts and users are able to connect to the Internet. Finally, the bastion hosts (computers that can communicate with both the Internet and the intranet) that are responsible for delivering e-mail to the internal network should be able to successfully locate and communicate with the appropriate internal servers through the firewall.

This situation implies the use of a *split DNS* configuration. A split DNS configuration requires two sets of name servers for the same namespace. For example, suppose that a set of DNS servers in the DMZ contains records for the hosts, such as the A records for the Web servers, the MX and A records for the mail servers, and the NS and A records for the DNS servers in the DMZ. Another set of authoritative DNS servers that contains records for internal hosts is configured in the internal network for the same namespace. The DNS servers configured on the internal network are not accessible to external clients for name resolution.

Internal clients should also be able to gain access to the company's publicly available Web servers. Depending on the configuration of the infrastructure, this can create some challenges. For example, if the company is using ISA Server as its firewall and making a Web server in its DMZ available to external clients via Web server publishing rules, internal clients might not be able to connect to the internal Web server if the internal DNS uses an A record for the Web server that points to an external address. Supporting this kind of con-

figuration requires that the internal DNS servers use A records that point to the internal IP address of the Web server and not the external IP address that is used to publish the Web server for external clients. In other words, the A records for the Web server will differ in the internal and external DNS servers that are authoritative for the zone. Figure 6.16 shows a possible configuration for a split DNS to allow internal clients to connect to the publicly available Web server.

Figure 6.16 Split DNS Configuration to Allow Internal Clients to Connect to the Web Server in the DMZ



NOTE

Supporting a split DNS configuration involves more effort on the part of the DNS administrator. For example, the DNS administrator might need to manually update separate DNS servers that are authoritative for the same zone. In addition, the DNS administrator must ensure that no records for the internal network appear in the publicly available DNS server.

Interoperability with WINS

In a mixed environment that includes downlevel clients such as Windows NT 4 and Windows 95, you must continue to support NetBIOS name resolution. The primary mechanism for supporting NetBIOS name resolution in a segmented network is through WINS, which allows clients on different subnets to register and resolve NetBIOS computer names on WINS servers. In some situations, it might be necessary for UNIX clients, which do not

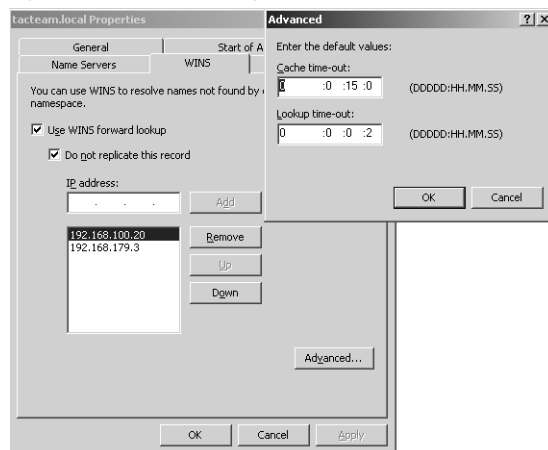
support NetBIOS, to connect to Windows NT 4 computers. In order to resolve the Windows NT 4 computer names, the UNIX hosts must use DNS. However, if the Windows NT 4 server is configured with a static IP address, it will not be able to dynamically register its host name and IP address in DNS.

One way to support DNS resolution for NetBIOS computer names is to integrate WINS with DNS through WINS forward and reverse lookup records. When a DNS zone is configured with WINS forward or reverse lookup records, it will consult a WINS server to resolve host names for records that are not present in its zone data.

For example, suppose that a UNIX host needs to send a print a job to Windows NT 4 server named PServer1.tacteam.local. The UNIX host sends a query for PServer1.tacteam.local to the DNS server authoritative for the tacteam.local zone. The DNS server does not find a record for PServer1 in its zone data, so it performs a WINS lookup to the IP address of the server listed in its WINS forward lookup record. After receiving a reply from the WINS server, it sends the information to the UNIX host. The DNS server that performs the NetBIOS resolution will keep the record in its cache for a configurable interval (the default is 15 minutes), so that if it receives a query for the same name within the interval, it can resolve the name from its cache.

As a result of this integration with WINS and DNS, it is not necessary for the DNS administrator to manually update the DNS zones with A records for NetBIOS computers that are incapable of updating DNS data on their own. The configuration of WINS forward and reverse lookup records is performed on a per-zone basis. To configure WINS lookup records, go to the forward or reverse lookup zone for which you wish to configure WINS integration, go to the property pages for the zone, and click the **WINS** tab. Figure 6.17 shows the **WINS** tab property pages.

Figure 6.17 WINS tab for a DNS Forward Zone
Showing Advanced Configuration Options



There are a few things to note about the configuration shown in Figure 6.17:

- Two WINS servers are specified to improve fault tolerance in the event that the first WINS server does not have the record or is unreachable.
- The check box for **Do not replicate this record** is selected. The purpose of this configuration is to prevent the replication of WINS records to BIND secondaries that might encounter data errors or fail to load the zone if they encounter the proprietary WINS record in the replicated data.
- **Cache time-out** and **Lookup time-out** values are configured in the **Advanced** properties of the WINS tab. The **Cache time-out** value indicates the length of time the DNS server will cache WINS records. The **Lookup time-out** value indicates the length of time the DNS server will wait for a response from a WINS server.

The WINS forward record has the following format in the zone file:

```
@           WINS      LOCAL  L2 C900 (192.168.100.20 192.168.179.3)
```


The @ is a kind of shorthand used in DNS files to indicate the domain name, also known as the *origin* for the domain, in this case `tacteam.local`. The LOCAL label indicates that the record should not be sent to secondary servers as part of zone replication. The L2 label refers to the lookup timeout value of two seconds. The C900 label indicates the cache timeout value of 900 seconds, or 15 minutes. Both of these represent the default values. If you have a relatively static environment, it can be advantageous to configure a longer cache timeout value of perhaps an hour or more.

WINS Reverse Lookup Records

Reverse lookup zones are used to resolve IP addresses to host names, rather than host names to IP addresses, as is the case with forward lookup zones. WINS records are not indexed by IP address. Therefore, the WINS server cannot do a reverse lookup. Consequently, in a reverse lookup zone, a WINS-R RR will cause the DNS server to issue a remote adapter node status query using the `nbtstat` command to determine the NetBIOS name associated with an IP address.

Configuring a WINS-R record in a reverse lookup zone is similar to configuring a WINS record. Figure 6.18 shows the property pages of the **WINS-R** tab for a reverse lookup zone.

Figure 6.18 The WINS-R Tab for a DNS Reverse Lookup Zone Showing Advanced Configuration Options



As with WINS forward lookup records, you have the option of preventing the WINS-R record from replicating to secondary servers. This will prevent problems with BIND secondaries encountering this record in the zone data.

Note that the values in the WINS-R record are different. Instead of specifying the IP address of a WINS server, you specify the domain name that should be appended to the reverse lookup query response. Also, in the **Advanced** property page, you can check a box to **Submit DNS domain as NetBIOS scope**. This option should be used *only* if you are using NetBIOS scopes on a subnet. When this option is selected, DNS uses the host name as a NetBIOS computer name to query the remote adapter node status, but submits the domain name as a NetBIOS scope identifier.



NOTE

NetBIOS scopes are used in certain, rare circumstances when it is necessary to isolate legacy computers from communicating with other groups of computers on the same subnet.

A WINS-R RR has a similar format to a WINS forward record in the zone data file:


```
@                WINSR    LOCAL  L2  C900 (tacteam.local. )
```

The @ indicates the origin of the domain, in this case the 100.168.192.in-addr.arpa reverse lookup domain. The tacteam.local. value is the domain name that will be appended to the host name.

WINS Referral Zones

In a mixed DNS infrastructure where you are not replicating WINS RRs to secondaries, clients will get varying answers to queries if they query a secondary zone for a WINS record. To get around this problem and to provide a means of organizing and distinguishing between WINS and DNS records, you should configure a WINS referral zone. A WINS referral zone is a delegated child subdomain of the parent domain. The WINS child domain contains only the SOA for the child domain and the WINS RRs. For example, if the parent domain is `tacteam.local`, you would configure a child domain named something like `wins.tacteam.local`. If you have a large network with multiple WINS servers for different locations, you could use multiple child domains, such as `dallas.tacteam.local` and `edmonton.tacteam.local`. However, in order for this configuration to work in your environment, you need to populate the DNS suffix search list on your DNS clients so that they will append the domain name of the WINS referral zone to unqualified queries (queries that do not use the FQDN). Figure 6.19 shows a possible configuration of a DNS client to support WINS referral zones.

Figure 6.19 DNS Client Suffix Search List Configured to Support WINS Referral Zones



You should note that this configuration overrides the default configuration, which is to **Append primary and connection specific suffixes** and **Append parent suffixes of the primary DNS suffix**. The default configuration allows a client to send a query for an unqualified host name based on the suffix configured for it in the properties of **My Computer** and to *devolve* the domain name to the suffix of the parent domain. For example, if the client FQDN is `host1.dev.research.tacteam.local`, and it issues a recursive query to resolve the name `PServer1` to an IP address, it will first append `dev.tacteam.local`

to the name query. If the query fails, it will subsequently devolve the suffix to the parent domain and append `tacteam.local` to the name query.

Overriding the default settings for the DNS suffix search list increases administrative effort. However, you can reduce the administration of DNS client settings by using Group Policy settings to supply the clients with a DNS suffix search list. You cannot use DHCP options to specify a custom DNS suffix search list because Option 015 (DNS Domain name), which is used to specify the DNS domain name to append to unqualified queries, allows only one value. If you are implementing a custom DNS suffix search list, you should keep this list as small as possible to reduce DNS traffic on your network.

EXAM
70-293
OBJECTIVE
2.7.4

DNS Security Issues

Security should always be a primary consideration in the deployment of any network service. This is also true of the implementation of a DNS infrastructure. DNS is an open standard that is used throughout the Internet. Over the years, a number of exploits have appeared that can compromise an unsecured DNS infrastructure. When DNS is compromised, hackers can learn information about your internal network that they can subsequently use to launch other attacks. Furthermore, if a DNS server is vulnerable to DoS attacks, hackers can prevent name resolution from occurring for critical servers such as your Web and mail servers. Finally, an unsecured DNS server can be compromised with the addition of false records that redirect traffic to bogus Web and mail servers.

Security measures that you can take to mitigate risk to your DNS infrastructure include those available to standard DNS implementations, such as disabling recursion on Internet-facing servers, as well as those available to Windows Server 2003 DNS only, such as using Active Directory-integrated zones for zone transfers and secure dynamic updates.

As with developing any security policy, it is important to understand the nature and likelihood of the threats involved to determine the cost to the organization if a particular threat is realized, and then compare this cost with that of implementing countermeasures to mitigate the risk to the organization. Certain trade-offs need to be considered. For example, to completely secure your DNS infrastructure from attacks launched from the Internet, the only completely reliable countermeasure is to not have an Internet connection. Obviously, many organizations could not survive without Internet access, so this particular countermeasure is not appropriate.

In the next section, we will take a look at common threats to a DNS infrastructure. Then we will review the standard and Windows-specific countermeasures you can take to mitigate the risk from these threats.

Common DNS Threats

An unsecured DNS infrastructure is vulnerable to a number of common threats. These include footprinting, redirection, and DoS attacks. These threats are described in the following sections.


Footprinting

Footprinting is the process whereby attackers gain information about your internal DNS RRs and are subsequently able to use this information to infer the identity and purpose of servers on your internal network. Attackers can use this information in a variety of ways to compromise the organization. For example, an attacker can use this information to launch *data modification* attacks using *spoofed* IP addresses to compromise critical servers and data on the internal network. Another possibility is that, because host names are often informative, the attacker could use this information to infer confidential information about the internal operations of the company, such as products that are under development.

Footprinting often occurs when zone transfers are not secured and the attacker is able to perform a name dump from authoritative servers using the **nslookup** command with the **ls** option or the **dig** command with the **afxr** option—both of these commands initiate a zone transfer from the target domain.

To mitigate the risk from footprinting, it is important to ensure that zone transfers are secured. At the very least, zone transfers should be allowed to only a predetermined list of IP addresses that can be configured in the properties of the primary zone on the DNS server, as shown in Figure 6.20. You should also remember to secure your secondary name servers from unauthorized zone transfers, not just your primary server. Keep in mind that a secondary name server can also transfer zone information. However, even this configuration is vulnerable. For maximum security of zone transfers, you should ensure that zone transfers occur only within Active Directory-integrated zones. If you must transfer zone information over the Internet, you should also consider the use of VPN tunnels or IPSec to secure this traffic.

Figure 6.20 Configuring a Primary Zone with a List of Secondaries Authorized to Do Zone Transfers



Redirection

A *redirection attack* occurs when an attacker is able to modify DNS records to redirect Web server or other traffic to servers under the attacker's control. This attack occurs when an attacker is able to write information to the zone file. For example, this might happen if dynamic updates are enabled on a zone that is located on an Internet-facing DNS server. For this reason, it is always prudent to disable dynamic updates on zone files that are accessible to clients on the Internet.

Another common cause of redirection attacks is *cache pollution* (also called *cache poisoning*). Cache pollution can occur when a DNS server queries another DNS server and receives a reply from the queried DNS server that is outside the domain namespace in the original query. Unless countermeasures are taken, the DNS server will store this referral information in its cache, even though it did not originally request the information. For example, suppose that your DNS server issues a query for the MX record in the `sampledomain.com` domain. The authoritative DNS for the `sampledomain.com` server responds with the MX record, but it also replies with a bogus record for `a.root-servers.net`, listing its own IP address for the A record. Your DNS server now has a bogus record for a root-level DNS server in its cache.

DNS servers are vulnerable to cache pollution if an answer to a DNS query can be falsified. The consequences of cache pollution can be severe. Imagine what might happen if the poisoned cache of a DNS server redirected users to bogus Web site that contained malicious code designed to install Trojan viruses on client computers.

When cache pollution protection is enabled, the DNS server will discard from its cache the records it receives in response to queries if those responses contain information unrelated to the domain subtree of the requested resource. In our example, if protection against cache pollution is enabled, the DNS server will cache the MX record for the mail server in `sampledomain.com`, but will not cache the record for the `a.root-servers.net`, since it is not part of the queried domain subtree. Cache pollution protection is a DNS server-wide setting (**Secure cache against pollution**) and is enabled by default on Windows Server 2003 DNS servers (see Figure 6.15 earlier in this chapter).

Another way to mitigate the risk of cache pollution is to disable recursion on the DNS server. An attacker can use recursion to query the DNS server for resources in the attacker's domain. The recursive name server is then forced to query DNS servers in the attacker's domain that might attempt to pollute the cache of the recursive server.

DoS Attacks

A DoS attack occurs when a DNS server is deliberately flooded with traffic to the extent that it cannot respond to legitimate requests. DoS attacks on a DNS can be in-band on UDP and TCP port 53 (the ports used for DNS queries and zone transfers), or they can be out-of-band. In the case of an in-band attack, DNS servers are flooded with recursive queries to the extent that they become unable to handle legitimate queries, or the DNS service is subjected to a buffer overflow attack specific to the DNS service. In an out-of-

band DoS attack, the DNS server is the victim of an attack that is not specific to the DNS service, such as buffer overflow, SYN, and Smurf attacks. When a DoS attack occurs on a DNS server, mail servers and Web servers become unavailable as well because the host names for these servers cannot be resolved to IP addresses.

One approach to mitigate the risk of DoS attacks against your DNS server is to eliminate single points of failure by having multiple DNS servers that are located on separate subnets served by separate routers. Also, you can arrange to have secondary servers hosted offsite by a third party, such as your ISP.



NOTE

Recently, Microsoft's own DNS servers were the victims of a DoS attack that made a number of Microsoft Web sites inaccessible. The reason that Microsoft's DNS servers were vulnerable is that all of them were placed in the same physical location behind a single router, hence exposing a single point of failure.

To provide further protection against in-band DoS attacks, you can disable recursion on Internet-facing DNS servers. Recursive queries take a relatively long time to process, making a DNS server that performs recursion vulnerable to a DoS attack that involves sending a large number of recursive queries to the DNS server. When you disable recursion on a DNS server, it will not respond to recursive queries issued by DNS clients. DNS clients will not be able to use this server to resolve names on the Internet. However, the DNS server will still respond to iterative queries issued by other DNS servers. This means that it will respond to queries for resources in zones for which it authoritative.

Recursion is a server-wide DNS setting and is enabled by default. (You can also disable recursion for forwarding servers on a per-domain basis.) If you disable recursion for the entire DNS server, you will not be able to use that DNS server as a forwarder. You can see the **Disable recursion (also disables forwarders)** option in Figure 6.15, shown earlier in the chapter. On internal DNS servers, it is often not desirable to disable recursion. In this case, these DNS servers need to be protected by firewall access rules that prevent their use by DNS clients on the Internet.

To provide further protection against both in-band and out-of-band DoS attacks, it is important to ensure that you apply the latest service packs and harden the servers as much as possible. In addition, your firewall access rules and packet filtering should be configured to prevent any external traffic that is not related to the DNS service from reaching the DNS server. For example, a firewall that is in front of a DNS server in a DMZ should allow traffic to reach the DNS server only on TCP and UDP port 53.

Securing DNS Deployment

In the preceding section, we identified some of the common threats to the DNS infrastructure and provided a number of countermeasures such as securing zone transfers, disabling

recursion, and enabling protection against cache pollution. However, securing a DNS infrastructure requires more than just fine-tuning settings of the individual DNS servers.

Securing the DNS infrastructure starts with the design and implementation of your DNS namespace, and continues with the implementation and configuration of the DNS servers themselves, along with the implementation and configuration of firewalls, routers, and other network devices that can serve to protect individual servers and the network itself. It is possible, for example, to use a private root zone on your intranet and tightly control DNS query access to the Internet. Using a private root in combination with a DNS security policy that restricts DNS queries to the Internet can result in enhanced security for your organization.

DNS Security Levels

To assist in the secure deployment of a DNS infrastructure, Microsoft has published guidelines on its Web site and within the Windows Server 2003 help files that categorize three basic levels of DNS security: low level, medium level, and high level. In the following sections, we will discuss each level in more detail. In considering these models, you should assume that they represent a set of ideal guidelines for the purposes of conceptualization and example. Many organizations do not want to slavishly abide by the models in their purest form.

Low-level DNS Security

The low level of DNS security is precisely that: *low*. In fact, some of the default security configurations of DNS are removed entirely. The effective security is none at all. As the Windows Server 2003 help files state, this kind of configuration should be used only when there is no concern for the integrity of your DNS data or there is no threat that the DNS data on a private network is accessible from the Internet. The characteristics of low-level security are as follows:

- The DNS infrastructure is fully exposed to the Internet.
- All the DNS servers in your network use standard DNS resolution.
- All DNS servers are capable of performing queries to the Internet using root hints that point to the root servers for the Internet.
- Zone transfers are allowed to any server, which represents a removal of the default setting to allow zone transfers only to servers listed in the Name Servers tab.
- The default setting to prevent cache pollution is disabled on the DNS server.
- Multihomed DNS servers (servers with multiple IP addresses) are configured to listen for DNS queries on all configured interfaces.
- All zones are configured to accept dynamic updates from DNS clients.

- UDP and TCP port 53 are open on the firewall for both the source and destination address (that is, the firewall allows any DNS traffic to traverse your firewall, regardless of whether it is initiated by an external or an internal host).

Some organizations may have such a deployment; however, it would be extremely unwise to deploy something like this yourself. Turning off cache pollution protection, in particular, exposes your DNS infrastructure to an unacceptable level of risk, relative to the cost of leaving the default configuration enabled.

Medium-level DNS Security

The medium level of DNS security takes advantage of the countermeasures that are available in a DNS infrastructure where zone data is stored in standard primary or secondary zone files. The security features available through Active Directory-integrated zones are not employed here. The characteristics of medium-level security are as follows:

- Exposure of your DNS infrastructure to the Internet is minimized.
- Internal DNS servers are configured to use a limited list of forwarders when they cannot resolve names locally.
- The default configuration to limit zone transfers to DNS servers listed on the Name Servers tab is left in place.
- In the case of multihomed DNS servers, the DNS servers are configured to listen on only specified IP addresses.
- The default setting to prevent cache pollution is left in place.
- No dynamic updates are allowed on any zones.
- The firewall is configured to limit the traffic traversing the firewall to a limited set of source and destination addresses. Only the external DNS servers under your control are allowed to communicate with internal DNS servers.
- Only the external DNS servers in front of your firewall are configured with root hints to perform recursion.
- All name resolution required by a host on your internal network is performed by proxy servers or gateways.

This represents a more reasonable and prudent approach to mitigating risk to the DNS infrastructure than is offered by the low level, with a low cost of implementation relative to the advantages gained.

High-level DNS Security

A high-level security policy starts with the medium-level security policy and further enhances security by leveraging the security available with Active Directory-integrated zones. Furthermore, the high-level security policy assumes that there is no DNS communi-

cation with the Internet. This is an unlikely configuration but something like it might be implemented by organizations that have strict security requirements, and the risk of connectivity to the Internet is deemed to be too great. The characteristics of a high-level security policy are as follows:

- No DNS communication is allowed between the Internet and internal DNS servers.
- The internal DNS infrastructure deploys a private, internal root namespace and is authoritative for all zones.
- The root hints file on all DNS servers points to only the IP addresses of the internal DNS servers that are authoritative for the private root zone.
- Zone transfers are limited to specific IP addresses, rather than just servers listed on the Name Server tab.
- DNS servers are configured to listen on specific IP addresses.
- All DNS servers run on domain controllers, with discretionary access control lists (DACLS) configured to allow only specific authorized individuals to perform administrative tasks on the DNS servers.
- All DNS zones are configured as Active Directory-integrated zones, with DACLS configured to allow only specific authorized individuals to create, modify, or delete DNS zones.
- All RRs stored in Active Directory-integrated zones have DACLS to allow only specific individuals to create, delete, or modify zone data.
- No dynamic updates are allowed on the root and top-level domains.
- Only secure dynamic updates are allowed on the child domains.

For many organizations, none of these models will be adequate. The cost, for example, of not allowing DNS communication with the Internet, and expense of connectivity, might be too great. The reality is that many organizations will want to develop and deploy a DNS security model that is hybrid of the medium-level and high-level security models.

General DNS Security Guidelines

In planning for the security of your DNS infrastructure, you will want to take into account the design of your DNS namespace, the number and type of DNS servers and zones you plan to deploy, and whether the DNS servers will be serving internal or external clients. You will also want to take into account the security already present or needed in your current infrastructure, such as the location, type, and configuration of firewalls that protect your network.

In most cases, it is desirable to maintain a set of DNS servers that serve the internal network only and a separate set of external DNS servers that allow DNS clients on the

Internet to be able to resolve the names for your Web, mail, and other publicly available servers. Each set of DNS servers would have different security configurations, depending on their role. Furthermore, it is desirable to further enhance security of these two sets of DNS servers by maintaining either a split DNS configuration if you choose to use the same namespace for the intranet and Internet, or a split DNS namespace for the intranet and Internet. If your internal namespace includes a private root zone, you can further enhance the security of the DNS infrastructure.

Security Guidelines for an External DNS Infrastructure

Integrity and availability of DNS data are primary considerations for an external DNS infrastructure, and your design should be informed by these considerations:

- Place all DNS servers in a DMZ or a perimeter network to ensure that access rules and packet filtering on firewalls and routers tightly control source and destination addresses and ports. If possible, configure single-purpose DNS servers and allow traffic on only UDP and TCP port 53 to reach these servers from the Internet.
- Uninstall all unnecessary services from these servers, install current service packs, and harden the servers as much as possible.
- Eliminate single points of failure by hosting DNS servers on different subnets served by different routers. Consider hosting a secondary server at your ISP, for example. This will help mitigate the risk of DoS attacks.
- Consider using a stealth primary server to update read-only secondary servers that are registered with ICANN.
- Allow zone transfers to only a specific set of IP addresses and consider using IPSec or VPN tunnels to enhance the security of zone transfer traffic.
- Do not enable dynamic updates on Internet-facing DNS servers.
- Enable protection against cache pollution on Internet-facing DNS servers.
- Disable recursion on Internet-facing servers.
- Regularly monitor DNS logs and Event Viewer.

Security Guidelines for an Internal DNS Infrastructure

Confidentiality, integrity, and availability of DNS data are primary considerations for an internal DNS infrastructure. The following are security guidelines to consider:

- Consider using a separate, internal namespace to enhance security.
- Do not allow external access from the Internet to your internal DNS servers.

- Consider using a proxy server or a gateway to manage Internet DNS requests for internal clients.
- Use Active Directory-integrated zones and allow only secure updates to these zones.
- Specify and limit the servers that are able to receive zone transfers.
- Eliminate single points of failure and consider how internal DNS clients will resolve names in the event that the primary DNS server in their TCP/IP configuration fails.
- Consider that delegating authority of child domains can involve a security trade-off if different administrators are responsible for the authoritative DNS servers.

DNS Ports

For configuring firewall access rules, keep in mind that DNS uses both TCP port 53 and UDP port 53 for DNS communications. UDP is generally used for normal query traffic, whereas TCP is used for zone transfers. However, if the DNS server cannot deliver a response to the query using UDP because the response is too large, it will ask the resolver to switch to TCP port 53. This should occur only rarely (or never) if the DNS records are properly configured. The most common cause is an excessive and improper use of records used for round-robin name resolution or an excessive number of name server records.

You can use EDNS0 (discussed earlier in this chapter) to increase the default size for UDP packets. However, you would want to do this on your internal network, not the Internet. If you want to load-balance your Internet Web servers, you might want to consider using NLB or a third-party product such as BigIP from F5. In any event, it should be safe to block inbound traffic on TCP port 25 to prevent zone transfers from all but authorized DNS servers.

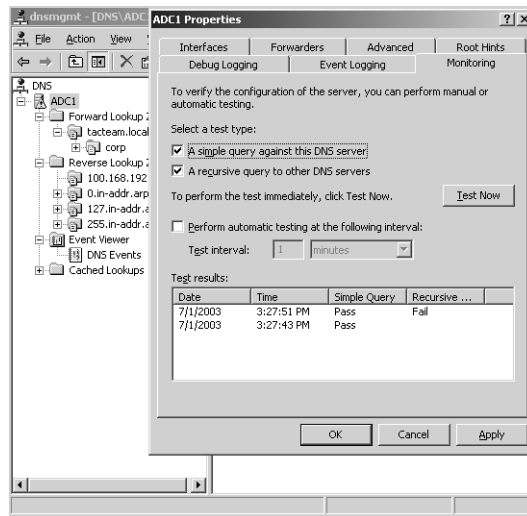
Monitoring DNS Servers

An important task in maintaining a DNS environment is monitoring the DNS servers to ensure that they are resolving names and IP addresses properly, and to ensure that they have sufficient resources to handle their workload. Windows Server 2003 and the Windows Server 2003 DNS service provide a number of tools for monitoring DNS servers. These tools include the Monitoring tab on the DNS console, DNS debug logging, DNS event logging, and DNS Performance Monitor counters, as well as command-line tools such as NSLookup.exe, Dnscmd.exe, and DNSLint.exe. In this section, we will briefly cover the use of these tools to monitor a DNS server environment.

Testing DNS Server Configuration with the DNS Console Monitoring Tab

The DNS console provides a simple but effective tool for ensuring that the DNS service is working properly. To use this tool, click the **Monitoring** tab of the properties for the DNS server, as shown in Figure 6.21.

Figure 6.21 Performing Simple and Recursive Queries Using the Monitoring Tab of the DNS Server Properties

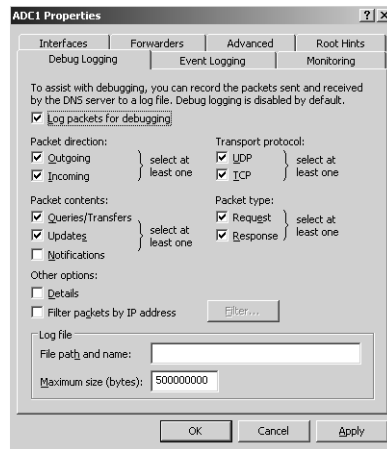


The Monitoring tab allows you to perform a simple and a recursive query test to ensure proper operation. A simple query test uses the DNS client installed on the DNS server to send a local query to the DNS server. A recursive query test uses the local DNS client as well. However, in this case, the DNS client requests that the DNS service use recursion to resolve an NS-type query for the root zone. Failure of this test usually indicates a problem with network connectivity or incorrectly configured root hints. (In the example in Figure 6.21, the recursive query test failed because the network adapter was unplugged before the test was run, and the DNS server could not connect to the servers listed in the root hints file.) When a DNS server fails one of these tests, a warning symbol is displayed on the DNS server in the DNS console. Note that you can set up automatic simple and recursive query testing in the Monitoring tab. It is a good practice to use these tests after you have set up a DNS server or have made a configuration change on a current DNS server.

Debug Logging

If you need to analyze and monitor the DNS server performance in greater detail, you can use the optional debug tool that you can enable in the **Debug Logging** tab of the DNS server property pages. Because debug logging consumes significant resources, it is not enabled by default and should be enabled only on a temporary basis, such as when you're trying to troubleshoot a problem with DNS. Figure 6.22 shows the configurable properties for DNS debug logging.

Figure 6.22 Debug Logging Properties



As you can see in Figure 6.22, you have a lot of flexibility and control with regard to the filtering of DNS traffic you wish to include in the debug logs. You can choose to log packets based on the following:

- Their direction, either outbound or inbound
- The transport protocol, either TCP or UDP
- Their contents: queries/transfers, updates, or notifications
- Their type, either requests or responses
- Their IP address

Finally, you can choose to include detailed information.

Let's assume you were trying to troubleshoot a problem with dynamic updates. You could configure the debug utility to log any update packets but exclude queries/transfers and notifications. This configuration would exclude information that isn't relevant to the problem. You could further refine the information contained in the logs by monitoring for either requests or responses or for incoming or outgoing packets.

Depending on the amount of DNS traffic the server processes and the logging options you select, the log files can grow quite large. You should configure logging to occur on a

separate hard drive. When the log file reaches the maximum size or the hard drive runs out of room, newer events will overwrite older events.


Event Logging

By default, the DNS service will log all DNS events to the *DNS Event log*. In Windows Server 2003, DNS events are kept in a separate system event log that can be accessed from either the DNS console or Event Viewer. The **Event Logging** tab on the properties of the DNS server allows you to configure the events you would like to log in the DNS Event log. There are four options on the **Event Logging** tab: **No events**, **Errors only**, **Errors and warnings**, and **All events**. The default is to log **All events**, which include informational messages that indicate service startup, a new version number for a zone file, and so on. On a busy DNS server, the default size of the event log might not be large enough. You should consider increasing the size of the DNS Event log in this case.

Monitoring DNS Server Using the Performance Console

The Windows Server 2003 Performance console provides a means of monitoring DNS performance, either in real time through the System Monitor or as events logged over a period time by Performance Logs. When the DNS service is installed on Windows Server 2003, more than 60 performance counters are installed for measuring performance of the DNS service. Figure 6.23 shows some of the DNS performance counters that you can select in System Monitor.

Figure 6.23 DNS Performance Counters



Because the DNS is a critical service, you should log its performance over a period of time using Performance Logs to establish a baseline for normal operating conditions. Once you've established a baseline, you can then use this information to predict effects of planned changes to the infrastructure, such as adding or removing other DNS servers or adding more DNS clients. Performance baselines also help you to optimize services on your net-

work by providing real-world data about the performance of your servers and your network.

Having a baseline also allows you to detect and troubleshoot problems with your DNS and network infrastructure. For example, if the number of **Secure Update Failures** suddenly increased, you might be prompted to investigate further to determine the cause of the problem.

In choosing DNS counters to monitor, you should consider the role(s) of the DNS server:

- If the DNS server is installed on a domain controller and configured for secure only updates in Active Directory-integrated zones, you should monitor counters that are relevant to dynamic updater performance and security, such as **Secure Update Failure**, **Dynamic Update Written to Database/sec**, **Dynamic Update Received/sec**, and so on.
- If the DNS server is used to perform recursion on behalf of clients, you should monitor counters that are relevant to the performance of recursive queries, such as **Recursive Queries/sec** or **Recursive Query Failures/sec**. If you have disabled recursion on your server, a spike in the number of recursive queries the DNS server receives could warrant further investigation.
- If the DNS server replicates zone data with other servers, either as a primary or secondary server, you should monitor counters relevant to zone transfers, such as **AFXR Requests Received**, which would indicate that a number of secondary DNS servers are requesting a full, rather than incremental, zone transfer. A sudden increase in the number of zone transfer requests could indicate the presence of an attacker trying to footprint your DNS records.



EXAM WARNING


Microsoft exams will not test you on how to use utilities. Rarely does the test candidate need to remember commands and switches. However, many questions that you encounter are often informed by an understanding of the capabilities of a particular utility. The Performance Monitor counters are no exception. Often, your ability to answer a question correctly will depend on your understanding of the capabilities of the utility.

Command-line Tools for Maintaining and Monitoring DNS Servers

Windows Server 2003 provides three command-line utilities for maintaining and monitoring DNS servers:

- **NSLookup** This is a standard tool used for monitoring and troubleshooting DNS servers. It provides a means to obtain detailed results for queries performed against a DNS server. NSLookup has two modes: interactive and noninteractive. Interactive mode allows you enter more than one command at an NSLookup prompt. Noninteractive mode is invoked as a single command with options from a command prompt. For NSLookup to work properly, the DNS server that NSLookup is pointing to must have a PTR record for it in a reverse lookup zone.
- **Dnscmd** This utility is found in the \Support\Tools folder on the Windows Server 2003 installation CD. The Dnscmd tool can be used as an alternative to the DNS MMC. With Dnscmd, you can create and delete zones, view records, update zone records, and perform other administrative tasks that you would normally perform using the DNS console. Dnscmd can be used to script batch operations and perform remote administration, providing an efficient way to manage multiple, remote DNS servers.
- **DNSLint** This utility is found in the \Support\Tools folder on the Windows Server 2003 installation CD. DNSLint is new to Windows Server 2003. Its primary purpose is to assist in troubleshooting problems arising from lame (incorrect) delegations and common AD DNS problems, such as verifying records for AD replication. A key advantage of the tool is that it can examine multiple servers in a single operation and display the output as an HTML file. For example, if you were trying to troubleshoot a problem with delegation, you would need to traverse the DNS namespace in multiple steps. With DNSLint, you can diagnose the problem in a single operation. You can also use DNSLint with the /c switch to test well-known e-mail ports on all e-mail servers that it finds in the zone records of the DNS servers it checks in the domain.

These tools can be used for a variety of purposes, such as verifying the presence of RRs, checking for lame delegations, checking for missing AD replication records, configuring DNS server settings on multiple servers, and so on.



Planning for NetBIOS Name Resolution

In a Windows 2000 or Windows Server 2003 environment, DNS is the primary method for name resolution. However, even in these environments, NetBIOS name resolution might still be extensively used. For example, if the network consists of older clients such as Windows NT 4 and Windows 9x clients, you must still support NetBIOS name resolution. Also, certain applications such as Microsoft Exchange Server, still rely on NetBIOS for their functionality. So, even if the domain is upgraded to AD and all of the clients on the network are upgraded to Windows 2000 or Windows XP, it might still be necessary to support NetBIOS name resolution.

Planning for NetBIOS name resolution means designing a NetBIOS name resolution infrastructure that ensures clients can always resolve NetBIOS names in a fault-tolerant and

timely manner. The primary means for ensuring fault-tolerant and timely name resolution is through the implementation of WINS. Through its ability to replicate information with other WINS servers, WINS provides a distributed database that allows NetBIOS clients to register their NetBIOS names to ensure uniqueness and to resolve NetBIOS-to-IP address mappings consistently throughout the network infrastructure. Because WINS servers are capable of replicating database information to one another, this means that multiple WINS servers can provide both fault tolerance and availability of records for NetBIOS resolution to even very large networks that involve many different sites.

Understanding NETBIOS Naming

A NetBIOS name is a 16-character string that is used to identify computers, groups, or services on the network. The first 15 characters of the name are configured on the computer by the user or the administrator. This name serves as a mnemonic to assist users in connecting to computers and resources on the network. The sixteenth and last character of a NetBIOS name identifies a resource type associated with the NetBIOS-related services that are running on the computer. For example, Windows XP and Windows 2000 clients will register a name for the messenger service, the workstation service, and the server service. Additionally, the resource type can indicate whether the computer is a domain master browser.

There are two kinds of NetBIOS names:

- **Unique names** These names are used to uniquely identify computers and the NetBIOS-related services running on them.
- **Group names** These names are used primarily to support browsing and browser elections.

Collectively, NetBIOS names comprise a *flat* namespace. This differs from a DNS namespace, which provides a hierarchical namespace. And, while it is possible to group NetBIOS names according to a workgroup or domain name for display by the browser service, NetBIOS names must be unique within the NetBIOS namespace.

NetBIOS Name Resolution Process

In a NetBIOS environment that does not employ LMHOSTS files or WINS servers, NetBIOS is completely dependent on broadcasts to register names and to resolve NetBIOS names to addresses. When one computer uses a NetBIOS application to communicate with another computer, it must first determine the address for that computer. For a client configured to use a WINS server, the typical steps for NetBIOS name resolution are as follows:

1. The computer checks its own name.
2. The computer checks its NetBIOS remote name cache to see if it has a mapping cached in memory.

3. The computer sends name resolution requests to WINS servers it is configured to use.
4. The computer checks its LMHOSTS files.
5. The computer sends a NetBIOS broadcast on the local subnet to find the address.
6. The computer can then use the hosts file, followed by DNS.

NetBIOS Node Types

The order of NetBIOS name resolution is actually governed by the NetBIOS *node type*. There are four possible NetBIOS node types:

- **b-node (broadcast node)** A b-node client will not contact a WINS server and will rely on broadcasts, an LMHOSTS file, and host name resolution.
- **p-node (peer node)** A p-node client will use a WINS server, but will not use broadcasts.
- **h-node (hybrid node)** An h-node client will contact a WINS server first before using an LMHOSTS file and broadcasts. The order of resolution for an h-node client is shown in the previous steps for NetBIOS name resolution, and it is the default configuration for a client configured to use a WINS server.
- **m-node (mixed node)** An m-node client will use broadcasts before it contacts a WINS server.

You can configure NetBIOS node types by specifying the use of a WINS server in the **WINS** tab of the **Advanced** TCP/IP property pages, through Registry settings, or through DHCP options. The node types are stored as a REG_DWORD at HKLM\System\CurrentControlSet\Services\NetBT\Parameters, as follows:

- b-node = 0x1
- p-node = 0x2
- h-node = 0x8
- m-node = 0x4

Without the presence of WINS or LMHOSTS files to assist in name resolution, NetBIOS name resolution would generate considerable broadcast traffic for name resolution, adding to the traffic generated to support NetBIOS registration and the browser service. Furthermore, since NetBIOS broadcast traffic normally does not cross routers, it would not be possible to resolve computer names on remote subnets.

To support NetBIOS name resolution on a segmented network, you can use two methods. The first relies on the deployment of LMHOSTS. The second relies on the

deployment of WINS servers. A third method, opening routers to forward NetBIOS broadcast traffic, is neither a sensible nor a viable solution in most instances, with the exception of small networks that must use separate networks because they are in different physical locations. That said, it is generally recommended that routers not be configured to forward NetBIOS broadcast traffic and that LMHOSTS files or WINS servers be used to support NetBIOS name resolution across subnets.

EXAM
70-293
OBJECTIVE
2.8.2

Understanding the LMHOSTS File

The LMHOSTS file is similar in purpose to the hosts file used for host name resolution. Both are text files that provide static mappings of names to IP addresses. However, the LMHOSTS file has added functionality that can be implemented with the use of special “tags” that can provide additional information. For example, you can use the #PRE tag (it is case-sensitive) to instruct the computer to always cache the record in the NetBIOS name cache. You can use the #DOM tag (it is also case-sensitive) to identify the computer as a domain controller.

LMHOSTS files are simple to create. They are created in the same folder where the hosts file is located. The complete path to the LMHOSTS file is `%systemroot%\system32\drivers\etc\lmhosts`, where `%systemroot%` is a variable used to identify the folder where the operating system is installed, such as `C:\Winnt` or `C:\Windows`. When you open the `%systemroot%\system32\drivers\etc` folder, you will notice that a file named `Lmhosts.SAM` has been provided for you as a sample. This file contains instructions and sample records for creating an LMHOSTS file. You can modify this file for use as your LMHOSTS file, but it is important to remember that you must save the file without the `.SAM` extension in order for the file to work for NetBIOS resolution. As with the hosts file, the LMHOSTS file should have no extension.

LMHOSTS files are parsed from top to bottom, so you can optimize the operation of LMHOSTS files by placing any entries with the #PRE tag at the bottom of the file. These entries need to be read only once to be cached.

It is also possible to use a centrally located LMHOSTS file to provide name resolution for clients through the use of the #INCLUDE, #BEGIN_ALTERNATE, and #END_ALTERNATE tags in the LMHOSTS file. The #INCLUDE tag is used to indicate the NetBIOS UNC share name and path to the LMHOSTS file. Multiple #INCLUDE statements can be listed between the #BEGIN_ALTERNATE and #END_ALTERNATE statements.

In order for the client computers to be able to access the share specified in the #INCLUDE statement, the computers hosting the remote LMHOSTS files must have support for NullSessionShares enabled, which allows anonymous connections to the share. This is a weakening of file sharing security, so you need to be careful when using LMHOSTS files in this way.

LMHOSTS files are a good solution in small environments that have a segmented network. In addition, they can be useful in situations where you want some computers to communicate with others across a WAN link, but you do not want to combine the

NetBIOS namespace of the offices on either side of the link. However, in large environments, LMHOSTS files are difficult to maintain. An LMHOSTS file must be present on each computer that needs it for name resolution. You can create centralized LMHOSTS files, but the client computers must first have an LMHOSTS file to gain access to the centralized LMHOSTS files. Also, you must manually enter NetBIOS name-to-IP address mappings, increasing the possibility for error. Finally, the use of LMHOSTS files is not possible in an environment that uses DHCP to assigned TCP/IP address configurations to client computers.

To support NetBIOS name resolution in a segmented network or one that uses multiple broadcast domains, a better approach than LMHOSTS files is to use WINS. If a network has been using LMHOSTS files extensively, it is relatively easy to migrate to WINS by importing LMHOSTS files to the database to create static mappings. However, you need to exercise care to ensure that these mappings can be overwritten by WINS clients that use dynamic mappings. We discuss this issue in more detail later in the chapter.



EXAM WARNING

Planning NetBIOS name resolution by using LMHOSTS files is part of the objective domain for Exam 70-293. Therefore, you might encounter questions that require a knowledge of the LMHOSTS file.

Understanding WINS

In segmented network environments that use DHCP, the best solution to allow for proper NetBIOS name resolution is to use WINS servers. A WINS server provides a database that NetBIOS clients can use to register their NetBIOS names and resolve NetBIOS-to-IP address mappings. Furthermore, WINS traffic is unicast-based. This means that instead of relying on broadcasts to register and resolve names, WINS clients will send unicast messages directly to the WINS server, whether on the same or different subnet. The use of unicast helps to reduce the amount of broadcast traffic on the network that is the result of NetBIOS name resolution.

The ability for different WINS servers to replicate database information with each other is another significant advantage in that the replication of this information ensures that clients can resolve NetBIOS names regardless of their location or the WINS server they contact.

WINS Name Registration, Renewal, and Release Process

When a WINS-configured client starts up, it will attempt to register its NetBIOS names (that is, its computer name and other resource types indicated by the sixteenth byte of the NetBIOS name) using a unicast message to the primary WINS server that is configured in the **Advanced** properties of the TCP/IP protocol object. (If the client cannot contact the

primary WINS server, it will contact the alternate WINS servers it is configured with. Up to 12 WINS servers can be specified on Windows 2000 and Windows XP clients.) The registration will be accepted or denied based on the presence of a preexisting registration for the same name.

If the name does not exist in the WINS database, the registration is accepted with a new version ID, given a timestamp, and marked with the owner ID of the WINS server. The timestamp is determined by adding the renewal interval (by default, six days) to the WINS server's current date and time. The WINS server then sends a positive response to the client with a TTL value that is based on the timestamp.

If the name already exists in the WINS server, the WINS server will perform a number of steps based on whether the record is active and whether the IP address is the same or different as the record in the database:

- If the name registration is for the same IP address and the WINS entry is marked as active and owned by the WINS server, a new timestamp is assigned and a positive response is sent back to the WINS client.
- If the name registration is for a different IP address and the record is marked as being inactive, or *tombstoned*, or is owned by a different WINS server, the registration is treated as a new registration.
- If the name registration is for a different IP address but the record is still marked as active, the WINS server will send a Wait for Acknowledgment (WACK) packet with a short TTL to the client trying to register the name. The WINS server will then issue a series of challenges in 500-millisecond intervals to the active node listed in its database (in the case of a multihomed computer, the challenges will be issued to every IP address listed for the node). If the node does not respond after three challenges, the WINS server treats the registration attempt as a new registration and sends a positive response to the client. However, if the client node in the WINS database responds to the challenge from the WINS server, the WINS server will send a negative response to the client attempting to register the duplicate NetBIOS name.

When a NetBIOS registration is successful, the WINS server stores the name mapping with the following information:

- **Record Name** The NetBIOS name of the computer, group, or service registered in the database.
- **Type** The resource type associated with the name. Common resource types are [00h] for the workstation service, [03h] for the messenger service, [020h] for the server service, and [1Ch] for the domain name.
- **IP Address** The IP address for the registered name.
- **State** The state of the registration, such as Active, Released, or Tombstoned.

- **Static** Indicates if record is a static mapping (column entry marked with an *x*).
- **Owner** Indicates the owner (specific WINS server where the record was registered) of the record.
- **Version** Indicates the version ID of the record.
- **Expiration** Indicates the date and time the record will expire, if it is not refreshed.

NetBIOS clients that do not use a WINS server must constantly register and defend their NetBIOS names. However, the registration occurs by means of a broadcast and, thus the registrations are local to the subnet. Whenever a computer attempts to register a duplicate name, it will receive a negative response from the computer that actively possesses the name. Whenever NetBIOS name registration fails, the computer receives an error message and NetBT won't initialize on the computer until the problem is resolved.

Whenever a WINS client is gracefully shut down, it sends a name release request to the WINS server. The WINS server marks the entry as *released* and gives it a timestamp of the current time plus the *extinction interval*. This interval indicates the time that must elapse before the record is marked as extinct (or tombstoned) and can be scavenged from the database. However, if the WINS server that receives the release request is not the original owner of the name registration, it will immediately mark the record as *extinct*. The reason for this is to prevent inconsistencies between primary and secondary servers configured as replication partners.

Tombstoned Records

A *tombstoned* record is one that is marked as extinct and can be scavenged (removed) from all WINS databases in which the record appears; that is, both the local and all replicated instances of the object will be removed. Tombstoning, which is a feature introduced with Windows NT 4.0 Service Pack 4 and carried over to Windows 2000, is a method by which an inactive record is allowed to remain in the WINS database so that its inactive state can be replicated to WINS replication partners. In the past, manual deletion of dynamic records was difficult and required extra administrative effort to prevent the record from reappearing on the original WINS server as a result of replication.

For manual deletion of dynamic records, WINS gives you the option of deleting only the local copy of the record or tombstoning the record; that is, replicating the deletion of the object to other WINS servers. To ensure that the record is properly removed, it is recommended that the manual tombstoning of the record take place on the WINS server that owns the record.

What's New for WINS in Windows Server 2003

The WINS service was significantly improved in Windows 2000, which introduced a number of significant enhancements, such as enhanced filtering and searching for records display and burst handling of WINS registrations. The WINS service in Windows Server 2003 maintains these improvements and adds two more improvements:

- **Filtering records** It is now possible to combine filters used to query the WINS database. The available filters include record owner, record type, NetBIOS name, and IP address with or without subnet mask. You can also cache the results of queries in the RAM of the local computer from which the query is performed, improving performance for subsequent queries and reducing network traffic.
- **Accepting replication partners** For pull replication, you can configure the WINS server to either accept or reject only the name records that are owned by a list of predetermined WINS servers. (Figure 6.24 later in this chapter shows the Advanced replication property page where accepting replication partners are configured.) In Windows 2000, it was possible to create only a list of replication partners to block.

Planning WINS Server Deployment

Planning for a WINS server deployment means designing a WINS server infrastructure that provides a fault-tolerant and highly available solution for clients to register and resolve NetBIOS computer names. Clients can be configured to use multiple WINS servers (up to 12) to register and resolve computer names in the event that the primary WINS server is unavailable.

WINS servers can be configured to replicate with one another to ensure both fault tolerance and the availability of records for name resolution in a distributed environment where WINS clients are registering with different WINS servers. The number of WINS servers that should be deployed depends on several complex factors such as the number of WINS clients that will be registering and resolving names and the network topology (in particular, the presence, number, and capacity of WAN links). If you require WINS replication to meet the goals of fault tolerance and availability, the WINS replication topology needs to be designed to ensure optimal performance in the replication of WINS records and the currency of those records, as well as to ensure optimal query response times, given the constraints of the network topology.

Server Number and Placement

The number of WINS servers you deploy will be determined by the number of WINS clients and the network topology. In general, you should try to design the WINS infrastructure to minimize the number of WINS servers. Having too many WINS servers can complicate network problems, so Microsoft recommends a conservative approach in determining the

number of WINS servers and Microsoft further recommends one primary and one backup WINS server set up as replication partners to each other for every 10,000 clients.

In capacity testing for WINS, Microsoft configured a dedicated WINS service on Windows Server 2003 machine running on a single-processor Pentium II with 128MB of RAM and a standard IDE hard drive. The server was able to process approximately 300 registrations per second and 350 queries per second. WINS is a disk- and CPU-intensive service, and these estimates will vary depending on the hardware that is in place. A dual-processor CPU will increase WINS performance by up to 25 percent. Placing the WINS database on its own separate disk spindle, using fast disks, and using RAID arrays can improve the performance of disk I/O required for WINS.

The WINS traffic between clients and servers is relatively small, about 40 bytes, which is the average size of a WINS registration for a client. So, 10,000 WINS records would be approximately 400KB. Through a 56 Kbps WAN link, this amount of data would take a minute to transfer, assuming you could transfer at this rate. It is more likely, though, that the effective bandwidth is somewhat lower than 56 Kbps, and the transfer of 400KB of information would take longer. This delay might or might not be acceptable. However, it is unlikely that 10,000 WINS registrations or 10,000 updated WINS records would need to be transferred simultaneously across the WAN.

To determine the actual number of WINS servers to deploy, you will need to take into account server hardware, WAN links, the number of clients, and the need for redundancy and availability. You should also take into account peak-load conditions that could occur, for example, if the power to client computers were suddenly terminated and then resumed, resulting in a large number of simultaneous registration requests.

EXERCISE 6.02

INSTALLING THE WINS SERVICE

In order to follow along with the discussion of WINS, you might find it helpful to install the WINS service on a Windows Server 2003 server. If you haven't already installed Windows Server 2003, see Exercise 6.1 for some advice on how to get an evaluation copy of Windows Server 2003 and how to run it in a virtual machine.

Before installing the WINS service, you need to ensure that the server meets a couple of prerequisites:

- The server needs to be configured with a static IP address.
- The WINS TCP/IP properties must be configured so that the server uses itself as a primary WINS server.
- You should not configure a WINS server to register with any other WINS server.

Here are the steps for installing WINS service:

1. Select **Start | Control Panel | Add or Remove Programs**.
2. Select **Windows Components**.
3. In the Windows Component Wizard window, scroll down the list of Windows components, Highlight **Networking Services**, and then click **Details**.
4. In the Network Service Services dialog box, click **Windows Internet Name Service (WINS)** and then click **OK**.
5. If prompted, insert the Windows Server 2003 source CD to provide the installation files for the DNS service.

After you have installed the WINS server, make sure that the WINS server has a registration for itself. You can do this by following these steps:

1. Open the **WINS** console by selecting **Start | Administrative Tools | WINS**.
2. Expand the WINS server node, right-click **Active Registrations**, and select **Display Records** from the context menu.
3. In the **Display Records** page, click **Find Now**. (Note that by not entering any query information, you are causing all records to be displayed. On a large network, you would want to limit the scope of the result set by entering query filters.)

If you don't see a record, you can force the registration of the server in the WINS database. Assuming that the WINS TCP/IP properties have been configured so that the server points to itself as the primary WINS server, open a command prompt and enter the command **nbststat -RR** (you must use uppercase *Rs* to achieve the desired results). Executing this command will cause the Windows Server 2003 to register itself in the WINS database without needing to restart.



NOTE

It is not necessary to use the WINS console to configure WINS servers and to view records. Anything that can be accomplished using the WINS console can be accomplished from a command line using the Netsh utility. You can find more information about the Netsh utility in the Windows Server 2003 help files.

EXAM
70-293
OBJECTIVE
2.8.1

Planning for WINS Replication

Even in simple environments, it is a good idea to have two WINS servers configured as replication partners to provide greater fault tolerance and availability. This kind of replication is relatively easy to set up, for example by specifying configuration, and requires little planning.

In larger and more complex networks, however, the replication topology will need to be carefully designed to ensure an optimal *convergence time* for the replication of WINS records, given the size, topology, and available bandwidth of the network. Convergence time refers to the maximum amount of time it takes for an updated record to be replicated to all WINS servers in the infrastructure. Generally, convergence time is a function of replication pull schedules and the number of hops in a given replication path that a changed object must travel.

In reality, the amount of time it will take for a record to be synchronized depends on a number of factors. For example, a WINS server configured to send push notifications immediately upon receiving an update will replicate the record faster than the time determined by the replication schedule of the WINS servers that have configured it as a pull partner. Factors that can affect convergence time include the following:

- The kind of replication partnership that is configured. A push/pull replication partnership is more efficient for replicating records than a limited partnership (push or pull only).
- The settings for pull and push partnerships that determine the frequency of replication between servers.
- The ability to use persistent connections for push and pull replication partners (this setting is the default and found in the **Replication** properties pages of the WINS server console.) The ability to use this setting depends on the presence of stable, high-speed links.
- The particular replication model. The longer the replication path that a replicated object must travel, the longer the convergence time. In a complex replication environment involving multiple WINS servers, a hub-and-spoke replication topology provides the shortest replication paths.
- The particular kind of update that occurs in the WINS database. For example, a name release update will not propagate as quickly as a name registration update because it is more common for a name-to-IP address mapping to be reused by the same computer, even in an environment that uses DHCP. Since this kind of update is not as urgent as a new name registration, the WINS server provides it with a greater *latency* period for replication.

Along with these factors, you will need to take into account the network topology in planning for WINS replication. Your WINS replication design should, if possible, mirror your network topology. For example, if the organization has a centrally located head office

that is connected by high-speed, persistent WAN links to satellite branch offices, you should consider a hub-and-spoke replication model with full push/pull partnerships. In this model, the WINS server in the head office receives replicated records through the push/pull partnership and then propagates updates it receives from its own WINS clients and those changes it receives from the individual branch offices to the other branch offices.

To plan for WINS replication, it is important to first understand partnership agreements and the settings that can be configured on them.

WINS Change-Only Replication

After the initial replication that occurs between two WINS servers that have been set up as replication partners, all subsequent replication between WINS servers replicates only the updates to the database since the last replication cycle, reducing the load on the network and the WINS servers.

An update is defined by an increment in the highest version ID in the local WINS database for records the WINS server owns. An incremental update will occur, for example, when a new registration is added to the database or a previously registered client updates its IP address (a refresh of a previously registered name with the same IP address does not cause a change in the version ID). WINS replication partners use these version IDs to determine how many WINS records need to be exchanged in order to synchronize databases using incremental updates.

Unlike the version number that is found in a DNS SOA RR and applies to the entire zone file (rather than individual RRs), a version ID is associated with every WINS record. Version IDs are incremented across the entire set of records that a WINS server owns. For example, imagine a WINS server that owns client records with the following version IDs: Client-A, version ID = 64; Client B, version ID = 65; and Client C, version ID = 66. The highest version ID on the WINS server is 66. (Version IDs are stored as hexadecimal numbers, so these numbers represent 100, 101, 103, and 104 in decimal, respectively.) If Client B registers a new IP address with the WINS server, it is assigned the next highest incremental value in the database, in this case 67. WINS servers configured as replication partners would compare the version ID they had for this WINS server and request only those records that would bring them up to date.

WINS servers also maintain an owner-mapping table to store the highest version IDs associated with various owners of records of which it is aware. This table is built dynamically and stored in memory at startup. WINS servers use this table when communicating with replication partners to determine how many records to transfer.

Replication Partnership Configuration

In order for WINS servers to replicate WINS records with each other, a replication partnership must be configured between them. There are three possible kinds of replication part-

nerships that can be configured between WINS servers: *push/pull* (also known as *full*), *push-only*, and *pull-only* (also known as *limited*). You can set up a replication partnership manually or implement it automatically.

Automatic Partner Configuration

Automatic partner configuration is an option that can be implemented on small networks to eliminate the administrative effort for configuring replication partnerships between WINS servers. When the automatic partner configuration is enabled, the WINS server will send announcements using the multicast Internet Group Messaging Protocol (IGMP) address at 224.0.1.24, which is the well-known multicast address for WINS servers. When the WINS server discovers other WINS servers that are announcing themselves, the WINS server will automatically configure a partnership agreement between itself and the discovered WINS server. (Both must be enabled for automatic partner configuration.) When the WINS server discovers another WINS server, it will add the server to its list of replication partners, configure push/pull replication between the servers, and set the pull replication interval for every two hours.

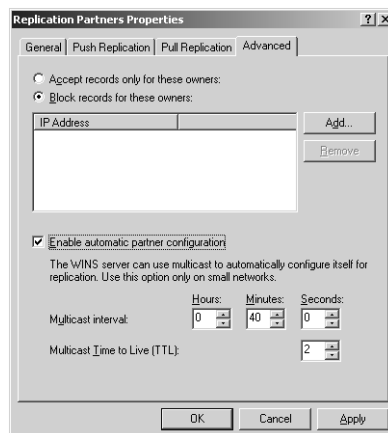


NOTE

Automatic replication partners are removed from the list of replication partners when they are shut down gracefully. If you want the replication partnership to persist across restarts, you should manually configure the replication partnership.

Figure 6.24 shows the **Advanced** tab of the **Replication Partners Properties** dialog box, which contains the **Enable automatic partner configuration** option. To view this page, you need to view the properties of the **Replication Partners** node in the **WINS** console.

Figure 6.24 Enabling Automatic Partner Configuration



Normally routers do not forward IGMP traffic, so this configuration is best used on small, unsegmented LANs. However, it is possible to configure routers to forward this traffic, allowing automatic partner configuration to be used in a routed environment. If there are only a few routers in the environment, the amount of multicast broadcast traffic should be minimal.



NOTE

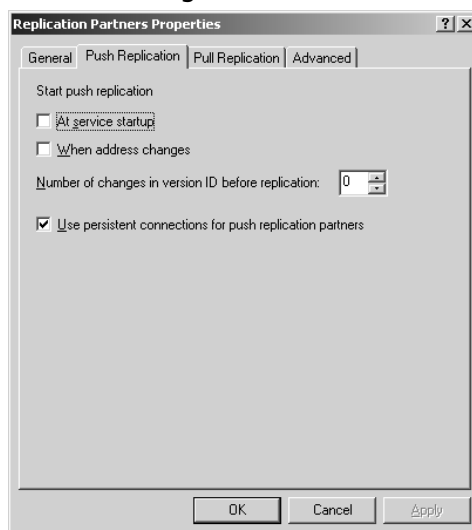
By default, automatic partnership configuration is limited to two hops in Windows 2000 and Windows Server 2003, which is reduced from the default six hops for WINS autodiscovery in Windows NT 4. You can extend this limit by changing the `McastTtl` Registry entry. You can find this entry at `HKLM:System\CurrentControlSet\Services\Wins\Parameters`. Note that this entry is not visible unless you enable automatic partner configuration.

Push Partnerships

As the name implies, when a push partnership is configured, changes in the WINS database are *pushed* to the remote WINS server. More accurately, a WINS server with records to replicate sends a push notification to target servers (those configured to use it as a pull partner), alerting them that it has records to update on the target WINS servers. The push notification includes an owner table that lists the owner IDs and the highest version ID for each owner. The target servers compare this information with their own owner tables to determine which records to replicate. The target servers reply to the push notification with a pull request, and the transfer of records takes place. Accordingly, since a transfer of records will not take place until a pull request has been received by the server that sent the push notification, pull replication is the single mechanism for replication. The process for push replication occurs as follows:

1. The source WINS server receives updates to its database and, based on a configurable threshold, sends a push notification to the destination WINS server (its push partner), indicating that it has updates to replicate.
2. The destination WINS server for the notification (the push partner) responds by initiating a pull request to its pull partner (the WINS server that sent the notification), and the replication is initiated between the replication partners.

Push replication is not schedulable according to an interval of time. Rather, the WINS administrator configures an update threshold that will trigger a push notification. For example, the WINS server could be configured to send a notification to its push partner after it has received 100 updates. Figure 6.25 shows the **Push Replication** tab of the **Replication Partners Properties** dialog box with the default settings for push replication.

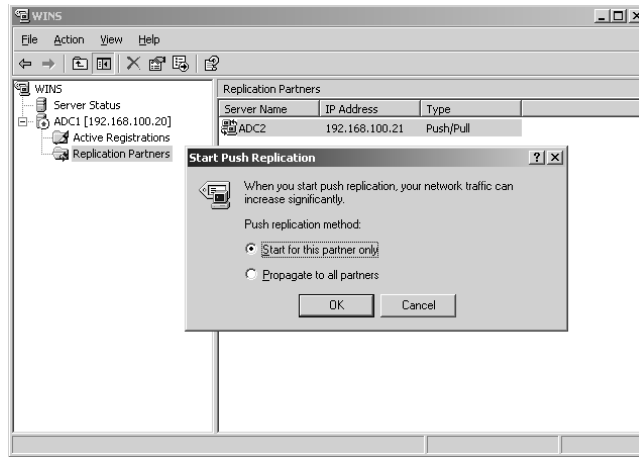
Figure 6.25 Push Replication Settings

The settings that you enter here will determine the threshold trigger for the push notification. In the configuration shown in Figure 6.25, a push notification is sent to the replication partner as soon as an update occurs in the WINS database. This is the result of setting the value for **Number of changes in version ID before replication** to **0**. However, the value could be set to a higher number, such as 100. It is also possible to configure a push notification to occur when the service starts up or when there is an address change in a WINS registration.

The setting to **Use persistent connections for push replication partners** allows the connections between WINS servers to remain open. This is a useful feature when the WINS servers are connected by a high-speed LAN. Earlier versions of WINS would close the connection after each replication cycle. Opening the connection to initiate a new replication cycle could cause delays, however modest, that are not acceptable in an environment where records need to be synchronized as soon as possible.

It is also possible to manually initiate the push notification, as shown in Figure 6.26. When you manually initiate the push notification, you can choose to push the notification to the replication partner or to trigger the replication to send a notification to all of its partners as well. As an example, consider a replication topology where three WINS servers are configured as push replication partners. WINS-A replicates to WINS-B, which replicates to WINS-C. So, if you manually sent a push notification from WINS-A to its replication partner, WINS-B, you could force WINS-B to also send a push notification to its other replication partner, WINS-C.

Figure 6.26 Manually Starting Push Notification



In certain rare situations, it might be desirable to use a *push-only* replication partnership for one-way replication, for example, from a head office to a branch office. For example, suppose that WINS-A in the head office configures WINS-B in the branch office as its push-only partner. (WINS-B should also configure WINS-A as its pull-only partner.) When WINS-A receives updates to its records, it notifies WINS-B, which sends an update (pull) request to WINS-A for the changed records since the last replication cycle. In this scenario, WINS-B never sends its updated records to WINS-A.

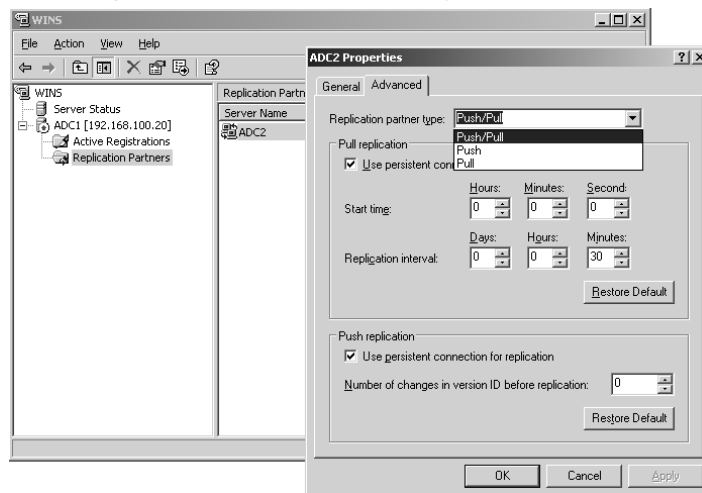
Push partnerships are generally configured in LAN environments where bandwidth is not an issue and it is not necessary to schedule replication to occur during off-peak hours. In general, you should use push replication partnerships in the following situations:

- There is ample bandwidth over LAN or WAN connections.
- There is a need to ensure that updates are replicated as soon as possible and the frequency of replication traffic is not a consideration.

Pull Partnerships

Pull replication differs from push replication in that the replication frequency is defined as an interval of time. At regularly scheduled intervals, a pull partner requests updates from other WINS servers (those configured to use it as a push partner) for updated records that have a higher version ID than the ones it currently has in its database.

Pull replication is configured similarly to push replication. The primary difference is that the WINS administrator schedules the times that the pull replication will take place. Figure 6.27 shows the settings for pull replication that an administrator can configure for individual replication partners.

Figure 6.27 Choosing Replication Partnership Type and Push/Pull Settings

In some situations, it might be desirable to configure *pull-only* replication between replication partners. Usually, this configuration is implemented where WAN links are operating close to capacity and there is a need to schedule WINS replication during off-peak hours. Pull-only replication has an advantage over push-only replication in that the replication schedule can be known in advance. With push-only replication, replication is triggered by reaching a configured threshold of updates, and you can only estimate when this would occur based on experience with the network. However, a disadvantage of pull-only replication is that the WINS server could potentially have acquired a large number of updates to replicate between cycles.

In general, you should use pull replication partnerships in the following situations:

- There is limited bandwidth between WINS servers that requires replication to be scheduled during off hours.
- There is a need to consolidate updates and reduce frequency and amount of replication traffic.
- There is a need to exercise finer control over the timing and frequency of replication traffic.

Push/Pull Partnerships

A push/pull partnership is the default when you configure replication between WINS servers. In fact, Microsoft recommends a push/pull partnership as a best practice and it further recommends that all WINS partnerships be set up this way unless there is an overriding need to implement a limited partnership. The only need that Microsoft cites for a

limited partnership is the presence of a large network connected by relatively slow WAN links. Microsoft often stresses the need for simplicity in a WINS environment.

With a push/pull partnership, a WINS server will be configured both to send push notifications and to make pull requests to its replication partner. The replication partner will also be configured in a similar way. Such a configuration helps to ensure that synchronization among WINS server is optimal, depending on the pull schedule and the configured threshold for push notifications, among other factors. For example, suppose that a WINS server that suddenly experiences a large number of updates immediately sends a push notification to its push partner. The push partner would immediately request these updates, without waiting for the request to be triggered by its pull schedule. Conversely, a WINS server always pulls up-to-date records from its pull partner according to the replication schedule, regardless of how few records have been updated on the pull partner WINS server. You should always try to deploy a push/pull partnership, unless there is an overriding concern that requires the implementation of a limited partnership.

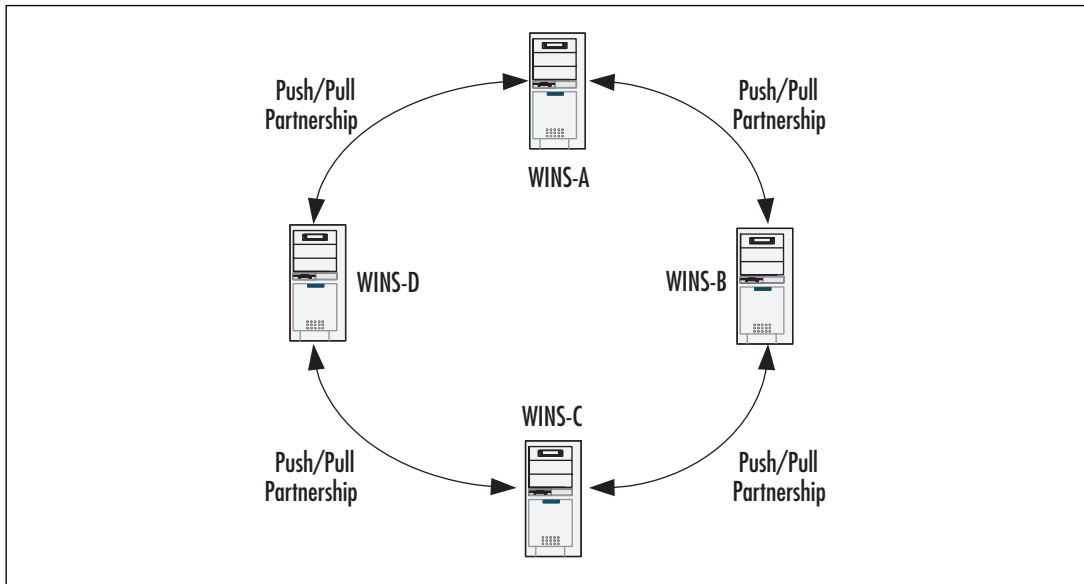
Replication Models

As we mentioned earlier, the replication model you design will have an effect on the convergence time for replicated WINS records and fault tolerance for replicated records. A replication model that is appropriate for your network topology will ensure the shortest convergence time for replicated WINS records. Where possible, it is recommended that your replication model mirror your network topology and that you keep this model as simple as possible.

In WINS environments where there are three or more WINS servers, you can employ either a *ring* replication model or a *hub-and-spoke* replication model. In more complex environments, these models can be combined to ensure optimal convergence time and fault tolerance for a given network topology. In the following sections, we will discuss each of these models in more detail.

Ring Model

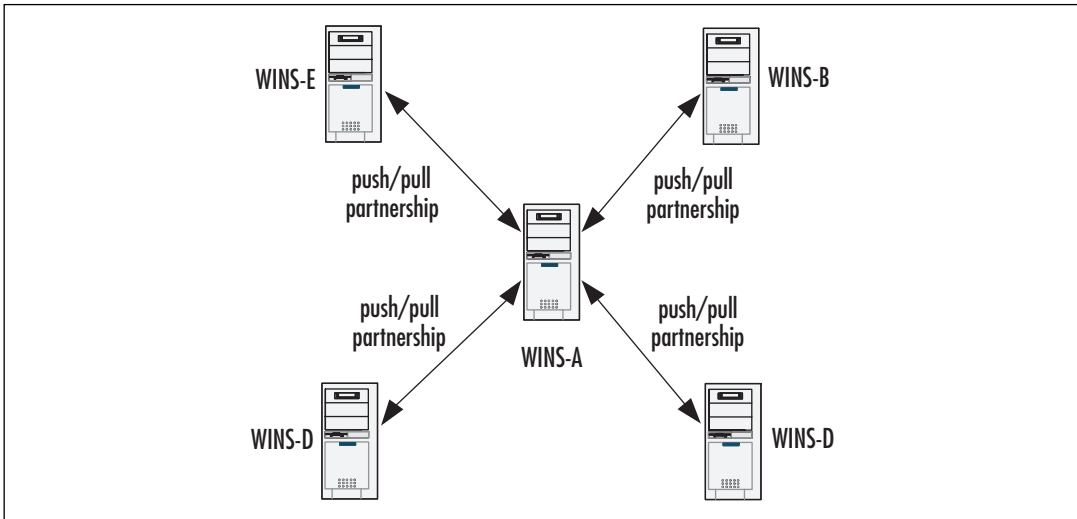
In a ring model, three or more WINS servers are configured to replicate with one another in a circular fashion. The ring model provides for good convergence times for all replication partners when there are no more than four WINS servers. Figure 6.28 shows a ring replication model.

Figure 6.28 Ring Replication Model for WINS Servers


In this model, fault tolerance for replication of WINS records is given priority. Imagine that a record is updated on WINS-A. The record must travel through either WINS-D or WINS-B before it is replicated to WINS-C. However, suppose that the WAN link connecting WINS-A and WINS-D fails. The updated record can still arrive at WINS-C and WINS-D (via WINS-C). Conversely, a record created on WINS-D can still be replicated to WINS-A via WINS-C and WINS-B.

Hub-and-Spoke Model

In the hub-and-spoke model, all WINS servers replicate with a centrally located hub WINS server. The hub-and-spoke model provides for the shortest convergence time in a replication environment that comprises five or more WINS servers, because it provides for the shortest replication paths between any two WINS servers. Furthermore, by implementing a hub-and-spoke model, you reduce the number of replication partnership agreements that you need to maintain. Figure 6.29 shows a typical hub-and-spoke implementation.

Figure 6.29 Hub-and-Spoke Replication Model for WINS Servers

Even though there are five WINS servers that replicate information, there are only four replication agreements to maintain. Furthermore, no server is more than two hops from any other server, regardless of the number of servers added to the topology.

A disadvantage of this model is that it is not as fault tolerant as the ring model. If WINS-A fails, no WINS server will be able to replicate its records to other WINS servers. Furthermore, depending on the average number of records the spoke WINS servers need to replicate and the settings for the push and pull triggers, WINS-A can be continuously replicating with other servers and processing updates. It should be well connected to the other WINS servers and have the capacity to handle the load.

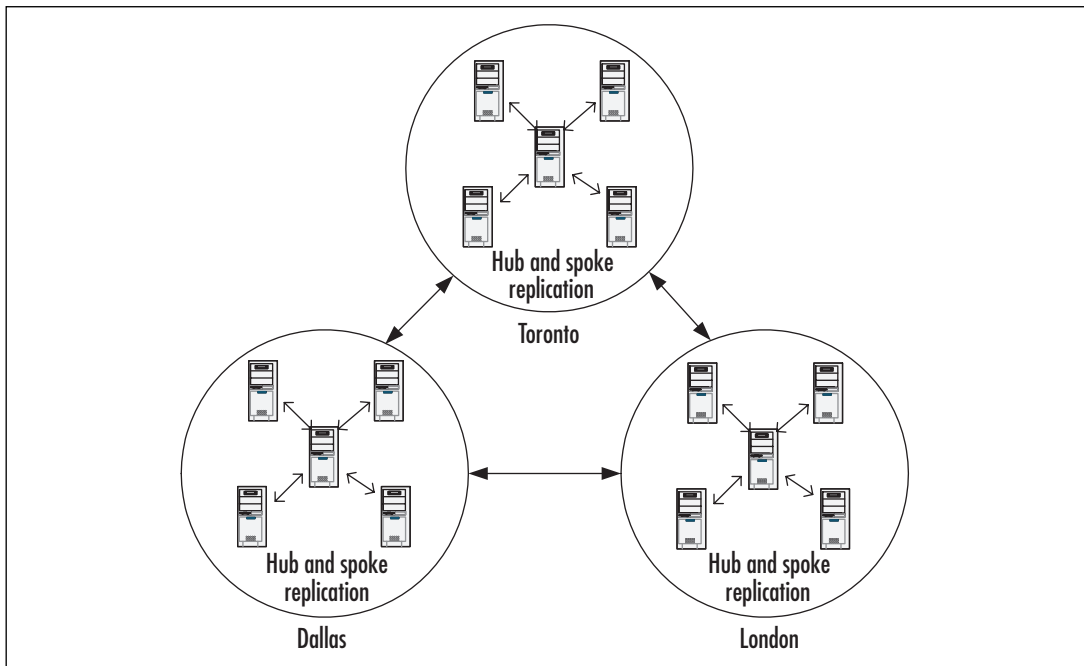
To enhance fault tolerance in this situation, you could set up a backup WINS server in the same location as WINS-A and configure a replication partnership agreement between them. This solution, however, increases administrative complexity for the maintenance of replication partnerships. An alternative solution that still provides a high degree of availability is to use Windows clustering for the hub WINS server.

A Windows cluster gives you the ability to set up separate WINS servers, known as *cluster nodes*, that use the same database located in a shared SCSI or Fibre Channel device. When the WINS server that is the active node in the cluster fails, the services will *failover* to another node. Failover is the process of taking resources offline in one node and bringing them online on a new node. The primary advantage of using a Windows cluster is that in the event of a failure of a WINS server, no subsequent replication needs to occur to synchronize records when the failed server is brought online, since only a single database is used. Windows Server 2003 Standard, Enterprise, and Datacenter editions support clustering. For more information about clustering, see the Windows Server 2003 help files.

Hybrid Replication Model

In many situations, it is desirable to combine replication models. As an example, consider a large organization that has three divisions in different geographic locations. Each of these divisions has a number of branch offices that are connected to their respective divisional offices. It might be advantageous to use a ring model of WINS replication among the divisional offices and use hub-and-spoke replication for replication between the divisional offices and their respective branch offices. Figure 6.30 shows a conceptual representation of a hybrid model. Many other variations are possible. A hybrid replication model can employ any mixture of full and limited replication partnerships, driven by the contingencies of the network topology.

Figure 6.30 Hybrid Replication Model



WINS Issues

After establishing the need for WINS planning for the replication topology of the WINS infrastructure, the WINS servers need to be installed and configured. In this section, we will look at the various configuration issues that a WINS administrator needs to be familiar with to ensure an efficient and secure WINS infrastructure, such as handling static WINS entries, client configurations, database maintenance, and WINS infrastructure security.



EXAM WARNING

Exam 70-293 is mainly concerned with planning a WINS replication environment. However, it is important that you understand the WINS configuration issues involved with WINS replication.

Static WINS Entries


One of the advantages of using WINS is that it provides a way to dynamically register NetBIOS names, eliminating the need for static entries in LMHOSTS files. However, there are situations that require the use of static mappings in the WINS server database. For example, if you have non-WINS clients that are running NetBIOS applications, you might find it desirable to have entries for these clients in the WINS database so that you can allow WINS clients to resolve the NetBIOS names of those clients. Static mappings are superior to entries in an LMHOSTS file because they can be replicated throughout the WINS infrastructure.

The use of static mappings can create problems on your network. Unlike dynamic mappings, static mappings stay in the WINS database until they are manually removed. (The expiration date for the static mapping entry in the WINS database is labeled as *infinite*.) Furthermore, unless the migrate on setting is enabled, static mappings are not overwritten by dynamic mappings. For example, a client computer might be given a static mapping in the WINS database, or an LMHOSTS file might be imported to the WINS database, creating a number of static WINS entries. If the clients that are associated with the static mappings were later configured as WINS clients, they would not be able to perform dynamic registration of their NetBIOS names unless the migrate on setting was enabled. Figure 6.31 shows the **Replication Partners Properties** dialog box where the **Overwrite unique static mappings at this server (migrate on)** setting is enabled.



NOTE

Even though the Migrate On setting can prevent a number of problems associated with the ability to overwrite static entries, this setting does not affect all NetBIOS record types. For example, the domain [1Ch] record type is never overwritten, regardless of this setting.

Figure 6.31 Configuring Static Entries to Be Overwritten

In general, static entries should never be created for WINS-capable client computers. However, for security purposes it is sometimes desirable to use static entries for mission-critical servers to prevent redirection.

Using Static Entries to Prevent Redirection

Unlike with Active Directory-integrated DNS zones, you cannot restrict clients from dynamically registering names according to Windows group memberships. The only mechanism by which WINS prevents clients from registering duplicate names is to send a challenge to the IP address of the active record. If the client responds to the challenge, the duplicate name registration is denied. However, during periods when WINS clients are offline for maintenance or are being rebooted, a rogue computer could register the same name as the original computer with the malicious intent of redirecting traffic to the rogue computer. In high-security environments, it might be desirable to enter static mappings for critical computers and to ensure that the **Overwrite unique static mappings at this server (migrate on)** setting is disabled.

Multihomed WINS Servers

A multihomed WINS server is one that has more than one active network connection. You should avoid this configuration of a WINS server. Name resolution and replication problems are often the result of using a multihomed computer as a WINS server.


The problem is that the WINS service binds to a single IP address. In theory, WINS is supposed to bind itself to the primary IP address of the first adapter listed in the binding order for adapters bound to the TCP/IP. However, Microsoft does not guarantee this behavior. It is important that all adapters be properly configured with routable IP addresses to avoid problems that might arise if WINS binds to an adapter with an improperly configured IP address.

Multihomed WINS servers also introduce potential problems with replication. The issue is that when a replication partner sends a pull request to the multihomed server, it uses a WinSock interface and specifies that the reply can come from any IP address. If the multihomed WINS server replies from an IP address that isn't in the list the replication partner has for the multihomed computer, replication will fail. This is a consequence of specifying the name of the multihomed computer in a replication partner configuration. If the replication partner specifies the NetBIOS name of the multihomed WINS server, it will associate this name with an IP address. If the response to the pull request comes from a different IP address than the expected one, replication will fail. We can't rely on the binding order to determine what IP address will be used on the reply, since WINS will consult its routing table to determine which IP address to use for sending the packet. This behavior is by design. The workaround for this is to configure all replication partners of the multihomed WINS server to replicate with each of its configured IP addresses.

Client Configuration

WINS client configuration is accomplished either through a DHCP server or manually by configuring the settings in the **WINS** tab of the **Advanced TCP/IP Settings** property pages of the TCP/IP properties. Figure 6.32 shows these settings for a Windows XP client.

Figure 6.32 Advanced TCP/IP Settings for WINS Client Configuration



With the configuration shown in Figure 6.32, the client will use an LMHOSTS file if WINS name resolution fails. This is the default configuration. Also, the client is configured to get information about whether NetBT should be enabled from the DHCP server. (This

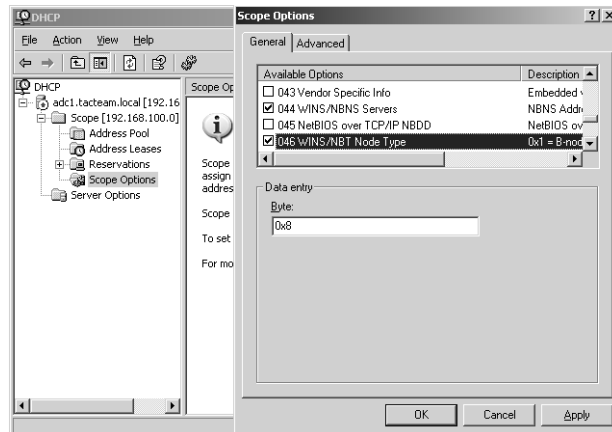
information is provided by a special Microsoft-specific DHCP option.) If the client does not get this information from the DHCP server, it will default to enabling NetBT. You would disable NetBT only when you need to, such as on the Internet-facing interface of a multihomed server or if you have determined that the interface does not need to be able to provide access to NetBIOS applications.

If you are using a DHCP server, you do not need to specify static addresses for WINS servers in this dialog box. If you do configure WINS addresses here, these settings will override those that are supplied by the DHCP server. If you are using DHCP to supply the client settings, you will need to configure two DHCP options:

- **Option 044 WINS/NBNS Servers** You use this option to provide DHCP clients with the IP addresses of WINS servers to contact.
- **Option 046 WINS/NBT Node Type** This option governs the order of NetBIOS name resolution mechanisms that will be used. The hybrid setting option (0x08) is the one most commonly used. With the hybrid node option specified, the WINS client will contact a WINS server before using broadcasts and other methods to resolve names. The mixed node setting option (0x04) forces WINS clients to use broadcasts before contacting the WINS server. This setting is useful in situations where there is a single subnet in a small branch office connected by a slow WAN link. In this case, you might want broadcasts to resolve local NetBIOS names before contacting the remote WINS server. (NetBIOS node types and name resolution were discussed earlier in this chapter.)

Figure 6.33 shows the configuration of DHCP server options to support WINS client configurations.

Figure 6.33 DHCP Options for WINS Client Configurations



Using the nbtstat Command to Verify and Troubleshoot WINS Registration

You can use the `nbtstat -n` command to verify client registration with the WINS server. The Status column in the output of the command will show whether the name is registered, registering, or in conflict with another name. Here is an example of using this command:

```
C:\>nbtstat -n
```

```
Local Area Connection:
```

```
Node IpAddress: [192.168.100.20] Scope Id: []
```

NetBIOS Local Name Table

Name	Type	Status
ADC1	<00> UNIQUE	Registered
WORKGROUP	<00> GROUP	Registered
ADC1	<20> UNIQUE	Registered
WORKGROUP	<1E> GROUP	Registered
WORKGROUP	<1D> UNIQUE	Registered
..__MSBROWSE__.	<01> GROUP	Registered

```
C:\>
```

It is not necessary to restart a client computer to register NetBIOS names. Registration of NetBIOS names can be accomplished through the use of the `nbtstat -RR` (Refresh Release) command, which is invoked from a command prompt. This is useful in situations where you have just configured a WINS address in the TCP/IP properties or changed the TCP/IP address and wish to register the computer name immediately without rebooting.

The `nbtstat` command is also useful for troubleshooting problems with NetBIOS name resolution. You can use it to look at the contents of the NetBIOS remote name cache or to purge and reload the cache. The following are the options that can be used with the `nbtstat` command:

- `nbtstat -a <name>` List the NetBIOS name table for the remote computer specified.

Continued

- **nbstat -A <ip address>** Same as above, except the IP address is specified.
- **nbstat -c** List the names in a NetBIOS remote name cache.
- **nbstat -n** List the NetBIOS names in use by the computer. This will show the NetBIOS names that have been registered by WINS.
- **nbstat -r** Show statistical information and a list of names that have been resolved by broadcast and the WINS server.
- **nbstat -R** Purge and reload the name cache.
- **nbstat -RR** Refresh and renew NetBIOS registrations.
- **nbstat -s** List the currently open NetBIOS sessions by name.
- **nbstat -S** List the currently open NetBIOS sessions by IP address.

The `nbstat` commands are case-sensitive, so you need to make sure that you use the correct case to receive the desired results.

Multiple WINS Server Addresses

If you specify multiple WINS server addresses in the client configuration, the client will try to use the first WINS server in the list for registration and name resolution requests. If the primary server fails to respond, the client will then attempt to contact the alternate WINS servers in the order listed. Up to 12 WINS servers can be specified for Windows XP and Windows 2000 clients.

WINS Proxy Agent

For the rare case where you have a NetBIOS client that is not capable of querying a WINS server for NetBIOS name resolution, you can set up a WINS proxy agent. The WINS proxy agent is a WINS client that is set up on a subnet to provide limited WINS support for b-node and non-WINS NetBIOS clients. A WINS proxy agent listens for name registration and name query broadcasts, and it forwards these to its configured WINS server. This process ensures that the b-node client does not initialize with a duplicate name that is already registered in the WINS database and provides name resolution on behalf of the b-node client.

A common misconception is that a WINS proxy client will register the name on behalf of the non-WINS client. This is not the case. The WINS proxy merely contacts the WINS server to verify that the non-WINS client name does not exist. If there is a duplicate name in the WINS database, the WINS proxy client will send a negative response to the b-node client.

The proxy agent will use its NetBIOS name cache to temporarily store the results of responses to its queries to the WINS server. Performance of the WINS proxy agent

could therefore, be potentially improved by lengthening the amount of time an entry would persist in cache beyond the default 10 minutes.

Preventing Split WINS Registrations

A WINS server is also a WINS client to itself. A common configuration error is to specify a different WINS server to use as an alternate WINS server. The problem with doing this is that during startup of the WINS server, it will try to register its names with the WINS servers configured in the TCP/IP properties. Because the WINS server service won't start until NetBT has initialized, this causes the WINS server to attempt to register its [00h], [03h], [20h], and other mappings (depending on any additional services present on the WINS server) with the alternate or secondary WINS. However, it will continue to try to register these mappings in the local WINS database. Once the local WINS database is available, the WINS server will switch to it to register the remaining mappings. This results in what is known as a *split registration*, where name mappings are registered in and owned by two different WINS servers. The remote WINS server configured as a secondary might own the [00h] and [03h] name mappings, whereas the local WINS server owns the [20h] name mapping.

A split registration for WINS servers can cause intermittent problems with WINS name resolution. For example, suppose that one WINS server owned the mapping for the file server service [20h] and another WINS server owned the mapping for the workstation service [00h]. Clients that used the first WINS server would be able to connect to file shares on the server, while users who used the second WINS server would not be able to do this. If the WINS servers replicated with one another, eventually the records would converge and the symptoms would disappear. However, the underlying problem would remain.



NOTE

In configuring the WINS client properties on the WINS server, you should specify only the IP address of the WINS server itself as the primary and secondary WINS server. However, even if you specify no WINS server IP addresses in the WINS client configuration, WINS will eventually register its mappings in the local WINS database.

Performance Issues

As mentioned earlier in the chapter, a typical WINS server can handle WINS registrations and name resolution requests for up to 10,000 clients, even if the WINS server has only modest amounts of CPU power and RAM. WINS traffic for each registration and name resolution request is relatively small. However, a number of factors can affect the performance of WINS server. These factors include the presence of other services running on the WINS server, the performance of database maintenance on the WINS server, various WINS

server settings, and the flooding of the network with NetBIOS name registration requests. In this section, we'll look at WINS server performance issues.

Hardware Considerations

Even a modestly powered computer can handle a large number of registrations and name resolution requests. However, WINS can generate intensive CPU and disk activity. WINS server performance is therefore, significantly improved by using fast disks and multiple CPUs. For mission-critical WINS servers that handle large amounts of data, it is a good idea to use RAID arrays to enhance fault tolerance and performance. If you don't use a RAID array, you should consider placing the WINS database on a separate hard disk from the operating system. You can specify the location of the database on the **Advanced** tab of the property pages for the WINS server (see Figure 6.34 in the next section).

Moving the database to a separate spindle will improve disk I/O and the performance of the WINS server for not only name registration and renewal, but also for the performance of other tasks related to the maintenance of the database. You should be aware, however, that if you move the database files to a different location they will lose their security protection, and you will need to take additional measures to secure access to the database. We will discuss this issue later in this chapter.

If the organization requires a very high degree of fault tolerance and availability, you should consider using a Windows Server 2003 server cluster for the WINS service. This will require at least two WINS servers configured as cluster nodes to use a shared SCSI or Fibre Channel storage device. On large networks, adding more WINS servers to distribute the load will improve response times for WINS queries, but will add more replication traffic and require more administration.

To verify the need for additional or more powerful hardware, it is useful to take baseline measurements of the WINS servers using the Windows Server 2003 Performance console. When you install a WINS server, approximately 15 performance-related counters are installed to gather data on the performance of your WINS server. You should regularly monitor these counters, along with those that measure overall system and network performance, to establish baselines that will help you determine the need to upgrade or add hardware. Looking at WINS counters *without* considering them in the overall context of system performance can potentially cause you to draw incorrect inferences with regard to the performance of the WINS servers. Some of the relevant WINS-specific counters you might wish to look at include the following:

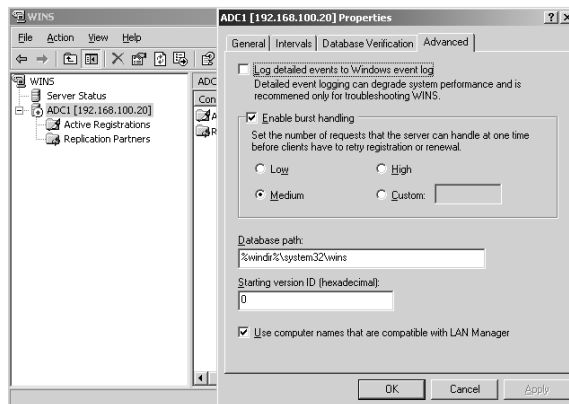
- **Queries/sec** Measures the rate at which the WINS server receives queries.
- **Total Number of Renewals/sec** Measures the rate at which the WINS server receives renewals and is the sum of unique and groups renewals per second.
- **Total Number of Registrations/sec** Measures the rate of registrations per second and is the sum of unique and group registrations per second.

Burst Handling

Since Windows NT 4.0 Service Pack 3, WINS servers have been capable of *burst handling*. Burst handling, which is enabled by default, allows the WINS server to handle a large volume of simultaneous registration requests. This situation can occur, for example, when power is suddenly returned to many computers after a power outage. With burst handling configured, the WINS server will respond positively to name registration and refresh requests before it writes them to the database. However, it will supply the WINS clients with varied and short TTLs for the name registrations to stagger the load for subsequent WINS client refresh attempts.

By default, burst handling occurs when the server has more than 500 requests in its queue. However, you can adjust this setting in the WINS console. As shown in Figure 6.34, the **Advanced** tab for the properties of the WINS server allows you to select a **Low** (300 requests), **Medium** (500), **High** (1000), or **Custom** (where you specify the number of requests). If the WINS server has more than 25,000 requests in its queue, it will start dropping queries.

Figure 6.34 Configuring Burst Handling



It's a good idea to ensure that burst handling is enabled for reasons other than improving the performance of the WINS server under peak-load conditions. When burst handling is enabled, it writes events to the event log. The presence of burst handling events can provide an indication that the WINS server hardware is not adequate. Furthermore, the presence of burst handling events can indicate a possible DoS attack on the WINS server. You can use Network Monitor or some other tool for analyzing network traffic to capture packets and track down the possible causes of the presence of burst handling events.

Scavenging of WINS Records

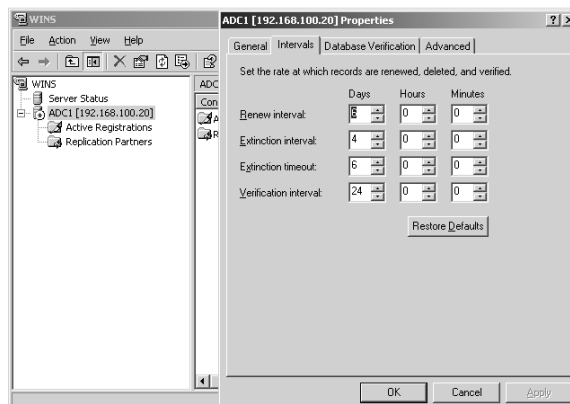
Performance of a WINS server will be affected by the settings that are used to determine how frequently clients refresh their registrations, how long it takes for a released or deleted registration to be removed from the database, and how frequently the database verifies its

records to ensure the integrity of data. *Scavenging* is the process by which WINS records are removed from the database. More specifically, scavenging is a preset process that periodically runs on the WINS server and either deletes or changes the status of WINS records based on their timestamps and current state. For example, the status of an active record that has an expired TTL is changed by the scavenging process to a status of released. The status of a released record, in turn, is changed to tombstoned after the extinction interval for the released record has expired. The tombstone record is deleted (scavenged) from the database after the extinction timeout period has elapsed. The settings that control these intervals are found in the **Intervals** tab in the property pages of the WINS server. It might be useful to change these settings to improve the performance of the WINS server:

- **Renewal interval** Governs the TTL of the client registration. WINS clients will attempt to renew registrations after half the renewal interval has elapsed. Increasing this interval will reduce the frequency of client renewal attempts and reduce the load on the WINS server. However, increasing the interval also makes the database less consistent with the network over time when computer names are changed. The renewal interval should be the same for all WINS servers when they are replicating with one another.
- **Extinction interval** Governs the period that must elapse from when a name is marked as released and when it is marked as tombstoned.
- **Extinction timeout** Governs the period that must elapse from when a name marked as tombstoned and is subsequently scavenged (removed) from the database.
- **Verification interval** Dependent on the previous values and governs when a WINS server must validate active records it does not own; that is, records learned of via replication with other WINS servers.

Figure 6.35 shows the Intervals tab with the default settings.

Figure 6.35 Interval Settings for Registration Renewal, Removal, and Verification



Database Compaction

When records are deleted from the WINS database, the space formerly occupied by them should be recovered to ensure optimal performance of the database. The process of recovering this space is referred to as *compaction*. The WINS service automatically and periodically performs online compaction of the WINS database. However, online compaction of the WINS database is not as efficient as offline compaction. It is, therefore, sometimes desirable for the WINS administrator to stop the WINS service (take the database offline) and perform a manual compaction of the database.

The WINS administrator can use `Jetpack.exe`, found in the `System32` folder, for manual database compaction. The `Jetpack` utility works by creating a temporary database in which to compact the records, and then replacing the original database with the compacted one. To manually compact the database, the WINS administrator must first stop the WINS service and then issue the **jetpack** command, using the following syntax:

```
jetpack %systemroot%\system32\Wins\Wins.mdb [name_of_temp_database.mdb]
```

After running the `jetpack` command, the WINS administrator can start the WINS service again.

Using the **net stop** and **net start** commands, the WINS administrator can automate offline compaction in a batch file. For example, you can create a simple batch file that contains the following three lines:

```
net stop wins
jetpack %systemroot%\system32\Wins\Wins.mdb [name_of_temp_database.mdb]
net start wins
```

Once you have created the batch file, you can configure it to run at preconfigured intervals using the Task Scheduler or the AT command-line utility. For example, you could configure the batch file to run once a month during off hours to ensure that the database uses space optimally.

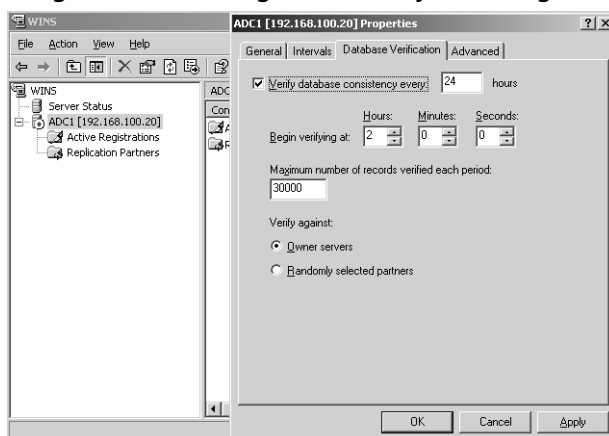
Scheduling Consistency Checking

In order to maintain database integrity in environments that employ replication, it is recommended that automatic periodic consistency checking be enabled. Consistency checking is the process whereby a local WINS server compares local entries that it has acquired by replication with the entries in WINS servers that own the record. If a WINS server detects that the records are identical between its locally stored copy and the remote database, it will update the record with a new timestamp. However, if the record has a lower version ID in the local database, it will pull the updated record from its replication partner and mark the original one for deletion.

Because consistency checking puts a significant load on the resources of a WINS server, it should be scheduled to run during off-peak hours. Figure 6.36 shows the

Database Verification tab of the property pages of the WINS server, where you can set the schedule for consistency checking.

Figure 6.36 Enabling and Scheduling Consistency Checking



In Figure 6.36, consistency checking is enabled and scheduled to run every 24 hours at 2:00 A.M. Additional settings allow the WINS administrator to specify the maximum number of records to verify and to select randomly selected replication partners. Consistency checking can also be manually initiated from the context menu of the WINS server in the WINS server console.

Security Issues

As with any service that you implement on your servers and your network, it is important for you to understand the service and take measures that mitigate the risk to the service and the network as a whole. These measures include setting up restricted ACLs, logging, auditing, and monitoring, as well as using VPNs or IPSec to secure WINS replication traffic. In the next section, we briefly examine issues related to the security of the WINS service.

NetBIOS Security Issues

With regard to NetBIOS in general and the WINS service in particular, administrators need to be aware that NetBIOS is an unauthenticated protocol. That is, users are not required to submit credentials before using the services provided by a WINS server such as name registration, renewal, release, and queries. This makes WINS susceptible to a number of different kinds of attacks, primarily DoS attacks and redirection attacks.

In a DoS attack, an attacker attempts to tie up the WINS service with a large number of requests that compromise the WINS server's ability to process legitimate requests. To mitigate the risk of DoS attacks, you should do the following:

- Secure the physical network from unauthorized access.

- Enable burst handling. When burst handling is enabled, burst handling events are recorded in Event Viewer, providing an alert to a possible DoS attack. (This is set in the Advanced tab for the properties of the WINS server, shown in Figure 6.34 earlier in this chapter.)
- Enable detailed WINS logging to provide more complete and specific logging of WINS events in the System log. (Also set in the Advanced tab for the properties of the WINS server, shown in Figure 6.34.)
- Use a protocol-analysis tool, such as Network Monitor, to analyze traffic in the case of a suspected attack.

In a redirection attack, an attacker tries to register a rogue computer that has the same name mappings as a previously registered computer. If the previously registered computer is down for maintenance or is otherwise unable to respond to the challenge from the WINS server (for example, if it is also a victim of a specific DoS attack), the rogue computer will be able to register the name mappings with its own IP address. WINS clients will then subsequently be redirected to the rogue computer. To mitigate the risk of redirection attacks, you should do the following:

- Identify mission-critical systems and assign them static mappings in the WINS database.
- Ensure that the migrate on setting is disabled to prevent the WINS server from overwriting the static mappings with dynamic mappings. (This is controlled by the **Overwrite unique static mappings at this server (migrate on)** setting in the **Replication Partners Properties** dialog box, shown in Figure 6.31 earlier in this chapter.)

Protecting the WINS Database and Log Files

The WINS databases and log files contain important information about your network. The information in these files could be used by attackers to glean confidential information about your company. For example, by an analysis of the number of computers, the attacker could learn the names of those computers, the NetBIOS applications running on the computers, and so on. Furthermore, the integrity and availability of the WINS database and the log files are critical to the operation of the network.

It is particularly important to note that, if you change the default location of the WINS database files, they will inherit the ACLs of the new destination folder, removing the effective security they inherited by virtue of being located in a subfolder of the System32 folder. Also, the WINS backup files inherit the ACL of the folder used to store the backup files. To mitigate the risks to the confidentiality, integrity, and availability of WINS databases, you should consider doing the following:

- Do not store WINS database files on anything except an NTFS formatted partition.
- Ensure that the ACLs for the WINS database, backup, and log files are restricted to allow access to only the Local System Account and the Administrators group.
- Enable file auditing on the WINS files to track attempts of objects that try to access these files.
- Ensure that the WINS server is physically secured from remote access.
- Do not transfer WINS database files over the network using FTP or other unsecured protocols.

WINS Users Group

Only members of the Administrators group can modify the settings for WINS servers. However, it is desirable, in some situations to provide read-only access to the configuration and database information of WINS servers. To provide users with read-only access to the WINS server, you can use the special WINS Users group. Users who are members of this group are able to query the WINS server database to find records and to view configuration information.

Planning for WINS Database Backup and Restoration

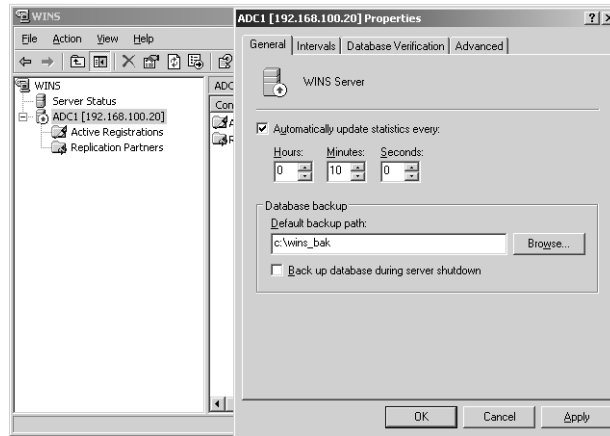
Although it is possible to rely on replicated copies of WINS records to restore a corrupted database, you should only do so as a method of last resort. It is far better to back up the WINS database, either manually or automatically, and restore WINS records from the backup, if necessary.

By default, the WINS database is not backed up automatically. To back up a WINS database automatically, you need to do the following:

1. Specify a backup directory in the **General** tab of the WINS server properties. (WINS will create a folder called **Wins_bak** under this folder to store the backups.)
2. Perform a manual backup of the WINS database to the specified location. You can do this by choosing the **Backup Up Database** option from context menu of the WINS server object in the WINS console tree.

After you have performed these steps, backups will occur every 24 hours or upon service shutdown (if so configured). In the event that the WINS service detects a corrupt WINS database upon startup, it will automatically restore the backed up version from the location that you specified for the WINS backups. You cannot use a network drive for this location. Figure 6.37 shows the settings for configuring automatic backups.

Figure 6.37 WINS Backup Configuration

**NOTE**

As part of a comprehensive restoration policy, you should also ensure that you have recent backups of the Registry, such as a backup of the System State.

To restore the WINS database, you can stop the WINS service, delete the original database, and restore the backup. If you do not stop the WINS service, the option to restore the WINS database will be grayed out.

EXAM
70-293
OBJECTIVE
2.5.2

Troubleshooting Name Resolution Issues

Proper name resolution is critical to the smooth operation of the network. When name resolution fails, for whatever reason, users might be inconvenienced, and connectivity to critical systems might be compromised. Usually, a name resolution failure requires immediate action on the part of the administrator, even if the failure is not widespread.

Often, problems that appear to be related to name resolution are, in fact, the result of problems that occur further down in the OSI or DoD networking models for TCP/IP. For example, if a router fails and the DNS or WINS servers are on the other side of the router from the client, clients will not be able to resolve names. The failure of a router occurs at Layer 3 (Network layer) of the ISO model.

As part of a prudent and successful troubleshooting strategy, it is important to troubleshoot from the bottom of the OSI model up, ensuring the following:

- The hardware is functioning properly.
- The computer is configured properly.
- The computer is able to communicate with hosts on the local subnet.


- The computer is able to communicate with the router configured as the default gateway.
- The computer can communicate with remote networks on the far side of the router.

Usually, this troubleshooting will rely on tools such as Ipconfig to verify configuration and PING (using IP addresses) to test communication.

Name resolution occurs further up in the OSI or DoD model. In a Windows environment that relies on both DNS and NetBIOS (which use the WinSock and NetBT interfaces, respectively), it is sometimes necessary to distinguish which interface used for name resolution is causing problems. You need to discover if is the failure specific to NetBIOS or DNS (host name resolution). To make this determination, you can test name resolution using applications that are specific to these interfaces. For example, to test NetBIOS name resolution, try to connect to shares using a UNC path or the net commands, invoked from the command line. As long as the name you are trying to resolve does not include dots or is less than 15 characters long, Windows clients will default to using NetBIOS resolution first. If name resolution is successful, but you encounter a delay, chances are that NetBIOS resolution failed but that the fallback host name resolution was successful. To test host name resolution, you can use WinSock applications, such as NSLookup, Telnet, FTP, HTTP, and so on.

You should also consider whether NetBT is enabled, either by means of a DHCP server option or as a configuration on the client computer. Some applications, such as Microsoft Exchange Server, still require that NetBT be enabled to work properly. In troubleshooting name resolution problems, you should therefore take into account the applications that may be involved and their dependence on either the WinSock or NetBT interfaces to work properly.

It is sometimes obvious that the problem is specific to either host name or NetBIOS name resolution. In any event, after you have made the determination, you can proceed to troubleshoot according to the interface (WinSock or NetBT) that is involved.



Troubleshooting Host Name Resolution

Assuming you have eliminated any antecedent causes that have to do with connectivity and communication on the network, troubleshooting host name resolution is easier if you can isolate whether the problem is caused by problems with client configuration or by problems with the DNS server. Problems with DNS clients include improperly entered addresses for the primary and secondary DNS server, in addition to improperly configured DNS suffix search list settings. Problems with DNS servers include improperly configured delegations; improperly configured restrictions on zone transfers; missing, incorrect, or stale resource records; and so on.

Effective troubleshooting of Microsoft DNS issues requires a familiarity with the process of DNS name resolution (for example, recursive versus iterative queries and authorita-

tive versus nonauthoritative responses), dynamic updates, zone transfers, stub zones, forwarding, and so on.

A familiarity with DNS-related troubleshooting tools such as NSLookup, Ipconfig, Dnscmd, and DNSLint, will help to ensure that you can trace the source of the problem effectively. Using NSLookup, you can request either an authoritative or nonauthoritative response from the DNS server, which can help you to narrow down the problem further; for example, to determine whether a stale record is coming from the DNS cache or not. Furthermore, with NSLookup, you can use debug mode to provide a great amount of detail in the output of the command. You can use DNSLint to check all your delegations and verify the correct configuration of well-known services such as SMTP. Issues with dynamic registration can sometimes be resolved by using the command **ipconfig /registerdns**. Incorrect entries in the client DNS cache can be resolved with the **ipconfig /displaydns** and **ipconfig /flushdns** commands. Some tools that are not specific to DNS, such as Nltest, can help to troubleshoot DNS-related issues with domain controllers.

Issues Related to Client Computer Configuration

Many problems with DNS resolution have their origins in the client configuration, so verifying the correct client configuration is a good place to begin. To troubleshoot problems with client configuration, use the **ipconfig /all** command to verify the DNS configuration, and then use PING and NSLookup to verify communication with the DNS server. If you can ping the DNS servers but an NSLookup query against them fails, the issue is most likely related to a problem with the DNS servers. (This is the kind of situation where troubleshooting from the bottom to the top of the OSI model really pays off in helping to narrow down the problem.)

One of the more common and serious problems with client configuration is an improperly configured FQDN, which is set up in the properties of My Computer on Windows 2000 and Windows XP clients. If the FQDN is not present or is incorrectly configured on the client computer, name resolution can fail for domains that would otherwise be searched according to a domain suffix search list. For example, if the computer is a member of the corp.tacteam.net domain and the user enters and uses an unqualified name (for example, PServer1) in a DNS query for a host in the tacteam.net domain, the query will fail unless the FQDN of the computer is properly configured. (An unqualified name is one that doesn't have a trailing dot.) By default, the DNS client uses a domain suffix search list based on the FQDN. So, if the host name cannot be resolved in the corp.tacteam.net domain, the suffix will be devolved to tacteam.net to find the host. To troubleshoot problems with domain suffix search lists, try to resolve the name at the client using the FQDN (that is, include the trailing dot). If this query succeeds but an unqualified query does not, the problem is related to the domain suffix search list.

Clients that have improperly configured FQDNs might also have problems with dynamic registration in the DNS zone. The client registration of a host record in a DNS requires that the primary suffix be properly configured. If the domain suffix is improperly configured, the client may be trying to register in a nonexistent or an unintended domain.

To troubleshoot and resolve this problem, verify that the client computer is correctly configured with a primary domain suffix and that it can reach a name server that is authoritative for the domain name (you can simulate this by using NSLookup to perform an SOA RR query type for the authoritative zone). If the client receives its TCP/IP configuration from a DHCP server, verify that the DHCP server option for the domain suffix and other settings are configured correctly. Also, the client needs to be configured with the IP address of a DNS that contains the primary zone in order for the updates to occur. Meeting these conditions and then using the **ipconfig /registerdns** command to register the host in the domain might solve the problem. However, if the problem is still not resolved, chances are the source of the problem is the DNS server; for example, an ACL on the RR may be preventing the update. Using Event Viewer on the client computer can help you determine the nature of the problem.

If the DNS clients are getting incorrect responses to DNS queries, the problem might be related to their DNS cache. To clear the cache and force a new query to a DNS server for the host name, use the command **ipconfig /flushdns**. (Alternatively, you can use NSLookup to request an authoritative response.) If you have cleared the cache and are still getting incorrect responses to queries, it is likely the source of the problem has to do with the DNS server; for example, there might be an incorrect entry in the cache of the DNS server, a problem with zone transfers, or a delay in AD replication.

EXAM
70-293
OBJECTIVE
2.9.1

Issues Related to DNS Services

If you have determined that the problem you are experiencing is unrelated to DNS client settings or the client's ability to communicate with the network, the problem is most likely related to DNS server configuration. To troubleshoot DNS server problems, you need to let the symptoms of the problem guide you to a likely cause and solution. Again, using a tool like NSLookup will help you get a clearer picture of the problem since it can provide more detailed information and be used to get either authoritative or nonauthoritative responses from DNS servers. The following is a brief list of guidelines to help troubleshoot problems with DNS servers:

- If clients cannot resolve names on the Internet or in domains for which their DNS servers are not authoritative using recursive queries, the problem might be related to the ability of the DNS server to perform recursion or to forward the query to a server that will perform recursion for it. In this case, check to make sure that the root hints file is present and the records are correct. If these settings are correct, your DNS server might be experiencing cache pollution. In this case, you should enable protection against cache pollution and restart the DNS service.
- If clients do not get correct records or are getting stale records, the cause could be the cache on the DNS client or DNS server. Examining and clearing the cache will eliminate the problem. However, if the problem is unrelated to cache, the cause could be failed or slow zone transfers to the secondary zones. If you are using Active Directory-integrated zones, the cause could be related to problems

with AD replication. Comparing the RRs in the various zone files and looking at events in Event Viewer will help to confirm zone transfers as being the cause or the problem.

- If some clients cannot get responses to DNS queries from a multihomed DNS server but others can, the problem might be related to the listener settings on the DNS server. These settings can restrict which interfaces the DNS service will use to respond to queries.
- If you are implementing a round-robin configuration to provide load balancing and are not getting the desired results, check the settings for subnet prioritization. Round robin is a kind of load balancing that can be configured using multiple host records that have the same name but different IP addresses. When round robin is enabled, the DNS server will rotate responses among all the records. However, if subnet prioritization is also enabled, the DNS server will try to respond with a record that is on the same subnet as the DNS client.
- If zone transfers are failing between DNS servers, the cause could be improper restrictions on DNS servers that are authorized to pull zone information from the primary server. By default, the DNS server will allow zone transfers to only those DNS servers listed as name servers for the zone. However, you might need to reconfigure these restrictions so that you specify the IP addresses of computers that are authorized to pull zone information.
- Another cause of failed zone transfers is the use of nonstandard characters in DNS names. By default, Microsoft DNS servers are configured to load the zone even if they encounter bad data. However, BIND servers are not as forgiving. In addition, WINS forward and reverse lookup records can cause problems if replicated to BIND servers. You can prevent WINS records from replicating to BIND servers. If you are replicating to BIND servers, you should use only standard DNS characters.
- Another common cause of zone transfer problems is an incorrect version number in the SOA of the primary or secondary zone. To determine whether to request a zone transfer from the primary server, the secondary server will compare the version number of the primary's SOA with its own. If the primary's number is higher, the secondary will request either a full or an incremental zone transfer. If the version number is reset on the primary so that it is lower than the version number in the secondary's SOA, the zone transfer will fail.
- If queries to subdomains are failing, the cause is most likely a lame delegation of authority. A lame delegation occurs when the name server and glue address records do not point correctly to the servers that are authoritative for the subdomain. NSLookup and DNSLint are useful tools in helping to troubleshoot problems with delegations.


- If dynamic updates are failing, the cause of the problem might be related to the security settings and ownership of RRs in their ACLs. For example, if a DHCP server is the original owner of a record and a client subsequently gets its IP address from another DHCP server, the dynamic update will fail. Another cause of failed dynamic updates is that the primary zone is, for whatever reason, unavailable. Dynamic updates can occur in only primary or Active Directory-integrated zones.

Troubleshooting NetBIOS Name Resolution

To avoid problems with NetBIOS name resolution in the first place, you should take very seriously the best practices that Microsoft recommends for the deployment of WINS servers and clients. In general, these best practices require the following:

- Be conservative in your estimates of the number of WINS servers you need.
- Use full replication partnership agreements.
- Use a hub-and-spoke replication topology to reduce convergence time in large environments.
- Do not install WINS on a multihomed server.

To troubleshoot problems with NetBIOS name resolution, you should first analyze the problem to determine whether it is a client configuration problem or a problem related to WINS server records or WINS server configuration, such as failed replication.



Issues Related to Client Computer Configuration

First, you should determine whether a problem that appears to be related to client configuration affects a single computer or a group of computers that all get their TCP/IP configuration from the same DHCP server or the same DHCP server scope. You should also verify the WINS server configuration by using the **ipconfig /all** command. The output of this command will list any WINS servers that are either manually configured or dynamically configured through DHCP. You should ping the IP addresses of the WINS servers listed in the client configuration to verify that communication is possible with these computers.

Another command you can use to troubleshoot problems related to client configuration is the **nbtstat** command. With the nbtstat command, you can cause a release and refresh of the NetBIOS registration for the client computer, view the remote name cache, view statistics on recent NetBIOS name resolution activity, and so on. Sometimes, a recently cached but incorrect entry in the remote name cache is causing a specific problem. You can also use the nbtstat command to clear the contents of the cache, except for those entries that are preloaded with the #PRE tag in an LMHOSTS file (these entries are obvious when you view the remote name cache using the nbtstat command).

In addition to verifying the correct configuration of the WINS server entries, you might want to consider whether the client is configured as an h-node, an m-node, a b-node, or a p-node client. For example, if the client is configured as an m-node client, it will use name query broadcasts before reverting to unicast name queries to the WINS server. If there is a duplicate NetBIOS name on the subnet, resolution to this name will occur first in the case of an m-node client.

Furthermore, you should consider the presence of an LMHOSTS file on the client computer and the order in which LMHOSTS will be used in name resolution queries. If the clients are using an LMHOSTS file and it appears that an LMHOSTS file is involved in the problem, you need to verify that the entries in the file are correct.

Issues Related to WINS Servers

In troubleshooting problems with name resolution that involve the WINS server, it is useful to first determine the scope of the problem. For example, does the problem involve dynamic or static name mappings, deleted records, replication, or a corrupted database? You should also consider any error messages that the NetBIOS client receives, such as “Network path not found” or “Duplicate name.” In addition, you should look at any events that are recorded in the System event log for the WINS service that might provide an indication of a corrupted database or problems with replication. Finally, you should confirm whether the problem affects one or multiple WINS servers. If the problem affects only one WINS server, you should first verify that the WINS service has started properly and that the database is not corrupted.

Problems Related to Static Mappings

You should avoid the use of static mappings except in situations where you need WINS to provide resolution to NetBIOS applications running on non-WINS clients or you want to provide static, permanent name mappings to mission-critical servers to mitigate the risk of redirection attacks. However, if you are using static mappings and the problem is related to these entries, you should do the following:

- Verify that the entries are correct and that they have replicated properly.
- If you have deleted the static mapping, you need to verify that the tombstoned record, in the case of a tombstone rather than simple deletion, has replicated properly.
- If the client error message refers to a “duplicate name” and there is a static mapping for the name, you need to ensure that the migrate on setting is enabled to allow the dynamic name registration to overwrite the manual registration.

Problems Related to Multihomed WINS Servers

Multihomed WINS servers are so often the cause of WINS-related problems that they deserve a special troubleshooting category. In general, you should avoid installing WINS on a multihomed server. Some of the problems you might experience with multihomed WINS servers can be hard to track down. If you are experiencing intermittent problems with name resolution or if you are having problems with replication, chances are that these can be traced to the configuration of the multihomed WINS server.

However, if you must use a multihomed WINS server and if you are experiencing problems, you should do the following:

- Verify that all network devices on the multihomed WINS servers are configured as routable interfaces with correct TCP/IP information. (You should never collocate the WINS service with the RAS service—that is just asking for trouble.)
- Verify that all TCP/IP configurations use the IP address of the WINS server for both their primary and alternate WINS servers. (You can leave this configuration blank if you like, because the WINS server will register itself without this configuration.)
- Verify that all the replication partners of the multihomed WINS servers are configured to replicate with *all* the configured IP addresses of the WINS server, and not the NetBIOS name.

Problems Related to Replication

Problems related to replication almost always are the result of not following Microsoft's recommended practices, such as installing too many WINS servers, installing WINS on a multihomed computer, or using limited replication partnerships. For example, installing too many WINS servers (more than 20, according to Microsoft) can cause intermittent and hard-to-locate problems with replication.

In troubleshooting replication problems, you should first consider whether the problem is related to network communication and name resolution to the replication partners themselves. Consider the following questions:

- Can you ping the IP address of the replication partner?
- If the replication partner is a multihomed computer, have you configured the replication partner settings with the IP addresses of the multihomed computer, rather than the NetBIOS name?
- Does the NetBIOS name of the WINS server resolve to the correct IP address?

If you are using limited replication partnerships (push-only and pull-only) replication partners, you should ensure that these partnerships are set up correctly. Also, you might achieve best results by setting up reciprocal partnerships on the push and pull partners. For example, a computer that is configured as a pull-only partner to another WINS server should also configure that WINS server as its push partner. To illustrate, WINS-A has configured WINS-B as its pull partner; WINS-B in turn should configure WINS-A as its push partner. (To ensure that records are never pushed and replicated strictly according to the pull replication schedule, you can set the push trigger threshold to a very high number that will never be reached between pull replication cycles.)

If replication partnerships are configured correctly and there is good connectivity, but you are still experiencing intermittent problems, the version IDs on some replicated records may not be correctly incremented. You can resolve this problem by entering a new starting version ID for the WINS database in the WINS console or using the netsh command.

Summary of Exam Objectives

In this chapter, we discussed planning and configuring a Windows Server 2003 network to support host name and NetBIOS name resolution. Name resolution is the process of resolving names to IP addresses to enable communication between computers. Because names are easier to remember than numbered IP addresses, name resolution makes it easier for users to connect to resources on the intranet or Internet.

Host name resolution occurs whenever name resolution occurs through the WinSock interface. This will occur whenever a WinSock application, such as Internet Explorer, FTP, and Telnet, requires name resolution or when a name that needs to be resolved to an IP address is longer than 15 characters or includes dots.

NetBIOS name resolution occurs whenever the NetBT interface is used. This will occur, for example, when a user connects to a file share using a UNC-formatted name such as `\\computer\sharename`, or when the user issues a net command, such as **net view**. Some applications are written specifically for the NetBIOS interface and will use NetBIOS name resolution.

DNS is a requirement for AD and is responsible for ensuring that resources, such as domain controllers, can be located by other computers. The DNS namespace is based on the domain tree of AD, and the hierarchical DNS domain namespace mirrors the hierarchical AD domain tree. To support AD, the DNS server must be able to use SRV records, which are a special kind of RR that allows computers to locate services running on the network. Ideally, the DNS server should also support dynamic updates and incremental zone transfers.

On the Internet, the domain name space starts at the root, usually represented by a dot (.). Below the root are the top-level domains, such as .com, .net, .org, .ca, .de, .biz, and so on. Below these domains are the second-level domains, such as `syngress.com` or `tacteam.net`. These second-level domains can have additional subdomains, such as `corp.tacteam.net`. The FQDN for a host in a particular domain, such as `www.corp.tacteam.net.`, represents the complete and contiguous path from the root to the host when read from left to right.

Authority for portions of the domain namespace is distributed among multiple DNS servers. A DNS server that is authoritative for a portion of the domain name space stores DNS RRs in a special file called a zone file. A zone file contains the RRs for hosts in the portion of the domain namespace for which it is authoritative. A zone begins with a special RR called an SOA record. The SOA record identifies the names and IP addresses of DNS servers that contain the authoritative RRs for their portion of the domain namespace. For example, two DNS servers could be authoritative for the entire `tacteam.net` namespace, including its subdomains.

It is important to understand the difference between a domain and a zone. A zone file can be authoritative for a single domain, or it can be authoritative for a parent domain and multiple child domains. In the latter case, the zone file contains all the records for the parent and child domain. If there is a delegation of authority from a parent domain to a

child domain, the DNS server that is authoritative for the parent domain will have (in addition to the RRs for the parent domain) only the name server and A records that serve to delegate authority to DNS servers authoritative for the child domain.

On large, complex networks, the DNS infrastructure can be optimized by delegating authority to multiple DNS servers that are authoritative for their respective child domains. However, incorrect delegation of authority, (sometimes referred to as *lame delegation*) is a common source of DNS-related problems. Also, delegation of authority of a portion of the domain namespace may involve a security trade-off in that administration of those zones may be handled by different administrators.

Regardless of whether the domain namespace is used for name resolution on the Internet or exclusively on your intranet, the domain name should be unique to your organization. Even if you intend the domain namespace to be used exclusively on your intranet, it should be based on a name that is registered for use on the Internet. By definition, names that are registered for use on the Internet through the agency of ICANN-approved registrar are unique. Choosing a unique name will mitigate any rare, but possible, problems that might arise if two organizations that use the same domain name merge.

The choice of a domain namespace for the internal network will have an effect on security of the network, as well as on the administrative effort required to maintain the DNS infrastructure. Regardless of the choice you make, the DNS infrastructure needs to be designed so that RRs that point to hosts on your internal network are never accessible to public DNS servers on the Internet. Your security needs may also include a requirement that DNS queries to the Internet are not allowed. Furthermore, firewalls and routers should be configured to restrict Internet access of internal DNS servers. The choice of a domain namespace can be used to enforce this requirement. In general, the choices are as follows: same domain name for intranet and Internet, subdomain of Internet domain name for intranet, or disjointed namespace for intranet.

To properly plan and implement a DNS infrastructure, it is important to understand the various DNS zone and server types that can be deployed on your network. Zone types include primary, secondary, stub, and Active Directory-integrated zones. Server types include caching-only, forwarding, and nonrecursive.

A carefully designed DNS infrastructure will ensure the fault-tolerant availability of DNS data, while at the same time, limit the bandwidth consumed by zone transfers and other DNS-related traffic, as appropriate. To ensure fault tolerance and eliminate single points of failure, it is recommended that at least two DNS servers (a primary and a secondary or two Active Directory-integrated zones) be configured for each zone of authority. Ideally, these servers should be on separate subnets that are served by different routers.

Standard zone transfers from primary to secondary zones can take either of two forms: the older, full zone transfer (AXFR) or the newer and more efficient incremental zone transfer (IXFR). During a full zone transfer, the entire zone file is copied from the primary to the secondary. During an incremental zone transfer, only the incremental changes since the last update are copied. Zone transfers are triggered either by a notify message sent from

the primary to the secondary when changes have occurred to the zone or according to the refresh interval in the SOA.

Active Directory-integrated zones do not use standard zone transfer mechanisms, unless a secondary server is configured to transfer zone information from an Active Directory-integrated zone. Active Directory-integrated zone transfers use multiple-master AD replication and as such provide a more efficient and secure mechanism for transferring zone information.

With Windows Server 2003, it is now possible to fine-tune the scope of replication for Active Directory-integrated zones through the use of the application directory partition. You can, for example, replicate Active Directory-integrated zones to only those domain controllers that have the DNS service installed. This kind of configuration, however, is only possible if all the domain controllers are running Windows Server 2003. Windows 2000 domain controllers store Active Directory-integrated zones in the domain partition and, consequently replicate this information to all domain controllers.

Security is an important consideration when planning for zone replication. Internal DNS data should be confidential. Active Directory-integrated zones provide the highest level of security because AD replication traffic is encrypted. To mitigate the risk of footprinting or data dumping, security is enhanced on DNS zones by specifying the IP addresses of computers that are authorized to initiate a zone transfer with the primary DNS server. VPN tunnels or IPSec should be used if zone transfers of internal DNS data take place over the Internet.

To increase fault tolerance of the forwarding servers, multiple forwarders can be specified in an ordered list of servers that the forwarding server will contact to resolve names. Recursion can be enabled or disabled on a per-domain basis if the queries to the forwarders fail. When recursion is disabled, the forwarding server will send a negative response to the client if none of the forwarders can resolve the query. Otherwise, the forwarding server will use recursion to attempt to resolve the name in the event of failure.

The Windows Server 2003 DNS is a standards-based implementation of DNS and will interoperate with other standards-based versions of DNS, depending on their respective capabilities. For example, BIND 9 and BIND 8, as well as Windows NT 4.0 with Service Pack 4 installed, can support SRV records and can be used for zones containing AD records.

Windows XP, Windows 2000, and Windows Server 2003 clients support dynamic updates of zone RRs, allowing them to create forward and reverse lookup records in zones configured to allow dynamic updates. The Windows 2000 and Windows Server 2003 DHCP server can be configured to update forward and reverse lookup records for DHCP clients in zones configured to accept dynamic updates, including downlevel clients that do not support dynamic updates. When the DHCP lease expires, the DHCP server can be configured to attempt to remove both the forward and reverse lookup records that have been created in their respective zones.

The ability to perform dynamic updates requires that the DHCP client service is running on the client computer, regardless of whether or not the computer is configured as a DHCP client. For downlevel clients, such as Windows NT 4 servers that are configured

with static addresses, the DHCP client service will not perform a dynamic update to the zone file. In these cases, it is necessary to either manually update the DNS RRs or configure the DNS server to do a lookup on a WINS server to resolve the computer name-to-IP address mapping.

A DNS server will perform a WINS lookup for names it cannot resolve authoritatively if a WINS RR is present. The WINS RR provides the DNS server with information on the IP addresses of the WINS servers and the length of time to cache the resolved query. Additionally, it is possible to specify that the WINS record not be replicated during a zone transfer in the event that non-Microsoft DNS servers are being used as secondaries. WINS records can be created for both forward and reverse lookup zones. In the case of a reverse lookup zone, the DNS server performs an `nbstat` command to resolve the IP address to a NetBIOS computer name to which it appends a configured domain name.

It is possible to create subdomains to host only WINS forward and reverse lookup records. However, to ensure proper name resolution, it may be necessary to change the client domain suffix search list to ensure name resolution for these domains.

Implementing a DNS infrastructure exposes the network to additional threats that are specific to DNS. These threats include footprinting, DoS attacks, and redirection. To ensure the confidentiality, integrity, and availability of DNS data, DNS administrators should follow recommended best practices. To assist in this, Microsoft has created three different categories of DNS security configurations: low, medium, and high security.

Another best practice for securing a DNS infrastructure includes ensuring that there is no unacceptable single point of failure in the DNS infrastructure, for example, by distributing secondary servers and Active Directory-integrated zones on different subnets served by different routers. Implementing a private root zone and choosing a disjointed namespace based on a registered domain name will also enhance the security of DNS data.

Windows Server 2003 and the DNS service provide a number of tools to monitor the operation, availability, and performance of DNS services in your infrastructure. These tools include the Monitoring tab on the DNS MMC console, DNS debug logging, DNS event logging, and DNS Performance Monitor counters. You can also use command-line tools such as `NSlookup.exe`, `Dnscmd.exe`, and `DNSLint.exe`.

NetBIOS methods of name resolution are invoked whenever a NetBIOS application needs to resolve computer name-to-IP address mappings. Unlike a host name that can be assigned to computer by means of a hosts file or DNS, a NetBIOS name is configured on the computer. A computer can be configured with only one NetBIOS name. Furthermore, NetBIOS computer names must be unique on the network. NetBIOS group names are nonexclusive and are used to identify groups of computers, such as workgroups or domains. They are primarily used to support browsing.

Because NetBIOS names are broadcast-based, special solutions were developed to support NetBIOS name resolution using routable protocols, such as TCP/IP, which uses logical addresses (IP addresses) to resolve the MAC addresses of network devices. NetBIOS is implemented as an interface, equivalent to the WinSock interface, called NetBT. Thus, NetBIOS name requests on a TCP/IP network are processed through the NetBT interface.

However, name resolution is still broadcast-based. Because these broadcasts normally do not cross routers, clients must use either an LMHOSTS files or a WINS server to resolve NetBIOS mappings that are on remote subnets.

The LMHOSTS file is very much like the hosts file that can be implemented for WinSock applications. Like the hosts file, the LMHOSTS file contains a static mapping of IP addresses to names. In order to implement a name resolution strategy using LMHOSTS files, these files must be installed and maintained on every computer that requires the ability to resolve names on remote subnets, creating a significant administrative burden if there are more than a handful of computers, as well as posing significant potential for human error.

To address the limitations of NetBIOS name resolution and the use of LMHOSTS files, WINS servers can be deployed. A WINS server provides a NetBIOS name service for the registration and resolution of NetBIOS names in a distributed database.

Communication between WINS servers and NetBIOS clients is unicast-based, rather than broadcast-based. This means that NetBIOS applications and services can register names and perform NetBIOS name queries with a WINS server that is located on a remote subnet.

How NetBIOS name resolution occurs depends on the specific node configuration of the NetBIOS client. There are four kinds of NetBIOS nodes. A b-node (broadcast node) client is exclusively broadcast-based and will not contact a WINS server using a unicast message. A p-node (peer node) client is exclusively unicast-based and will not use broadcasts to perform name registration or name queries. An h-node (hybrid node) will first try to contact a WINS server using unicast messages, before falling back to using broadcasts. An m-node (mixed node) client is like an h-node client except it will use broadcasts before falling back to unicast communication with a WINS server.

An h-node configuration is the default configuration for Windows clients. The order of name resolution for an h-node client is as follows: first, the computer will check its own computer name; second, the computer will check its NetBIOS remote name cache; third, the computer will contact a WINS server; fourth, the computer will send three broadcasts on the local subnet; fifth, the computer will check its LMHOSTS file, if so configured; and sixth, the computer will try host-based methods of name resolution (hosts file and DNS).

Although WINS servers are used for dynamic registrations, they can also store static records that are either manually created or imported from an existing LMHOSTS file. It is not recommended, except in some special circumstances, that static mappings be created on a WINS server. Static mappings never expire, and in cases where these records are replicated to other WINS servers, special care must be taken to ensure that they are deleted on all WINS servers by manually tombstoning them when they are deleted. More importantly, if a WINS client with a static mapping is later upgraded to register with WINS, the registration will fail unless the migrate on setting is enabled on the WINS server (it is disabled by default).

Even a modestly powered computer can provide an adequate level of name services for a large number of clients. Microsoft recommends approximately 10,000 clients per WINS server and goes out of its way to warn of the dangers of installing too many WINS servers on the network. Too many WINS servers can complicate and compound network problems

if things go awry. WINS is a relatively CPU- and disk-intensive service. So, if a WINS server needs to handle registrations for a large number of clients, performance can be improved by adding another processor and placing the database on a separate, dedicated hard drive.

The WINS service is cluster-aware and can be set up in a Windows cluster to provide a high degree of availability. In a Windows cluster, the WINS service can be set up on two or more computers, called cluster nodes, that share an external SCSI or Fibre Channel device where the WINS database is located. If the cluster node where an instance of the WINS service is running fails, the WINS service will start up on another node, providing continuity of service. The primary advantage of using a Windows cluster is that in the event of the failure of a WINS server, no replication is necessary to synchronize the database, since only a single database is used.

Using a Windows cluster may not be a desirable solution to provide the availability of the WINS service that you desire. To provide fault tolerance, high availability, and load balancing, you should set up multiple WINS servers as replication partners, keeping in mind the Microsoft recommendation to be as conservative as possible in estimating the number of WINS servers you require.

Replication between WINS servers is configured by means of replication partnerships. There are three kinds of replication partnerships: push-only, pull-only (both known as limited), and push/pull (also known as full), replication partnerships. WINS replication partnerships should be set up on reciprocal basis. For example, WINS-A is configured with WINS-B as its push partner; WINS-B should in turn have WINS-A configured as its pull partner. Ideally, however, both WINS-A and WINS-B will be both push and pull partners of each other. Microsoft recommends using the default push/pull partnerships and avoid using limited partnerships.

The primary difference between push and pull replication is that pull replication can be scheduled to occur at specific intervals or times. (It can also occur when the WINS service starts or when initiated manually by an administrator.) Configuring pull replication is useful in situations where replication needs to occur at off-peak hours over slow WAN links.

The amount of time it takes a replicated record to be replicated throughout the WINS infrastructure is referred to as convergence time. In designing the WINS infrastructure, you need to take into account the longest convergence that is acceptable and balance this with the bandwidth requirements of your network. The WINS replication topology you implement will have an effect on convergence times. If replicated WINS records need to travel through less than optimal paths, convergence times will be longer.

When there are multiple WINS servers in the infrastructure, the most efficient replication topology is the hub-and-spoke model. With this topology, a central hub WINS replicates with other spoke WINS servers using reciprocally configured push/pull partnership agreements. The longest path that any record needs to travel to reach the most distant WINS server is two: one hop to the hub server and a final hop to any spoke server.

Other topologies are possible, such as a ring topology, or hybrid topologies that mix ring and hop-and-spoke topologies. Also, various topologies could use a mix of limited and

full partnerships. The optimal topology for the network will be determined by the contingencies of the network infrastructure.

To avoid problems with replication, you should follow best practices, which, in general, advocate simplicity. You should avoid placing WINS servers on multihomed computers. Unless special care is taken in the configuration of the multihomed computer and its replication partners, replication will fail. WINS servers that have the multihomed computer as a replication partner need to configure multiple replication partnerships that point to all the IP addresses on the multihomed computer. It is also possible that each network interface of the multihomed computer could be configured to point to a WINS server other than itself. This can cause a problem known as split registration.

Split registration occurs when a part of a computer's name mappings are created in more than one WINS server so that the mappings are owned by different servers. If the split registration occurs for the name mappings of the WINS server itself, the likely result will be hard-to-find problems with name resolution. To avoid problems with split registration, you should ensure that the WINS server is configured to register with only itself as the WINS server. Do not specify a different secondary WINS server to register with.

The use of static records can cause problems in a replicated environment. The actual removal of either a static or a dynamic record from a WINS database occurs through a process known as scavenging. The renewal interval determines the TTL of a WINS registration. The default is six days. If a client does not renew the registration, it is marked as released. The extinction interval is the amount of time the must elapse before a released record is marked as tombstoned. The default is four days. The extinction time is the amount of time that must elapse from the time the record is marked as tombstoned and when it is deleted from the database. The default is six days. The verification interval depends on the previous values and is the amount of time that must elapse before a WINS server will validate active records it does not own. The default is 24 days. Changes to these settings will have an effect on the overall performance of the WINS servers in the environment. Other factors that can affect performance of a WINS server include the compaction of the database and when constancy verification is performed.

Another setting that can affect the performance of WINS servers is burst handling. Burst handling, which is enabled by default, gives the WINS server the ability to handle a very large number of registrations simultaneously. For example, if the power to a large number of computers failed and then was restored to all those computers at the same time, this would cause a huge volume of registration requests that could potentially overwhelm the WINS server. When burst handling is enabled and triggered by a surge in registration queries, the WINS server will send a positive response to the client with a very short TTL, forcing the client to renew the registration soon after the surge has abated.

Enabling burst handling has implications for WINS security. Because burst handling events are recorded in Event Viewer, they can provide an indication of a possible DoS attack launched against the WINS server. The only effective way to protect WINS against bogus registration requests is to secure the physical network.

Because NetBIOS is an unauthenticated protocol, it is possible for a rogue server to hijack a dynamic registration that belongs to another computer. To mitigate this risk, you should identify mission-critical servers and assign them static mappings that cannot be overwritten by dynamic registrations (you must ensure that the Migrate On setting is disabled).

To ensure the confidentiality of WINS data, you should physically secure the WINS servers. Also, if you move the WINS database from its default location, it will lose the security it inherits from the parent System32 folder. Thus, if you move the database, you should ensure that the ACL on the WINS folder is configured appropriately. In high-security environments or in situations where WINS data is replicated over a public network, you should encrypt replication traffic using VPNs or IPsec.

Although replication can help to ensure the integrity and availability of WINS data, it should not be considered sufficient for these purposes. It is therefore important that WINS databases be backed up on a frequent basis, either automatically or manually using the WINS backup utility.

In general, your troubleshooting strategy should be based on the OSI model, and you should troubleshoot from the bottom of the OSI model to the top to assist in isolating the precise nature of the problem. Many apparent name resolution problems are actually caused by problems with network communication. Utilities such as PING and Netdiag can help you eliminate network communications as the source of the problem.

Often, name resolution problems are caused by client configurations. Your first step in troubleshooting name resolution problems should be to verify the TCP/IP settings on the client using a utility such as Ipconfig or Netdiag. Your next step at the client computer should be to examine and clear the ARP, NetBIOS, and DNS cache using the arp, nbtstat, and ipconfig commands as appropriate. After you have verified the settings and cleared the cache, you should test network connectivity by first pinging the IP address of the default gateway, and then pinging the IP address of a remote host on your network, usually the WINS included in the DNS servers but could include the IP addresses of Internet hosts. Your next step should be to isolate whether the problem is with NetBIOS or host name resolution, if this is not already obvious.

Once you have determined that the problem is related to the DNS or WINS servers, you can take a troubleshooting strategy appropriate to each service, which should probably start with an examination of events in Event Viewer that would suggest possible causes such as failed zone transfers. Also, you should consider whether the problem is server-wide or related to a subset of name mapping records. If the problem is related to incorrect DNS records, you should consider whether the problem is related to cached or authoritative records. If there is a problem resolving names in an authoritative child domain, a likely cause is lame delegation. If the problem is related to resolution of unqualified names to a record in a parent domain, you should double-check the domain suffix search list on the client. Otherwise, you check to ensure that the DNS server in the child domain can contact a server in the parent domain to resolve names.

If the problem is related to dynamic updates of zones, you should review the client and DHCP server settings to ensure that update requests are sent to the appropriate primary

server for the domain. If you have enabled secure-only updates on an Active Directory-integrated zone, the ACL on the RRs may be preventing updates from occurring.

For problems related to NetBIOS records, the most likely causes are related to replication, such as slow convergence times or failed replication owing to network problems or misconfiguration. If a WINS replication partner is a multihomed server, multiple replication partnership agreements need to be configured to specify all the IP addresses of the multihomed server, rather than its NetBIOS names.

If the problem is related to the reappearance of a deleted static mapping, you should perform a manual tombstone deletion of the record on the WINS server that owns it. If you can't overwrite a static mapping with a dynamic mapping, you need to ensure that the migrate on setting is enabled.

Name resolution is a fundamental component of Windows Server 2003 networking. In this chapter, we have examined the many factors involved in developing the best name resolution strategy for your organization's network, not only to help you master the exam objectives, but to make it easier for you to perform the everyday duties of a network administrator.

Exam Objectives Fast Track

Planning for Host Name Resolution

- ☑ The design of your DNS namespace will have an effect on the security of your DNS infrastructure and the amount of effort required to administer it. At a minimum, the internal DNS namespace should either be registered or based on a registered name you own.
- ☑ The internal DNS namespace mirrors the AD domain tree. However, DNS and AD are separate from one another.
- ☑ The number of child domains or subdomains should be limited to five or fewer.
- ☑ Secondary zones can increase fault tolerance and availability, but zone transfer traffic can consume unacceptable amounts of bandwidth in some circumstances.
- ☑ Lame delegations are one of the most common sources of name resolution problems with a DNS infrastructure. As an alternative to using NS and glue address records to delegate authority, consider using stub zones or conditional forwarding.
- ☑ Conditional forwarding can reduce the amount of DNS referral traffic on the network.
- ☑ Conditional forwarding is a good alternative to using secondary or stub zones in many circumstances.

- ☑ DNS servers used for internal name resolution should never be accessible to Internet clients.
- ☑ Public DNS servers that are used to resolve name mappings for your Web and mail servers should not be able to perform recursion.
- ☑ Primary DNS servers should be configured to replicate only with a configured list of IP address or servers listed on the Name Servers tab.
- ☑ Cache pollution protection should be enabled on all DNS servers to protect against attacks.
- ☑ Publicly available DNS servers should be placed behind firewalls that have access rules controlling acceptable source and destination ports and addresses.
- ☑ Active Directory-integrated zones configured to accept authenticated updates only provide the highest level of security for dynamic updates.

Planning for NetBIOS Name Resolution

- ☑ WINS servers are capable of handling large numbers of client registrations; Microsoft recommends that as few WINS servers as possible be deployed to provide a desired level of service.
- ☑ To avoid problems with replication and name resolution, WINS servers should not be installed on multihomed computers.
- ☑ The TCP/IP stack on a WINS server should be configured so that the WINS server registers with itself.
- ☑ By default, WINS replication partnerships are set up as push/pull replication partnership. Limited partnerships (push-only and pull-only) are possible but should be avoided unless there is an overriding need to use them, such as extremely limited bandwidth.
- ☑ Push replication is triggered by a configurable number of updates in the WINS database. Push replication is used in situations where there is ample bandwidth, such as on a LAN or high-speed WAN.
- ☑ Pull replication is triggered by a configurable schedule. In general, pull replication is used in low-bandwidth situations where it is desirable to control the timing of replication traffic.
- ☑ Convergence time is the amount of time it takes an updated record to propagate to every WINS server.
- ☑ A hub-and-spoke topology is the most efficient for a replication environment involving multiple WINS servers.

- ☑ Enabling burst handling can alert administrators to the presence of possible DoS attacks because the events appear in Event Viewer.
- ☑ Static mappings should be avoided, unless they are used as a means to prevent redirection of name mappings of mission-critical servers.

Troubleshooting Name Resolution Issues

- ☑ Troubleshooting name resolution issues is more effective if a systematic approach is used to isolate the components and processes that may be causing the problem. Generally, this means troubleshooting from the bottom of the OSI model to the top.
- ☑ Client configurations are the most common source of name resolution issues and should be verified first.
- ☑ Before troubleshooting name resolution problems on the client, it is a good idea to clear the appropriate cache (DNS or NetBIOS) to eliminate that as the source of the problem.
- ☑ After the name resolution problem has been tracked down to the specific service—WINS or DNS—troubleshooting strategies appropriate to each can be employed.
- ☑ Troubleshooting tools for DNS include Ipconfig, Netdiag, NSLookup, Dnscmd, and DNSLint.
- ☑ Troubleshooting tools for WINS include Ipconfig, Netdiag, and the nbstat command.

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 new features of Windows Server 2003 DNS are most likely to be included in the exam?

A: Windows Server 2003 introduces a number of new features for DNS. These include stub zones, conditional forwarding, application directory partitions, support for EDNS0 to allow larger UDP packet sizes, and limited support for DNSSEC. Of these, the most widely used and implemented features will be stub zones, conditional forwarding, and application directory partitions. The use of these features can play a key role in optimizing a DNS infrastructure and making it more fault tolerant. For example, the use of stub zones can help to eliminate problems with lame delegations and can help to reduce DNS referral traffic. Conditional forwarding can be used to help reduce cross-referral traffic and can be used as an alternative to stub zones.

Q: What is the DnsUpdateProxy group and why is it important?

A: The DnsUpdateProxy group allows DHCP servers that are members of this group to register mappings in forward and reverse lookup zones so that the resulting ACLs on the resource records have no security configured on them. This makes it possible for other DHCP servers and clients to subsequently overwrite the resource records in Active Directory-integrated zones that are configured to accept authenticated updates only. Because this also creates a security hole, especially if the DHCP service is located on a domain controller, it is now possible in Windows Server 2003 to configure the DHCP server to create resource records using a user account created for this purpose.

Q: What is the difference between a DNS domain and a zone?

A: A zone contains the actual resource records that are used to provide name resolution for a particular DNS domain. A single zone file can contain records for both the parent and its subdomains. The DNS servers on which this zone file is located are said to be authoritative for the parent and the subdomains. If the zone file is too large or if there is some other need, authority for the subdomains can be delegated to other DNS servers. When this occurs, the DNS server that is authoritative for the parent domain is no longer authoritative for the subdomains, since its zone file no longer contains resource records for those zones.

Q: What is the difference between an authoritative and nonauthoritative response?

A: An authoritative response is one that comes directly from a DNS server that is authoritative for domain. A nonauthoritative response is one that comes from the cache of a DNS server that is not authoritative for the domain and has previously resolved the query.

Q: What is the difference between a recursive query and an iterative query?

A: A recursive query is performed by a DNS host to a DNS server asking the DNS server to find the answer or return an error. The DNS server assumes the responsibility for resolving the name. Usually DNS clients and DNS forwarders send recursive queries to DNS servers configured to perform recursion, that is, to send iterative queries to other DNS servers. Iterative queries are usually used for communication between DNS servers. If the DNS server cannot resolve a name mapping, it will contact other DNS servers in the DNS hierarchy and accept referral answers to find DNS servers that are authoritative for the domain where the host is located. For example, if the DNS server is trying to resolve `www.syngress.com` to an IP address, it might contact a DNS server that is authoritative for the `.com` namespace and accept a referral answer that provides the IP addresses of DNS servers that are authoritative for the `foobar.com` domain namespace. This process is referred to as *walking the tree*.

Q: What is the version ID, and why is it so important in WINS replication?

A: All name mapping records in a WINS database have a field that contains a version ID. Every time a registration is created or updated, the record is given an incremented version ID. The version ID is sequentially incremented across the entire set of records in the database. When WINS servers are set up as replication partners to each other, they store the highest version ID of the partner server in an owner table. When push or pull replication is triggered, the WINS servers compare the values of the version IDs in their owner table with values sent to them by their replication partner. By comparing version IDs, WINS servers can determine what incremental changes have occurred since the last replication cycle.

Q: What is tombstoning?

A: When a record is deleted in a local WINS database, the deletion of the record should be propagated to other WINS servers. The problem is this: How do you replicate a deletion? Well, you can't. But, you can mark the record in a special way that will let other WINS servers know that they should delete their own copy of the record. Tombstoning is the process of "marking" the record so that it will replicate and persist long enough to let other WINS servers know they should delete their copy of the record. When you manually delete a record in a WINS database, you will be prompted as to whether you want to perform a simple deletion or a tombstone deletion. When you choose a simple

deletion, only the local copy is deleted, and the record can be replicated back to the server and reappear. A tombstone deletion ensures the record stays deleted.

Q: What is the difference between push and pull replication?

A: Push and pull replication differ only in how they are initiated. A WINS server initiates push replication by informing its push replication partner that it has records it would like to replicate. The WINS server sends only a notification containing its owner table (a table listing owner IDs and version IDs) to its push partner. The push partner subsequently initiates a pull replication with the original server (its pull partner). Push replication is triggered when a configurable threshold of updates is reached. For example, a WINS server could be configured to send a notification after it had received 50 updates. Pull replication is initiated according to a configurable schedule on the WINS server.

Q: What is the most efficient replication topology for WINS server replication?

A: In most cases, the most efficient replication topology for replicating WINS records is a hub-and-spoke topology. However, this will vary depending on the circumstances. Most large, complex networks may find that the best topology is one that involves a mix of different topologies (such as ring and hub-and-spoke) and replication partnerships.

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 for Host Name Resolution

1. You are the administrator of a Windows Server 2003 network. Recently, your company made a sudden and unexpected announcement that it would be merging with another company called Syngress Industries, a large company that has more than 20,000 employees. You learn that, in the short term, communications between the two companies will need to take place over persistent VPNs using each company's respective connections to the Internet, both of which are operating at about 75 percent capacity. You will need to set up trust relationships between two AD forests. Furthermore, you plan to move significant amounts of data between the two networks. You learn the Syngress Industries uses a child domain of its Internet domain namespace for its AD forest root. The name of the internal domain is `ad.syngress.com`.

You want to ensure that your DNS infrastructure can resolve names for internal hosts of Syngress Industries. You also want to ensure that your solution is the most effective in terms of resource usage. What should you do to enable name resolution for internal hosts of Syngress Industries?

- A. Create a secondary zone for ad.syngress.com on your DNS servers.
 - B. Create a stub zone for syngress.com on your DNS servers.
 - C. Create an Active Directory-integrated zone for ad.syngress.com
 - D. Create a conditional forwarding configuration on your DNS servers for ad.syngress.com
2. You are the administrator of a Windows Server 2003 network. Your boss has just read an article on how DNS servers can be compromised so that they will redirect recursive queries to bogus Web sites that can cause potential harm. Your boss has asked you to ensure that the DNS servers in the DMZ have the highest level of protection possible against this and other types of common attacks on DNS servers. You have two DNS servers. DNS-A is used to resolve name mappings for your public Web and mail server. The other DNS server, DNS-B, is used by the internal proxy server to resolve Web site addresses to IP addresses. What actions should you take to carry out your boss's order to provide the highest possible security against common multiple DNS attacks? (Select the best answer.)
- A. Enable protection against cache pollution on DNS-B and disable recursion on DNS-A
 - B. Enable protection against cache pollution on DNS-A and disable recursion on DNS-B
 - C. Disable recursion on DNS-A and configure the firewall to not allow any inbound traffic with destination ports of TCP or UDP port 53 to reach DNS-B
 - D. Disable recursion on DNS-B and configure the firewall to not allow any inbound traffic with destination ports of TCP or UDP port 25 to reach DNS-A

3. You are the administrator of a Windows network that consists of a mixture of Windows NT 4, Windows 2000, and Windows Server 2003 servers, providing a mix of file, print, messaging, and other services critical to your network. You are currently running WINS, DNS, and DHCP services on your network. You have already enabled dynamic DNS on your forward and reverse lookup zones, but you want to ensure that all of your client computers can find the name-to-address mapping of all your servers using DNS. You want to minimize the administrative effort for this project. What action should you take? (Select the best answer.)
 - A. Place the DHCP servers in the DnsUpdateProxy group.
 - B. Enable DHCP to update forward and reverse lookup zones on behalf of all DHCP clients.
 - C. Manually enter the records for servers that have static addresses.
 - D. Create a WINS resource record in the forward and reverse lookup zones.

4. You are using ISA Server 2000 as a firewall and Web proxy server to protect your internal AD network and provide Web proxy and caching services for HTTP requests. You currently are using three DNS servers to support the DNS queries. DNS-A is used for your internal AD root. DNS-B is used to provide name resolution for Internet clients that want to connect to your public Web and mail servers. DNS-C is used to provide Internet name resolution. How should you configure the DNS and ISA Server access rules to provide the maximum security and functionality for your DNS infrastructure?
 - A. On DNS-A, remove the root hints file and enable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, remove the root hints file and disable recursion. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, enable recursion and update the root hints file. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 with a source port of ANY.
 - B. On DNS-A, remove the root hints file and disable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, remove the root hints file and disable recursion. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, enable recursion and update the root hints file. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 with a source port of ANY.
 - C. On DNS-A, remove the root hints file and enable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, remove the root hints file and disable recursion. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, enable recursion and update the root hints file. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 with a source port of ANY.

- D. On DNS-A, remove the root hints file and disable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, update the root hints file and enable recursion. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, disable recursion and update the root hints file. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 with a source port of ANY.
5. You are the administrator of a Windows Server 2003 network. Your company has recently merged with another company and you have set up trusts between the AD forests and have set up conditional forwarding on your DNS servers to resolve names in the AD forest of the newly merged company. You would like your users to be able to resolve names in the newly merged company with the least possible effort and typing on their part. You would like to implement a solution with the least possible effort on your part. What should you do?
- A. Using ADSI, create an msDS-AllowedDNSSuffixes attribute in the domain object container and include the domain suffix of the newly merged AD forest in the list of allowable suffixes.
 - B. Create a group policy that configures the DNS clients with a custom DNS suffix search list.
 - C. Configure the DHCP server option 81 to supply the name of the domain suffix of the newly merged AD forest to DHCP clients.
 - D. Configure a stub zone for a root domain of the newly merged company on your DNS servers.

6. You are a DNS administrator of a large, distributed Windows Server 2003 network. The AD domain tree consists of a number of child domains that reflect the geographic locations of the different offices of the company. You are responsible for the DNS root domain of the AD forest and the child domain of the office where you work. All administrative responsibility for the remaining child domains is performed by locally based administrators in their respective offices. The capacity of the WAN links connecting the various offices is showing signs of being insufficient. You want to ensure that DNS resolution for the child domains outside your administrative control will work company-wide in a fault-tolerant manner without adding additional strain to available resources. What should you do? (Select the best answer.)
- A. On the root DNS servers, configure conditional forwarding for the child domains.
 - B. On the DNS servers in the child domain under your control, configure secondary zones for the other child domains.
 - C. On the root DNS servers, configure stub zones for the child domains.
 - D. On the DNS servers in the child domain under your control, configure secondary zones for the other child domains.
7. You are the enterprise administrator of a Windows network that comprises a number of Windows 2000 and Windows 2003 domain controllers. You want to use Active Directory-integrated zones for your zone data to enhance security and optimize replication of zone data. What should you choose as the replication scope? (Select the best answer.)
- A. To all DNS servers in the forest
 - B. To all domain controllers in the AD domain
 - C. To all DNS servers in the AD domain
 - D. To all domain controllers specified in the scope of an application partition

Planning for NetBIOS Name Resolution

8. You are an administrator of a Windows Server 2003 network. You want to automate the backups of the WINS database. You want this backup to occur at least once every 24 hours. What should you do? (Select the best answer.)
- A. Configure the Windows Backup utility to back up the contents of the `%system-root%\System32\Wins` folder once every 24 hours.
 - B. Using the AT command scheduler, create a batch file that temporarily stops the WINS service, copies the WINS database to another location, and then restarts the service.

- C. Use a third-party backup solution that is capable of backing up open files and configure it to back up the contents of the `%systemroot%\System32\Wins` folder once every 24 hours.
 - D. In the WINS server console, configure a path to store backups of the database and initiate a manual backup.
9. You are the administrator of a Windows Server 2003 network. You are responsible for a number of WINS servers that are set up as push/pull replication partners to each other. You have a number of static mappings in your WINS database and want to remove one of these mappings from the WINS database. You want to ensure that the record is deleted on all servers with the least administrative effort. How should you delete the WINS static mapping? (Select the best answer.)
- A. On the owner server of the mapping, find the record and perform a simple deletion.
 - B. On the owner server of the mapping, find the record and perform a tombstone deletion.
 - C. On all of the WINS servers, find the record and perform a simple deletion.
 - D. On all of the WINS servers, find the record and perform a tombstone deletion.
10. You are the administrator of a Windows Server 2003 network. You have five WINS servers and need to reconfigure the replication topology as a result of some recent upgrades to your WAN links. All of your WAN links connecting the head office and your four branch offices now have ample bandwidth to handle additional traffic. You want to ensure the shortest convergence time of replicated records, while at the same time keep the number of replication partnership agreements to an absolute minimum. What replication topology should you choose? (Select the best answer.)
- A. Ring topology
 - B. Mesh topology
 - C. Hub-and-spoke topology
 - D. Hybrid of ring and hub-and-spoke topology

Troubleshooting Name Resolution Issues

11. You are an administrator of a Windows Server 2003 network. Your company, Syngress Industries, manages its own DNS for its public Web and mail servers. The primary DNS server for the syngress.com domain is located in a DMZ protected by ISA Server. Your ISP is hosting secondary servers for the syngress.com domain on its BIND 9 servers. While going through your performance logs, you notice a brief but sudden increase in the number of AXFR requests received and AXFR success sent events. Previously, these counters had values of zero in your logs. You suspect your ISP has changed the configuration of its BIND servers, but the ISP denies it and insists that the secondary zones are behaving optimally. You are concerned by these values and decide to investigate the issue and correct it, if necessary. What is the likely cause of the problem and what should you do? (Select the best answer.)
- A. A rogue DNS server is attempting to pollute the cache on your DNS server by sending bogus queries over TCP, rather than UDP. You should turn on debug logging to determine the source IP address and block all traffic from this address on ISA Server. You should also enable protection against cache pollution and inform the ISP.
 - B. A malicious user is issuing an **nslookup -ls** or equivalent command against your DNS server. You should configure the DNS server to allow zone transfers only to the IP addresses of the secondary servers at the ISP. You should also block all external requests destined for the primary DNS server on TCP port 53 with a source port of ANY, except for the external addresses of the secondary servers. You should inform the ISP managers and ask them to confirm an equivalent level of security on their servers.
 - C. A malicious user is attempting to launch a DoS attack on your DNS. You should disable recursion on the DNS server. You should also turn on debug logging to determine the source IP address of the attack and block the IP address at ISA Server. You should inform the ISP to be on the lookout for similar attacks against its DNS servers.
 - D. A malicious user is issuing an **nslookup -ds** or equivalent command against your DNS server to get detailed information. You should turn on debug logging to determine the source IP address. Once you determine the IP address, you should block it from all communication with your DNS servers at ISA Server. You should inform the ISP managers and ask them to confirm an equivalent level of security on their servers.
12. You are the administrator of a Windows Server 2003 network. Recently, a junior administrator has, on your instructions, rebuilt one of your WINS servers (WINS-A). You don't have a backup of the WINS database and need to restore the database

through reregistrations of WINS clients and replication with another WINS server, WINS-B. Both servers are configured as push/pull replication partners of each other. As soon as WINS-A is brought back online, users configured to use WINS-A as their WINS server immediately start to complain that they can't access file server shares on this server. By the time you hear about the complaints and try to reproduce the results, you find that that the problem has disappeared. However, you take the complaints seriously and investigate further. You examine the WINS database on WINS-B and see some data that strikes you as odd. Based on the data shown in the table here, what problem is indicated? (Select the best answer.)

Record Name	Type	IP Address	Owner	Version
WINS-A	[00h] Workstation	192.168.100.20	192.168.179.5	20D
WINS-A [20h]	File Server	192.168.100.20	192.168.179.5	20C

- A. There is a problem with the order of service registration. The workstation service needs to be registered before the file server service.
 - B. There is a problem with WINS replication that has caused the wrong owner to be associated with WINS-A.
 - C. The TCP/IP stack on WINS-A is configured with the IP address of WINS-B as its secondary WINS server.
 - D. The TCP/IP stack on WINS-B is not configured to register itself with a WINS server.
13. You are the administrator of a Windows Server 2003 network using DNS and WINS to provide name resolution services. You have two WINS servers that are set up with the default push/pull configurations. Users have been complaining for days about problems connecting to a server called File_Server2. You ping File_Server2 and get a response from the computer. However, when you issue a **net view \\File_Server2** command, you get an error message stating that a duplicate name exists on the network. What is the likely cause of the problem? (Select the best answer.)
- A. The underscore character cannot be used in a NetBIOS name. Rename the computer and reboot it.
 - B. There is a problem with the replication of the records for File_Server2. Manually initiate replication with the WINS server that is the owner of the record of File_Server2.
 - C. The WINS database is corrupt. Manually initiate consistency checking to restore database integrity.
 - D. The WINS server contains an incorrect name mapping for File_Server2.

14. You are the administrator of a WINS server. The WINS server has suffered a hardware failure, and you have subsequently been forced to reinstall Windows Server 2003 and the WINS service. Fortunately, you have a recent backup of the WINS database. You restore the database, but notice that none of the former WINS configuration settings are present. What should you do? (Select the best answer.)
- A. You need to use the `%systemroot%\system32\jetpack.exe` file to restore the WINS configuration after you restore the database.
 - B. You need to restore the original System State from the backup to the Windows Server 2003 server.
 - C. You need to invoke database consistency checking on the database.
 - D. You need to set up replication with a WINS server that was a replication of the former WINS server.
15. You are the administrator of a Windows Server 2003 network. After restoring the Windows Server 2003 domain controller that you had taken off the network for a few hours for maintenance, your Windows 95 and 98 users have begun complaining that they are unable to access resources on this computer. You remember seeing a message about a duplicate name on the network when you turned on the domain controller, but didn't think much of it at the time because you had changed the IP address of the domain controller before you took it offline. What action should you take?
- A. Create static mappings in the WINS database for the domain controller and disable the migrate on setting.
 - B. Create static mappings in the WINS database for the domain controller and enable the migrate on setting.
 - C. Have the users of Windows 95 and 98 computers issue an **nbtstat -RR** command.
 - D. Have the users of the Windows 95 and 98 computers issue an **ipconfig /flushdns** 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. D | 9. B |
| 2. C | 10. C |
| 3. D | 11. B |
| 4. A | 12. C |
| 5. B | 13. D |
| 6. C | 14. B |
| 7. B | 15. A |
| 8. D | |

MCSE 70-293

Planning, Implementing, and Maintaining a Remote Access Strategy

Exam Objectives in this Chapter:

- 3 Planning, Implementing, and Maintaining Routing and Remote Access
 - 3.3 Implement secure access between private networks.
 - 3.2 Plan security for remote access users.
 - 3.2.1 Plan remote access policies.
 - 3.2.2 Analyze protocol security requirements.
 - 3.2.3 Plan authentication methods for remote access clients.
 - 5.4.1 Create a plan to offer Remote Assistance to client computers.
 - 5.4.2 Plan for remote administration by using Terminal Services.
- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

In today's business world, users need access to the company network not only when they're on company premises, but from home and when they're on the road, as well. An important part of the network administrator's job is to design and implement a strategy for allowing authorized users to connect to network resources without compromising security. Windows Server 2003 includes technologies and features that make this easier.

This chapter examines the issues and procedures involved in devising a remote access strategy, including planning tasks such as analyzing organizational needs, analyzing user needs, and selecting the remote access types that will be allowed (dial-in, VPN, and/or wireless). We'll discuss design considerations related to dial-in access, such as the allocation of IP addresses, how to determine incoming port needs, and how to select the best administrative model based on your organizational needs and the functional level of your domain.

Next, we'll talk about design considerations related to VPN access. You'll learn how to select the VPN protocols to be allowed, based on client support, PKI requirements, and the need for data integrity and sender authentication. You'll learn how to install machine certificates, how to configure firewall filters, and how to create access policies governing VPN connections.

In the next section, you'll learn about the design considerations that relate to wireless remote access. We'll discuss the use of IAS for wireless connections, and how to configure remote access policies for wireless connections. We'll address the use of multiple WAPs and the advantages of placing a certification authority on a Virtual LAN (VLAN) for new wireless clients. We'll also show you how to configure WAPs as RADIUS clients.

Finally, we move on to planning overall security strategies for remote access connections. We'll discuss the best practices in selecting authentication methods that will be allowed, and the benefits of disallowing insecure password-based connections such as PAP, SPAP, CHAP, and MS-CHAPv1. We'll then look at the more secure methods such as MS-CHAPv2 and EAP, and discuss the advantages of using RADIUS/IAS rather than Windows authentication. We'll also address the selection of the data-encryption level and other security measures such as requiring callback, mandating operating system and file system choices, using managed connections, and using smart cards for remote access. We'll delve deeply into the subject of remote access Policies and show you how to authorize remote access by user or group, how to restrict remote access in various ways, and how to control remote connections.

Planning the Remote Access Strategy

Even if your network is small, chances are you have a need for remote access, whether for traveling employees, telecommuters, or remote branches. You can choose from several methods of remote access, including dial-in access, VPN access through the Internet, and wireless networking. Which methods you support and how you configure them will depend on the needs of your organization and its individual users.

EXAM
70-293

OBJECTIVE

3



NOTE

Wireless access to a network is not as remote as access by modem or VPN; in fact, most wireless technologies are limited to the area of a building or small group of buildings. But wireless access shares some features with these methods: the access is typically temporary and it can be managed in many of the same ways.

Analyzing Organizational Needs

Different organizations have different needs in a remote access strategy. The following are some of the organizational needs you might need to address:

- Security of dial-in and VPN connections
- Availability of modems and connections
- Determination of which resources or subnets need to be reachable remotely
- Determination of whether existing network servers can be adapted to provide remote access

Analyzing User Needs

You also need to consider the needs of individual users when planning a strategy for remote access. The following are some needs you might have to address:

- The bandwidth requirements of users, and what their modems or connections can support
- How frequently users need to connect to the network and how critical network availability is
- The types of operating systems and computers used by clients
- Whether clients have existing Internet connections that could be used for VPN access

Selecting Remote Access Types To Allow

When you plan which types of remote access to allow, you should consider how they meet your organization's needs and the needs of the users, the expense and administrative effort involved in implementing each one, and their relative levels of security. In the next sections, we'll look in more detail at those aspects of each of the remote access types mentioned earlier: dial-in, VPN, and wireless.

Dial-In

The traditional method of remote access uses a pool of modems and a server running the Routing and Remote Access (RRAS) service. Although there are alternatives, such as VPN access, modems still have some advantages:

- Dedicated modem lines don't require encryption and communications are more difficult to intercept. This is because the connection is direct and does not go over a public data network. In addition, you can use security features available only in the phone system, such as caller ID verification and callback security.
- Although modem access is slow, its speed is consistent and unaffected by Internet usage and other issues. Thus, it might be more reliable when high bandwidth is not needed. You can also use the multilink feature to combine multiple modem links into a faster connection. (You can also use ISDN "modems" for faster dial-in access. ISDN lines are highly reliable and provide for speeds of 128 Kbps, almost three times faster than the typical analog modem connection.)
- You might be able to use existing phone lines and modems rather than purchasing or configuring new equipment for VPN access.
- Adding phone lines for clients is an expense, but additional clients do not increase the bandwidth load on an Internet connection.

Dial-in access typically uses PPP (point-to-point protocol) for communication. This is an Internet-standard protocol for dial-in connections. PPP supports a negotiation process that authenticates and authorizes the user and can also assign an IP address, DNS server addresses, and other critical configuration elements for remote access.



NOTE

SLIP (Serial Line Internet Protocol) was the original protocol used for dial-in connections. While SLIP has largely been replaced by the more reliable and secure PPP, it is still used with some older equipment, and you can support it if necessary.

VPN

A VPN (virtual private network) uses encryption to create a virtual connection, or tunnel, between a remote node and your network, using a public network such as the Internet. VPN access has a number of advantages over dial-in remote access:

- More bandwidth is available, assuming the client can obtain a broadband Internet connection or has a high-speed dedicated leased line.

- The network can accept unlimited connections from clients through a single Internet connection, without the need to add equipment for additional clients.
- Clients and corporate networks often have existing Internet connections that can be adapted for VPN use with a minimum of effort and expense.

While VPN access is theoretically less secure than a dial-up connection, because data is transmitted over a public network, Windows Server 2003 supports strong levels of encryption to minimize this risk. You can also mandate a level of encryption so that clients that do not support your minimum encryption level cannot connect to the network.

Wireless Remote Access


Wireless network access is rapidly becoming more popular as a facet of remote access strategies. Wireless networks using the 802.11 standard enable a number of wireless users to connect to your network by connecting to a wireless access point, or WAP. While wireless networks typically span a room or building, they can also be scaled upward to cover several buildings, and systems of multiple WAPs have been configured to cover an area as large as a neighborhood or town.

The 802.11 standards do allow for security, but many wireless networks are not configured for maximum security, and allowing wireless access is always a security risk. You should plan for wireless access when your users will be within range of a WAP but without access to a wired connection, and when security is not the highest priority.



NOTE

A new standard, 802.1x, adds security to 802.11 wireless networks by making use of EAP (Extensible Authentication Protocol) instead of the authentication features of PPP. 802.1x enables you to connect through multiple access points without changing the configuration and is supported by Windows XP and Windows Server 2003.



Addressing Dial-In Access Design Considerations

When you plan a system for dial-in access, you need to consider a number of factors. These include the following:

- How IP addresses will be assigned to clients
- The number and type of incoming ports to configure
- The security or administration model you will follow

Allocating IP Addresses

When clients connect to RRAS, whether through a dial-in or VPN connection, the RRAS server assigns each client an IP address. You can configure the RRAS server to allocate IP addresses from a static address pool, or by using DHCP or Automatic Private IP Addressing (APIPA).

Static Address Pools

You can configure the RRAS server to assign IP addresses from a static pool of addresses specified in the RRAS server's configuration. This requires a range of addresses that are dedicated for this purpose. Although this is often the simplest approach, keep these considerations in mind:

- Make sure the static address pool does not overlap the range of addresses assigned by a DHCP server. Two machines with the same address will cause a conflict and a loss of connection for both.
- If you are using multiple RRAS servers with separate modem pools, you will need to define a static address pool for each one and make sure there are no conflicts between the ranges you assign.
- Be sure the address pool includes at least as many addresses as there are modems for incoming connections.

You can also assign a static address for a single user, group, or a particular type of connection using a remote access policy. This process is described later in this chapter.

Using DHCP for Addressing

Rather than using a static address pool, you can configure the RRAS server to request IP addresses from a DHCP (Dynamic Host Configuration Protocol) server. If you are using DHCP to assign addresses in the network already, this technique allows you to assign remote client addresses from the same address pool and eliminate the possibility of address conflicts. It also makes it easy to manage addressing with multiple RRAS servers, because you can configure them to use the same DHCP server.



NOTE

When a remote client is configured to obtain an IP address automatically, it does not act as a DHCP client. Instead, the remote access server is responsible for assigning an address to each client during the PPP connection process. The server can request these addresses from a DHCP server or use its own pool of addresses.

Using APIPA

Finally, you can configure the RRAS server to assign addresses using Automatic Private IP Addressing (APIPA). This system uses private addresses in the range of 169.254.0.1 through 169.254.255.254, a range reserved for use by Windows networks, and is usually used when a DHCP server is unavailable. APIPA provides some of the advantages of DHCP without a dedicated server, but is usually only suitable for small networks.

If you enable the DHCP option on the RRAS server but a DHCP server is unavailable on the network, it will automatically use APIPA to issue addresses to remote clients. Clients must be configured to obtain an IP address when they connect, rather than with a specific IP address, to use this feature.

Determining Incoming Port Needs

When you are designing a dial-in remote access solution, one of the most important considerations is the number of incoming ports (modems) you will need. The following are some of the factors you should take into account:

- An estimate of the number of users who will need to concurrently access the network remotely. Keep in mind that a single user who requires access for several hours a day will require an additional modem for reliable access, but several users who use the network for only a few minutes at a time could be easily served by a single modem.
- The bandwidth available on the RAS server's connection to the LAN. If the bandwidth of all the modems combined approaches this limit, dial-in users will experience slow connections.
- The number of IP addresses available. If an address pool or DHCP server is out of addresses, additional modems will not allow additional users.

Multilink and BAP

Another factor that can affect the number of incoming ports you will need is whether you will be supporting multilink connections. This is a Windows Server 2003 feature that enables two or more ports on the RRAS server to be connected to a single client and combined into a higher-bandwidth connection.

For example, if a client connects with two 56K modems and multilink enabled, their bandwidth with a perfect connection would be 112K. In practice, if you've spent time trying to get a single modem to work at 56K, you can imagine how unlikely this best-case scenario is, and few client computers have two modems installed. Nonetheless, multilink is sometimes the cheapest way to get a high-bandwidth connection. Multilink is also often used to combine two 64K ISDN channels into a single 128K connection.



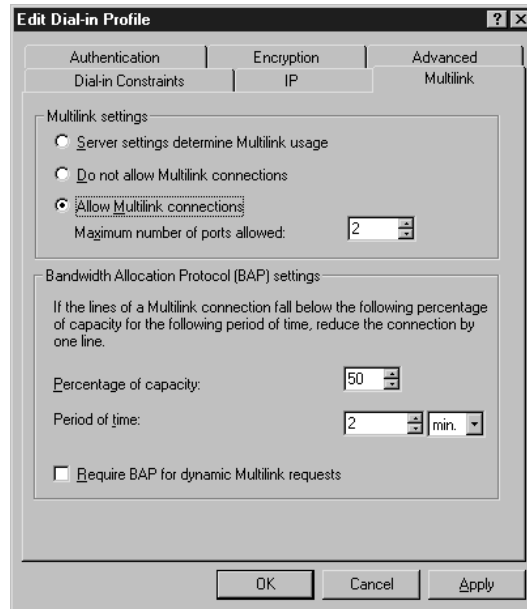
TEST DAY TIP

A basic rate ISDN connection consists of two 64 Kbps data (B) channels and a single 16 Kbps control (D) channel. The two B channels can be combined using hardware bonding, but this requires that all hardware support this feature. Multilink is a reliable and consistent way to aggregate the channels regardless of the equipment.

Windows Server 2003 also supports BAP (bandwidth Allocation protocol). This is a system that automatically disconnects one or more ports from a multilink connection if it is using only a small percentage of its capacity. This enables you to make the best use of multiple ports without relying on users to reconfigure their connections.

You can configure multilink and BAP settings as part of a dial-in profile. Remote Access Policies and profiles are described in detail later in this chapter. The **Multilink** settings tab for a dial-in profile enables you to enable or disable multilink and BAP and change their settings, as shown in Figure 7.1.

Figure 7.1 Multilink Options



Selecting an Administrative Model

There are two basic ways for you to control remote access to your network. You can configure individual user accounts to allow or disallow remote access, or configure one or more remote access Policies to control access based on users, groups, times of day, and many other criteria.

Access by User

You can allow or disallow remote access from the **Dial-in** tab of a user's **Properties** dialog box in the **Active Directory Users and Computers** console. Exercise 7.01 demonstrates how to enable remote access for a user account.


EXERCISE 7.01

ENABLING REMOTE ACCESS BY USER

Follow these steps to enable remote access for a user account:

1. From the **Start** menu, select **Programs | Administrative Tools | Active Directory Users and Computers**.
2. Click the **+** symbol next to the domain name node in the left column to display its contents.
3. Click **Users** in the left-hand column. A list of the domain's users and groups is displayed in the right-hand column, as shown in Figure 7.2.

Figure 7.2 Listing the Domain's Users and Groups



4. Click a user name to highlight it; then select **Action | Properties** from the menu or right click the user name and select **Properties** from the context menu.

5. The user's **Properties** dialog box is displayed. Click the **Dial-in** tab.
 6. The **Dial-in** properties are displayed. Select the **Allow Access** option in the **Remote Access Permission** section at the top of the dialog box.
 7. Click **OK** to exit the **Properties** dialog box and save your changes.
-

Access by Policy

You can also configure one or more Remote Access Policies for precise control of which users can reach the network through remote access. Whether a user is affected by policies depends on the setting you choose in the Dial-in tab of the user's Properties dialog box:

- **Allow access:** The user is allowed remote access regardless of policy settings.
- **Deny access:** The user is denied remote access regardless of policy settings.
- **Control access through Remote Access Policy:** Allows a Remote Access Policy to control whether the user has access.

Exercise 7.02 demonstrates how to enable remote access by policy for a user. You will learn how to create policies later in this chapter.

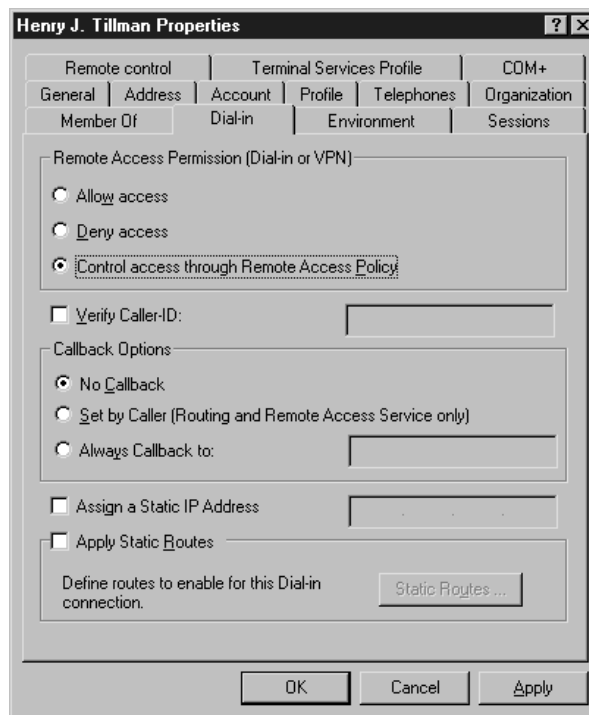
EXERCISE 7.02

ENABLING REMOTE ACCESS BY POLICY

Follow these steps to enable a Remote Access Policy for a user:


1. From the **Start** menu, select **Programs | Administrative Tools | Active Directory Users and Computers**.
2. Click the + symbol next to the domain name in the left column to display its contents.
3. Click **Users** in the left-hand column. A list of the domain's users and groups is displayed in the right-hand column, as shown in Figure 7.2.
4. Click a user name to highlight it, and then select **Action | Properties** from the menu or right-click the user name and select **Properties** from the context menu.
5. The user's **Properties** dialog box is displayed. Click the **Dial-in** tab.
6. The **Dial-in** properties are displayed, as shown in Figure 7.3. Select the **Control access through Remote Access Policy** option.

Figure 7.3 Dial-in Properties



7. Click **OK** to exit the **Properties** dialog box and save your changes.

After you have enabled remote access by policy for the user, you need to create one or more Remote Access Policies to control access. This procedure is described later in this chapter.



Addressing VPN Design Considerations

Rather than using individual modem or ISDN ports for remote access, you can configure a VPN (virtual private network) and enable any number of connections through the Internet. A VPN uses an encrypted tunnel to create a secure virtual connection and transmit private data over the public network.

Although using a VPN for remote access does not require any special hardware beyond an Internet connection for clients and the RRAS server, there are still a number of choices you must make when planning a VPN strategy. These include the VPN protocols you will support, the need for machine certificates, IP filtering, and remote access policies.

Selecting VPN Protocols

A VPN connection is created through the use of a tunneling protocol, (sometimes called a VPN protocol), supported by both the client and the server. Windows Server 2003 supports two tunneling protocols:

- PPTP (point-to-point Tunneling protocol) is an Internet standard for VPN connections based on PPP (point-to-point protocol). PPTP uses the MPPE (Microsoft Point-to-Point Encryption) system to encrypt data.
- L2TP (layer 2 Tunneling protocol) is a newer standard for a tunneling protocol, developed in cooperation between Microsoft and Cisco. L2TP is used with IPsec (IP Security) to provide encryption.

You can support one or both of these VPN protocols in your remote access strategy. Which protocols you support depends on the needs of clients, the requirements for public-key security, and whether you need the higher-security features of L2TP. These considerations are discussed in the following sections.

Client Support

Of course, a major factor in deciding which tunneling protocols you should support is the protocols supported by the client machines. The following is a summary of the VPN tunneling protocol support of Windows clients:

- PPTP is supported by Windows 95, Windows 98, Windows ME, Windows NT 4.0 and later, Windows 2000, Windows XP, and Windows Server 2003.
- L2TP is supported by Windows 2000, Windows XP, and Windows Server 2003.

If you are supporting non-Windows clients, you should determine which VPN protocols they support. The easiest way to support a wide variety of clients is to enable both VPN protocols at the server level; clients that support L2TP will use it, and other clients will use PPTP.

Data Integrity and Sender Authentication

The IPsec encryption used with L2TP supports two features that are not available with PPTP and MPPE encryption, along with the data confidentiality that is provided by both encryption protocols. You should make sure your network supports L2TP if you require either of the following:

- **Data integrity** L2TP over IPsec verifies the integrity of data by using hash algorithms (checksums).
- **Sender authentication** IPsec provides mutual authentication for the client computer and VPN server. This authentication is based by PKI (Public Key

Infrastructure) certificates and is in addition to the user authentication handled by protocols such as MS-CHAP v2 and EAP-TLS.

PKI Requirements

To support L2TP over IPSec for VPN connections, you need to install computer certificates at both the VPN server and the clients. If you do not have an existing certificate server configured on the network, this might require additional planning and configuration. PPTP does not require a PKI at all and is the only choice if you do not wish to install certificates.

Installing Machine Certificates

To use IPSec with L2TP, you need to install computer certificates at each client for encryption. Windows 2000 and Windows Server 2003 support auto-enrollment, a feature that automatically distributes certificates to computers the first time they connect to the network. If you are not using auto-enrollment, you can manually request a certificate for the computer. You can do this using the Certificates MMC snap-in or by connecting to the certificate server with a Web browser.



TEST DAY TIP

IPSec also supports user certificates. To complete a VPN connection with L2TP over IPSec, you will need a computer certificate and either a user certificate or smart card. Smart cards are described later in this chapter. Windows Server 2003 supports auto-enrollment for user certificates, unlike Windows 2000.

If you do not have a certification authority (CA) on the network, you can install Certificate Services on a domain controller. Exercise 7.03 explains how to set up the Certificates MMC snap-in and request a certificate.

EXERCISE 7.03


REQUESTING A COMPUTER CERTIFICATE

Follow these steps to configure the **Certificates** MMC snap-in and request a certificate.


1. Type **mmc** in the **Run** dialog box to start the management console.
2. An empty console is displayed. Select **File | Add/Remove Snap-in** from the menu.
3. Click **Add** and select **Certificates** from the list of snap-ins. Click **Add**.

4. A dialog box prompts you to choose whether the console will manage user or computer certificates. Select **Computer account** and click **Next**.
5. The **Select Computer** dialog box is displayed. Select **Local Computer** and click **Finish**.
6. Click **Close** and then **OK** to return to MMC.
7. Click the + icon next to **Certificates (Local computer)** to expand it. The certificate categories are now listed as folders, as shown in Figure 7.4.

Figure 7.4 Certificates MMC Snap-In



8. Select **Personal** in the left column, and then select **Action | All Tasks | Request New Certificate** from the menu.
9. The **Certificate Request Wizard** displays an introductory dialog box. Click **Next** to continue.
10. You are prompted for the type of certificate to request. Select **Computer** and click **Next**.
11. The **Certificate Friendly Name and Description** dialog box is displayed, as shown in Figure 7.5. Enter a name and description and click **Next**.

Figure 7.5 Certificate Friendly Name and Description

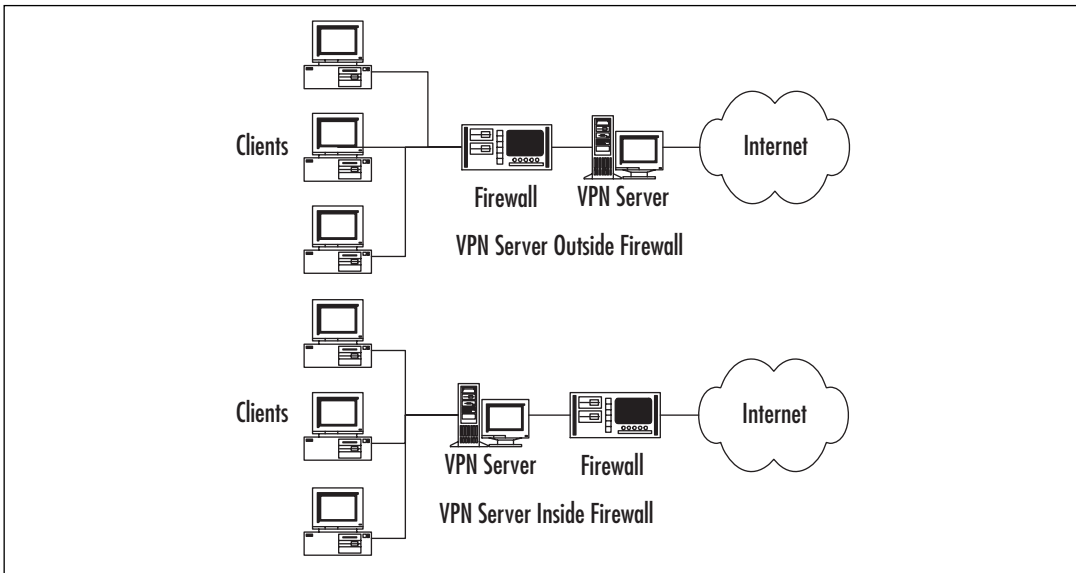
12. A completion message is displayed. Click **Finish** to exit.
-

Configuring Firewall Filters

Because a VPN server is connected to the Internet, it is often used in conjunction with a software or hardware firewall to prevent unauthorized traffic from the Internet from reaching the internal network. You can arrange the firewall and VPN server in one of two ways:

- The VPN server is directly connected to the Internet and the firewall separates it from the internal network.
- The firewall is connected to the Internet and the VPN server is behind the firewall.

Figure 7.6 shows these two configurations.

Figure 7.6 Firewall Configurations

The more common of the two arrangements is to connect the firewall to the Internet and keep the VPN server behind the firewall. In this scenario, you set up packet filters to allow all VPN traffic through the firewall. Since the VPN connection between the client and server handles authentication and security itself, this does not represent a security risk.

Creating Access Policies

After you have configured a VPN server and the required certificates, you can use Remote Access Policies to control access to the VPN. This enables you to restrict access by user or group, restrict the authentication methods allowed, control the encryption methods that can be used, and a variety of other settings. Remote Access Policies are explained in detail later in this chapter.

EXAM
70-293
OBJECTIVE
3

Addressing Wireless Remote Access Design Considerations

Wireless networks are fast becoming one of the most common network types. Although they are not cost-effective or efficient as a replacement for wired networking, wireless networks are a great choice for temporary networks, networking in areas where networking is normally difficult, or offering wireless access to customers or employees with portable computers.

Windows Server 2003's RRAS server can be used to manage wireless connections to the network. If you will be allowing wireless access, you will need to do the following:

- Configure remote access policies.
- Determine whether to use IAS for authentication.
- Configure the WAPs.

The 802.11 Wireless Standards

Today's wireless networks generally use one of the standards developed by the IEEE under the 802.11 working group and based on the original 802.11 protocol, which supported speeds of 2 Mbps in the 2.4 GHz radio spectrum. The newer standards support higher speeds and are popularly known as *Wi-Fi*. There are three current versions of 802.11 that define different wireless standards:

- 802.11b was the first standard to be widely accepted. It operates at 11 Mbps and has a range of about 50 meters. It uses the 2.4 GHz spectrum.
- 802.11a appeared in products in 2001. This standard uses the 5 GHz spectrum, has a theoretical maximum speed of 54 Mbps, but does not handle distance and obstacles as well as 802.11b.
- 802.11g is the latest standard, ratified in 2003. It uses the 2.4GHz band and is backward compatible with 802.11b equipment, but supports a theoretical throughput of 54 Mbps.



NOTE

The distance and speed ratings of wireless equipment tend to be optimistic. In practice, the speed will depend on the distance between equipment, any obstacles such as brick walls, interference from power lines or other sources, and other factors.

Using IAS for Wireless Connections

Many WAPs support RADIUS authentication. Because the security of normal wireless authentication with the 802.11 protocols is minimal, using RADIUS provides stronger authentication as well as a centralized source for authentication and accounting for all wireless access. IAS can be used for this purpose.

Because WAPs configured for RADIUS authentication rely on the presence of a RADIUS server, you might need to configure a second IAS server and specify it as a backup server in the WAP configuration. This ensures that wireless users can still connect if the primary IAS server is unavailable.

Configuring Remote Access Policies for Wireless Connections

To enable wireless connections, you need a basic remote access Policy to allow wireless users. This policy can restrict access to a group, require certificate-based authentication, and/or mandate a high level of encryption. Exercise 7.04 describes how to create this remote access policy.

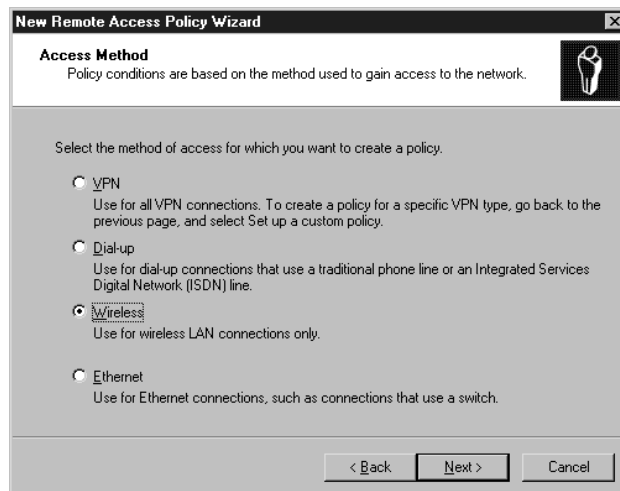
EXERCISE 7.04

CREATING A POLICY FOR WIRELESS ACCESS

Follow these steps to create a remote access Policy. For this example, the policy will enable access to the Domain Admins group, require certificate-based authentication, and require the highest level of encryption. Remote access Policies are described in more detail later in this chapter.

1. From the IAS or RRAS console, select **Remote Access Policies** in the left-hand column.
2. Select **Action | New Remote Access Policy** from the menu.
3. A welcome message is displayed. Click **Next** to continue.
4. The **Policy Configuration Method** dialog box is displayed. Select **Use the wizard to set up a typical policy** and enter **Wireless access** in the policy name field. Click **Next** to continue.
5. The **Access Method** dialog box is displayed, as shown in Figure 7.7. Select **Wireless** and click **Next**.

Figure 7.7 Access Method



6. The **User or Group Access** dialog box is displayed. Select **Group** and click **Add**. Enter **Domain Admins** and click **OK**, and then click **Next**.
 7. You are prompted to choose an EAP type to allow. Select **Smart card or other certificate** and click **Next**.
 8. A completion message is displayed. Click **Finish** to exit the wizard.
-

Multiple Wireless Access Points

You can support multiple WAPs for wireless access using RRAS or IAS for authentication. Because each access point covers only a limited area, it is common to have multiple WAPs. Keep the following considerations in mind when planning to deal with multiple WAPs:

- IAS authentication will enable all WAPs to use a central server for authentication.
- Each WAP will need to be added to the IAS server's list of clients and configured to use RADIUS authentication.
- There are several variations of the 802.11 protocols and not all devices are compatible. Be sure all WAPs and clients support the same protocols.

Placing CA on VLAN for New Wireless Clients

Wireless clients typically use certificate-based authentication, either using the EAP-TLS protocol with a user certificate or using a certificate stored in a smart card. Each client also needs a computer certificate installed in order to use EAP-TLS authentication. You need to configure a certificate server to issue certificates to wireless clients.

For new clients that might not have a certificate already, one strategy is to create a virtual LAN (VLAN) and place a certification authority (CA) on the VLAN to issue certificates. You can use a remote access policy to restrict new wireless clients to this VLAN so they will be unable to access other network resources and to limit their connection time. After a client successfully connects to the VLAN and is issued a certificate, it can reconnect using the standard wireless access policy and gain full access.

Configuring WAPs as RADIUS Clients

For WAPs to use the IAS server for authentication, you must configure both ends:

- In the IAS MMC snap-in, add each WAP as a RADIUS client.
- In the WAP's configuration, enable RADIUS authentication and specify the IAS server (or both servers, if you have a backup server configured.)

How you configure the WAP varies depending on the hardware in use. Consult the documentation provided by the manufacturer to find out how to do this.

Wireless Encryption and Security

Wireless networks are more vulnerable to snooping than wired networks, since it is difficult to control exactly how far the wireless radio signals travel. Due to this concern, several technologies have been developed to add security to 802.11 wireless networks: WEP (Wired Equivalent Privacy), 802.1x, and WPA (Wi-Fi Protected Access).

WEP (Wired Equivalent Privacy)

WEP was included as part of the 802.11 standard to provide a way to encrypt wireless data and reduce the possibility of intercepted signals. WEP is supported by most 802.11 hardware. When you activate WEP, wireless LAN cards and access points encrypt each packet before transmitting it over the antenna. WEP supports a 40- or 64-bit key for encryption, although some vendor-specific solutions support higher levels of encryption.

Although WEP provides basic privacy, it is not considered highly secure. It uses manually configured fixed keys rather than PKI, and all the equipment must be manually configured with the same key. Because the keys are rarely changed, a dedicated hacker can eventually collect enough data to break the encryption.



TEST DAY TIP

WEP only encrypts the data when it is sent between wireless equipment; it does not affect data sent over the wired network, so you should use IPSec if you need end-to-end encryption.

802.1X

The newer 802.1x standard enables EAP (Extensible Authentication Protocol) authentication methods to be used for wireless networks using the 802.11 standards. Along with supporting stronger authentication, this system can be used to exchange keys and create a more secure encrypted connection than WEP supports.

802.1x usually requires a RADIUS server to work, and is supported by Microsoft IAS in Windows 2000 and Windows Server 2003. Windows XP supports 802.1x natively for wireless connections, and support can be added to other client operating systems.



NOTE

Although 802.1x goes a long way toward improving wireless security, it is not by any means perfect—several exploits using forged packets have already been

demonstrated. For now, you should always consider wireless networks less secure than a well-secured wired network.

WPA

At this writing, the IEEE is working on a standard known as 802.11i, a replacement for earlier wireless standards with higher security. Some vendors have already implemented an early version of 802.11i known as WPA (Wi-Fi Protected Access). WPA requires the use of 802.1x authentication and a key exchange and supports stronger encryption than WEP. Windows Server 2003 and Windows XP SP1 include support for WPA. To use this feature, it must also be supported by all the wireless hardware.

EXAM 70-293
OBJECTIVE
3.2.2
3
3.2
3.2.1

Planning Remote Access Security

Windows Server 2003 includes a number of security features for remote access, including some new features that were not available in Windows 2000. When you plan a strategy for remote access security, you need to take several things into account:

- The functional levels of your domains
- The methods you will use for data encryption and authentication
- Whether you will use advanced security features such as callback security and smart cards

These items are discussed in the following sections.

Domain Functional Level

Domains hosted on Windows Server 2003 computers can have one of several different *domain functional levels*. The functional level of your domain affects which remote access security features you can use. Depending on your needs, you might need to raise the functional level of the domain to take advantage of new security features.



TEST DAY TIP

In Windows 2000 terminology, a domain's *mode* was either Windows 2000 Native or Mixed-mode. Windows Server 2003's *domain function levels* include these two options as well as options for domains with support for .NET and the new Windows Server 2003 security features.

Determining the Function Level

The domain functional level indicates whether the domain supports new security features added in Windows 2000 and in Windows Server 2003, and also whether support is available for older operating systems to participate in the domain. The following functional levels are possible:

- **Windows 2000 Mixed:** Supports a domain containing a combination of Windows 2000, Windows Server 2003, and Windows NT 4.0 servers. Active Directory support is limited and many of the security features of Windows 2000 and Windows Server 2003 are not available.
- **Windows 2000 Native:** The same as the native mode supported by Windows 2000. Supports all Active Directory features, but not the latest changes in Windows Server 2003. This level allows Windows 2000 and Windows Server 2003 DCs.
- **Windows Server 2003 Interim:** Supports a domain containing Windows Server 2003 and Windows NT 4.0 DCs, but no Windows 2000 DCs.
- **Windows Server 2003:** Supports all features of Windows Server 2003. Windows 2000 and earlier domain controllers are not supported.


You can determine the current functional level of a domain by examining its properties in the Active Directory Domains and Trusts console. Exercise 7.05 guides you through this process.

EXERCISE 7.05

CHECKING THE DOMAIN FUNCTION LEVEL

Follow these steps to check a domain's functional level:

1. From the **Start** menu, select **Programs | Administrative Tools | Active Directory Domains and Trusts**.
2. Highlight the domain name in the left-hand column.
3. Select **Action | Properties** from the menu.
4. The domain properties dialog box is displayed, as shown in Figure 7.8. The text in the lower part of this dialog box shows the current level for the domain and for the Active Directory forest.

Figure 7.8 Domain Properties**NOTE**

Along with the domain functional level, the domain controller also keeps track of the functional level of the Active Directory forest. This is the minimum functional level supported by all the domains within the forest.

Raising the Domain Functional Level

If you have determined that your domain is operating at a lower functional level than you need, you can raise the functional level. However, after this is done, you cannot lower the level. Exercise 7.06 shows the steps to follow to raise a domain's functional level.

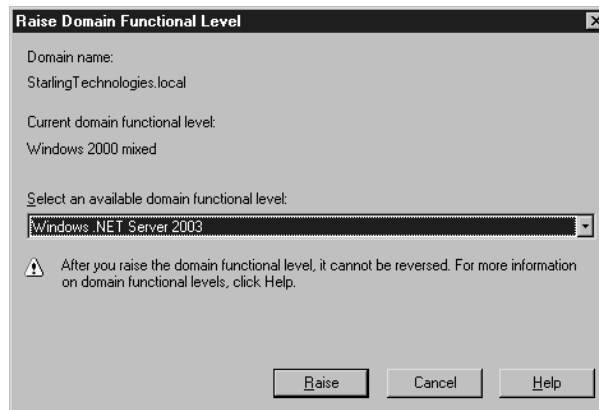
EXERCISE 7.06

RAISING THE DOMAIN FUNCTIONAL LEVEL

Follow these steps to raise a domain's functional level:

1. From the **Start** menu, select **Programs | Administrative Tools | Active Directory Domains and Trusts**.
2. Highlight the domain name in the left-hand column.
3. Select **Action | Raise Domain Functional Level** from the menu.
4. The **Raise Domain Functional Level** dialog box is displayed, as shown in Figure 7.9. Select the new level from the drop-down list and click **Raise**.

Figure 7.9 Raising the Functional Level



5. A dialog box warns you that the change will affect the entire domain and cannot be reversed. Click **OK** to confirm your choice.
6. After the process is completed, a dialog box indicates that the level was raised successfully. Click **OK** to exit.

After you have raised the domain's functional level, the change is replicated to all the domain controllers. This process can take several minutes.

EXAM
70-293
OBJECTIVE
3.2.3

Selecting Authentication Methods

When a user attempts to connect to a remote access server, one or more protocols are used for authentication, verifying the user's identity. After the user is authenticated, the RRAS server can determine what resources the user is authorized to access.

When you configure a remote access server you can select which authentication methods will be allowed. You should choose authentication methods based on their relative levels of security. Additionally, the methods you choose will depend on the client operating systems and the authentication methods they support.

Disallowing Password-Based Connections (PAP, SPAP, CHAP, MS-CHAP v1)

A number of the available authentication methods use simple user names and passwords for authentication. The simplest of these is PAP (Password Authentication Protocol). In PAP, the client transmits the user's password as unencrypted text. To ensure a secure network, you should disable SPAP, a variation of the same protocol that is used by Shiva clients.



NOTE

Shiva Corporation manufactured some of the most popular routers used in early LANs. Shiva was acquired by Intel and renamed Intel Network Systems, and was later acquired by a different company and renamed Shiva Corporation in 2002. Shiva still makes routers and VPN products, although they support modern authentication methods rather than SPAP.

CHAP (Challenge Handshake Authentication Protocol) improves security by creating an encrypted challenge and enabling the client to create a response using the password. This avoids sending the password over the network. However, CHAP stores passwords using reversible encryption, and is therefore also considered insecure. MS-CHAP v1, Microsoft's adaptation of CHAP, improves security but is superseded by the more secure version 2.

To ensure secure remote access, you should disable the less-secure authentication methods. Exercise 7.07 explains how to disable these methods and enable the more secure methods.

EXERCISE 7.07


DISABLING PASSWORD-BASED AUTHENTICATION METHODS

Follow these steps to disable PAP, CHAP, and MS-CHAP v1 authentication:

1. From the **Start** menu, select **Programs | Administrative Tools | Routing and Remote Access**.
2. Highlight the RRAS server name in the left-hand column.
3. Select **Action | Properties** from the menu.
4. The **Properties** dialog box is displayed. Click the **Security** tab.


5. The **Security** properties are displayed, as shown in Figure 7.10.
6. Click the **Authentication Methods** button.

Figure 7.10 Security Properties



7. The Authentication Methods dialog box is displayed. Uncheck the box next to Microsoft encrypted authentication (MS-CHAP).
8. Uncheck the box for Encrypted authentication (CHAP).
9. Uncheck the boxes next to Shiva Password Authentication Protocol (SPAP) and Unencrypted password (PAP). Figure 7.11 shows how the dialog box looks with all these options disabled.

Figure 7.11 Authentication Methods



10. Click **OK** to exit the **Authentication Methods** dialog box, and then click **OK** to exit the **Properties** dialog box and save the changes.
-



TEST DAY TIP

You can also restrict authentication methods by changing settings in the **Authentication** tab of the **Properties** dialog box for a Remote Access Policy. Policies are described in detail later in this chapter.

Using MS-CHAP v2

MS-CHAP v2 is a more secure version of MS-CHAP. This version uses stronger initial encryption keys, uses different keys for sending and receiving data, and supports *mutual authentication*—this means that after the server sends a challenge to the client and the client responds correctly, proving that it has the correct password, the client sends its own challenge to the server. The client disconnects immediately if the server responds incorrectly to this challenge. This enables the client to detect a server attempting to impersonate the legitimate server.

MS-CHAP v2 is supported by operating systems as old as Windows NT 4.0 and Windows 98, and is even supported by Windows 95 if the Dial-Up Networking upgrade is installed. This means that unless you are supporting very old computers, there is no need to risk security by supporting MS-CHAP v1.

Using EAP

EAP (Extensible Authentication Protocol) is not itself an authentication protocol, but provides a framework that enables authentication using a variety of different methods, known as *EAP types*. The following are the EAP types supported by Windows Server 2003:

- **EAP-MD5** A challenge-response protocol similar to CHAP. This method uses reversible encryption to store passwords, and is thus vulnerable to the same security problems as CHAP.
- **EAP-TLS (Transport Level Security)** A high-security protocol based on the SSL (Secure Sockets Layer) system used for Web server security. EAP-TLS uses encrypted certificates for authentication. It also supports mutual authentication, similar to MS-CHAP v2. This is considered the most secure authentication protocol supported by Windows Server 2003.



TEST DAY TIP

EAP-TLS is the most secure authentication method, but is not supported by all clients. Only Windows 2000, Windows XP, and Windows Server 2003 clients support this authentication method.

Using RADIUS/IAS vs. Windows Authentication

Windows Server 2003 supports RADIUS, an Internet standard for a centralized server to handle a network's authentication and accounting needs. Internet Access Server (IAS) is Microsoft's implementation of a RADIUS server, and is included with Windows Server 2003 but is not installed by default. You can install it through the **Add/Remove Programs** applet in Control Panel as a Windows component. When you configure an RRAS server, you can choose one of two authentication methods:

- **Windows Authentication:** The traditional method. Each RRAS server handles authentication itself, and you can configure the authentication methods supported in the Remote Access Policy section of the Routing and Remote Access MMC snap-in. Policies you create for one RRAS server apply only to that server.
- **RADIUS Authentication:** The RRAS server acts as a RADIUS client and contacts an IAS (or RADIUS) server to authenticate users. When RADIUS is in use, you configure authentication methods and other remote access security settings from the Remote Access Policy section of the Internet Access Server MMC snap-in. The policies you create for the IAS server apply to any RRAS server that authenticates using that server.



TEST DAY TIP

EAP supports an authentication type called EAP Over RADIUS. This is not an authentication method itself; instead, authentication requests are forwarded to a RADIUS server for processing. This enables you to install and configure EAP types on the RADIUS server and use them from any remote access server, without installing the types on each RRAS server.

Selecting the Data Encryption Level

In a VPN, you can control the level of encryption that is allowed for access. By disallowing unencrypted connections or those that use less-secure encryption, you can decrease the risk of network snooping. You can enable or disable the following levels of encryption:

- **No encryption:** Unencrypted connections, unsuitable for VPN use.

- **Basic encryption:** Encryption with a 40-bit key, considered relatively easy to break.
- **Strong encryption:** Encryption with a 56-bit key. In IPSec, this uses the DES standard for encryption. Although more secure, DES-encrypted data has been demonstrated to be breakable.
- **Strongest encryption:** Encryption with a 128-bit key for MPPE connections, or triple DES (3DES), which uses a 168-bit key (56-bit times three) for IPSec connections.

The Strongest Encryption option might not be available in international versions of Windows Server 2003 or US editions without the High Encryption Pack installed. You can enable or disable these encryption levels using remote access Policies. This process is described later in this chapter.

Using Callback Security

Callback security is a high-security system used for dial-in connections. When a client connects to a system using callback, the system disconnects and calls the client back at the client's phone number. There are two variations of callback:

- **Allowing the user to specify the callback number.** This does not provide a high level of security, but does ensure that the client's phone number can be logged and can be used to avoid long-distance charges being incurred by the client.
- **Using a callback number specified by the administrator.** This is very secure because it is difficult to impersonate a valid client, but it requires that a client always connect from the same number.

You can configure callback security as part of a remote access profile. This process is described in the final section of this chapter.

Managed Connections

For a user to connect to a remote access server via dial-in or VPN, the client computer must have the correct settings configured to match the server. Because this can be a daunting process for administrators, Windows Server 2003 supports two components to simplify the process of managing connections:

- **Connection Manager** is the client software Windows clients use to make a connection to a dial-in server or VPN server. Current versions of Windows include Connection Manager.
- **Connection Manager Administration Kit (CMAK)** is an administrator's tool that enables you to create a customized version of Connection Manager to dis-

tribute to clients. The customizations are stored in a dial-in profile and can include settings for your server, phone numbers, and even custom graphics, icons, and help files.

Connection Manager and CMAK are described in detail in Chapter 5.

Mandating Operating System/File System

Windows Server 2003 supports a new feature called Network Access Quarantine control. This feature enables you to restrict access to particular operating systems, file systems, and other aspects of the client's configuration. You use a script to accomplish this.

When Quarantine control is enabled, clients can connect normally to the RRAS server and are issued IP addresses. However, when a client first connects, it is put into quarantine mode and allowed only limited access to network resources. A script is then run through Connection Manager on the client machine to determine if the client's configuration matches the requirements. If it does, the quarantine is released and the client gains full access to the network.



TEST DAY TIP


Quarantine Control requires an IAS (RADIUS) server, a customized Connection Manager profile created with CMAK, and a custom script. It also requires that clients run Windows 98, Windows ME, Windows XP, Windows 2000, or Windows Server 2003.

Using Smart Cards for Remote Access

A smart card is a credit card-sized device that can store a public/private key pair or certificate for encryption. To use smart cards, you install card readers on client computers. Clients can request certificates from a certification authority (CA) and store them on the smart card. Because the encryption keys are not stored on client computers, this eliminates many potential security problems.

Smart cards are typically used with the EAP-TLS authentication method. Because IPSec encryption is used with L2TP VPN connections, smart cards can be used to encrypt a VPN connection that uses L2TP over IPSec.

Smart cards can store an encryption key with a large number of bits, making for highly secure communications. Their chief disadvantage is the smart card hardware; if it is damaged, a new card must be configured for the user, and if the card falls into the wrong hands, it can be used to gain unauthorized access to the network. However, smart cards use a PIN number to eliminate much of this risk.



Creating Remote Access Policies

You can manage the security of your remote access server by creating one or more Remote Access Policies. Depending on your configuration, you will need to create policies in one of these two places:

- If you are using Windows authentication, use the Remote Access Policies item under each RRAS server in the Routing and Remote Access MMC snap-in.
- If you are using RADIUS authentication, use the Remote Access Policies item under the IAS server in the Internet Authentication Service MMC snap-in.

Regardless of the type of authentication you are using, the policies you create will work the same way, and the dialog boxes for creating and modifying policies are the same.



TEST DAY TIP

Keep in mind that with RADIUS authentication you have exactly one set of remote access policies defined for the IAS server. With Windows authentication there is a separate set of policies for each RRAS server.

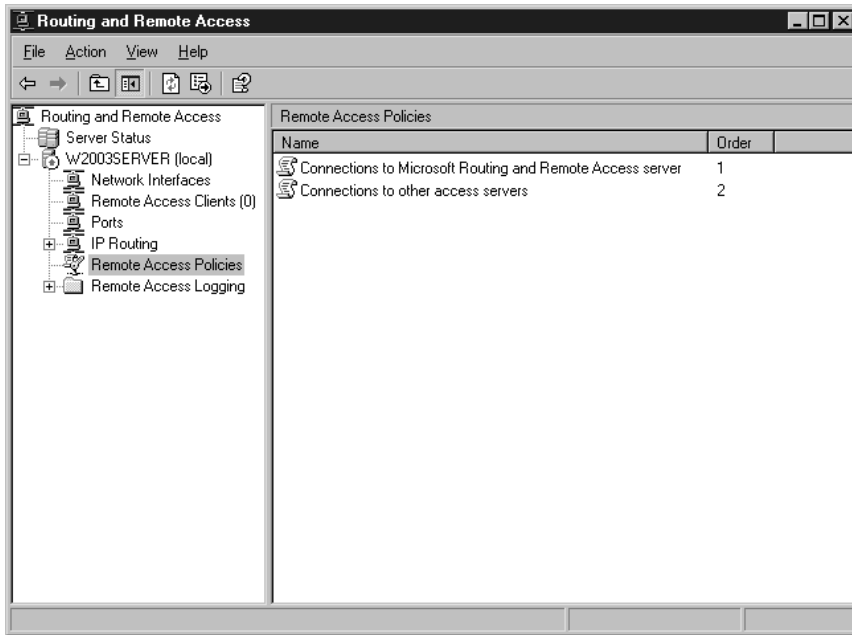
Policies and Profiles

Remote access security includes two key components:

- **Remote Access Policies** Determine which users can connect remotely and the connection methods they can use. You can have any number of remote access policies.
- **Remote Access Profiles** Provide further restrictions after the connection is established. Each policy contains exactly one profile.

Each remote access policy has an order number, or priority. You can define the order by using the Move Up and Move Down actions in the policy window. The list of policies in a default Windows Server 2003 RRAS installation is shown in Figure 7.12. Each policy can have various criteria against which connection attempts are checked. The policy can be set to either Grant or Deny access for users who match these criteria.

Figure 7.12 Remote Access Policies



When a user attempts to connect, his or her connection criteria are compared to each policy's conditions in order until a policy matches. The Grant or Deny setting of that policy then determines whether the user is allowed access. If a policy grants access, its associated profile is used to further restrict the connection.

In the following sections, you will learn how to make practical use of remote access policies and profiles to authorize or restrict remote access, and to control aspects of the connections using remote access profiles.

Authorizing Remote Access

The simplest use for a remote access policy is to authorize remote access for a particular user or group. Windows Server 2003 includes a wizard that you can use to quickly create these types of policies. After you have created a policy, you can modify the properties of the policy to make more specific settings or restrictions.

Authorizing Access By User

As described earlier in this chapter, you can use the Dial-in Properties page of a user account's **Properties** dialog box to explicitly allow or deny access to the user. This is the recommended way to authorize access by user. When you use the wizard to create a policy to authorize by user, it creates a policy that does not include any user restrictions. You can then use the user properties to allow or deny access. Exercise 7.08 shows you how to create a policy to authorize by user.

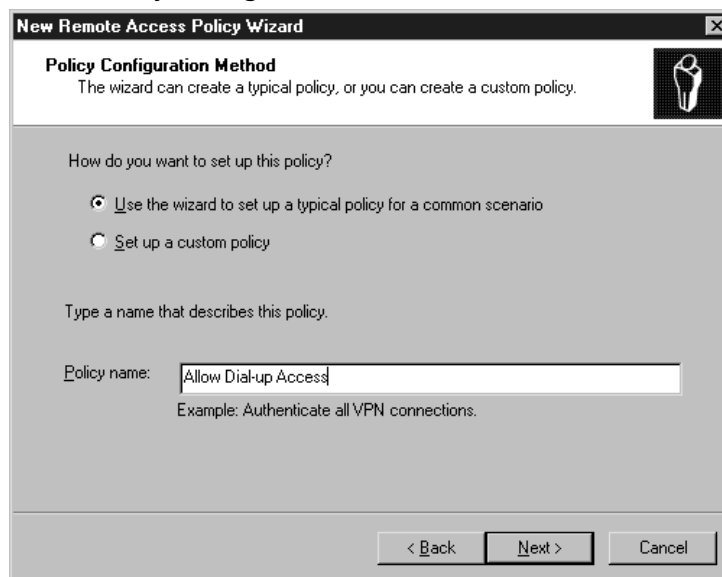
EXERCISE 7.08

AUTHORIZING REMOTE ACCESS BY USER

Follow these steps to create a policy to authorize access by user:

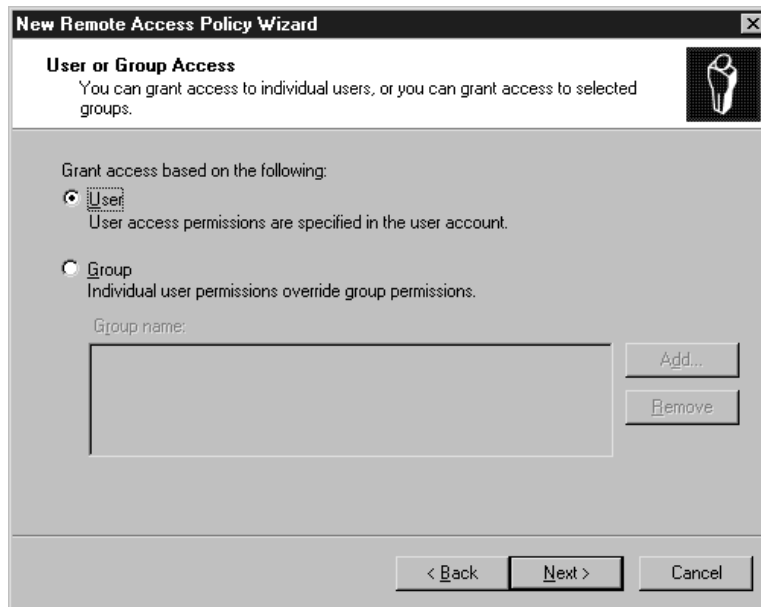
1. Select **Programs | Administrative Tools | Routing and Remote Access** from the **Start** menu. If you are using RADIUS authentication, select **Internet Authentication Service** instead.
2. Click **Remote Access Policies** in the left-hand column. A list of the current policies is displayed in the window.
3. From the menu, select **Action | New Remote Access Policy**.
4. The wizard displays a welcome message. Click **Next** to continue.
5. The **Policy Configuration Method** screen is displayed, as shown in Figure 7.13. Select the **Use the wizard to set up a typical policy** option and enter **Allow Dial-up Access** in the **Policy name** field. Click **Next** to continue.

Figure 7.13 Policy Configuration Method



6. The **Access Method** screen is displayed. You can select whether this policy will apply to Dial-up, VPN, Wireless, or Ethernet access. Select the **Dial-up** option and click **Next** to continue.
7. The **User or Group Access** dialog box is displayed, as shown in Figure 7.14. Select the **User** option and click **Next** to continue.

Figure 7.14 User or Group Access



8. The **Authentication Methods** dialog box is displayed. This dialog box enables you to choose the authentication methods this policy will accept. Click **Next** to continue.
9. The **Policy Encryption Level** screen is displayed. Select the encryption types to accept and click **Next**.
10. The wizard displays a completion dialog box. Click **Finish** to create the new policy.
11. You are returned to the **Remote Access Policies** window and your new policy has been added at the top of the list.

After you have created the policy with the wizard, you can use the **Move Up** and **Move Down** commands in the **Action** menu to change the policy order if you wish.

Authorizing Access By Group

Unlike user accounts, security groups do not include dial-in properties. If you wish to enable access for a group, you can use the wizard to create a remote access policy that includes a condition to check the user's group membership. Exercise 7.09 guides you through this process.


EXERCISE 7.09

AUTHORIZING REMOTE ACCESS BY GROUP

Follow these steps to create a policy to authorize access for the Domain Admins group:

1. Select **Programs | Administrative Tools | Routing and Remote Access** from the **Start** menu. If you are using RADIUS authentication, select **Internet Authentication Service** instead.
2. Click **Remote Access Policies** in the left-hand column. A list of the current policies is displayed in the window.
3. From the menu, select **Action | New Remote Access Policy**.
4. The wizard displays a welcome message. Click **Next** to continue.
5. The **Policy Configuration Method** screen is displayed. Select the **Use the wizard to set up a typical policy** option and enter **Allow Admin Access** in the **Policy name** field. Click **Next** to continue.
6. The **Access Method** screen is displayed, as shown in Figure 7.15. You can select whether this policy will apply to Dial-up, VPN, Wireless, or Ethernet access. Select the **Dial-up** option and click **Next** to continue.


Figure 7.15 Access Method



7. The **User or Group Access** dialog box is displayed. Select the **Group** option and click the **Add** button to add a group name.

8. The **Select Groups** dialog box is displayed, as shown in Figure 7.16. Enter **Domain Admins** in the **Enter the object names to select** field and click **OK**.

Figure 7.16 Select Groups



9. You are returned to the **User or Group Access** dialog box. Click **Next** to continue.
10. The **Authentication Methods** dialog box is displayed. Click **Next** to continue.
11. The **Policy Encryption Level** dialog box is displayed. Click **Next** to continue.
12. The wizard displays the completion dialog box. Click **Finish** to create the policy.


Restricting Remote Access

You can add any number of conditions to a remote access policy to restrict the users, connection types, and other criteria that can match the policy. Each policy can be configured to either allow access or deny access based on those criteria.

To restrict access, you can create a policy that denies access based on a set of criteria. Because each connection will use the first policy that it matches, be sure your policies for denying access are placed early in the list, before any other policy that might match the same users.

The current conditions for a policy are listed in its **Properties** dialog box. You can use the **Add** button to add a condition. There are a variety of attributes you can test to create a condition. For example, Figure 7.17 shows the **Properties** dialog box for a policy that checks the connection type and group membership.

Figure 7.17 Policy Properties



Restricting by User/Group Membership

You already used the wizard to create a simple policy to restrict by group membership earlier in this section. You can also add this condition manually to any policy using its properties. The attribute for group membership is **Windows-Groups**. You can specify one or more group memberships to match and set the policy to either grant or deny access.



TEST DAY TIP

You can restrict by user name using the **Dial-in** tab of the user's **Properties** dialog box, as described earlier in this chapter. Remote Access Policies do not include an option to restrict access by user name.

Restricting by Type of Connection

You can use the **NAS-Port-Type** attribute to restrict a remote access Policy to a particular type of connection. Connection types include modem, ISDN, wireless, VPN, and other network connections that can be used for remote access.

For example, suppose you were discontinuing the use of dial-in remote access and want to add a policy to prevent dial-in access. You would create a policy to deny access when the NAS-Port-Type attribute indicates a modem connection and place it at the top of the list to override other policies. Exercise 7.10 guides you through this process.

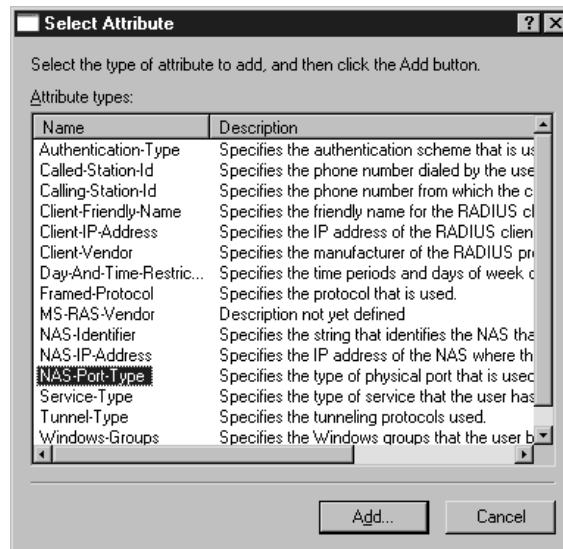
EXERCISE 7.10

RESTRICTING ACCESS BY CONNECTION TYPE

Follow these steps to create a policy that denies access to modem users:

1. Select **Programs | Administrative Tools | Routing and Remote Access** from the **Start** menu.
2. Click to highlight **Remote Access Policies** in the left-hand column.
3. Select **Action | New Remote Access Policy** from the menu.
4. A welcome message is displayed. Click **Next** to continue.
5. The **Policy Configuration Method** dialog box is displayed. Select **Set up a custom policy** and enter **Deny modem access** in the **Policy name** field.
6. The **Policy Conditions** dialog box is displayed. Click **Add** to add a condition.
7. The **Select Attribute** dialog box lists the available attributes, as shown in Figure 7.18. Select **NAS-Port-Type** and click **Add**.

Figure 7.18 Select Attribute



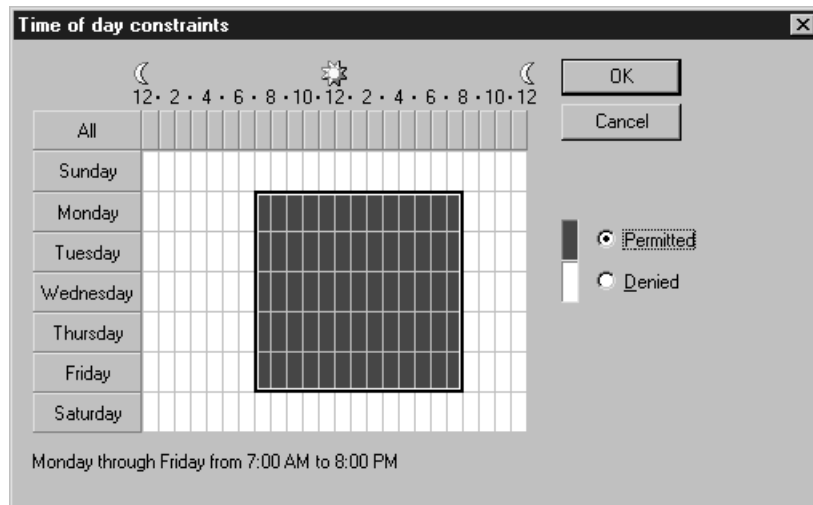
8. The available port types are listed in a dialog box. Select **Async (Modem)** and click **Add**; then click **OK**.
9. You are returned to the **Policy Conditions** dialog box. Click **Next** to continue.
10. The **Permissions** dialog box is displayed. Select **Deny remote access permission** and click **Next**.
11. The **Profile** dialog box is displayed. You can use the **Edit** button to make changes to the profile if you wish. Click **Next** to continue.
12. A completion message is displayed. Click **Finish** to create your policy.

Your new policy should appear at the top of the list by default and will prevent access by modem users regardless of other policies they may match.

Restricting by Time

You can use the **Day-and-Time-Restrictions** attribute to control the day of the week and times of day that a policy will be effective. You can use this feature to deny access at a specific time or day or to explicitly grant access at a certain time. To use this feature, use the **Add** button in the **Properties** dialog box to add a condition to a policy, and then select **Day-and-Time-Restrictions**. The **Time of day Constraints** dialog box, shown in Figure 7.19, enables you to allow or deny access for each hour of the day and each day of the week.

Figure 7.19 Time of Day Constraints



Restricting by Client Configuration

As mentioned earlier in this chapter, you can use the Network Access Quarantine Control (NAQC) feature to restrict connections based on aspects of a client's configuration: the operating system, file system, and even details of which security updates have been installed. You need to create a custom script or program to check the client's configuration to implement this feature.

NAQC is included with the Windows Server 2003 Resource Kit. It includes several components:

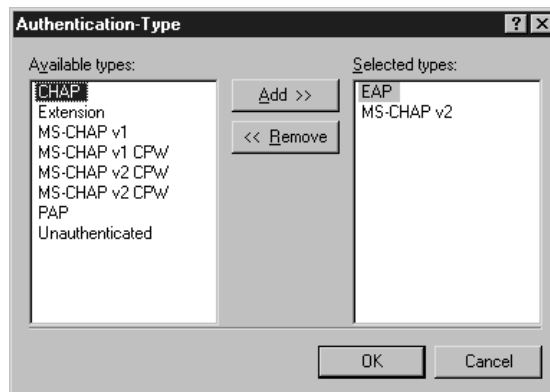
- The Remote Access Quarantine Agent service (RQS.EXE) runs on the RRAS servers.
- A custom script to check the configuration. The script can use RQC.EXE, included in the Resource Kit, to notify the quarantine agent whether the client passed its tests.
- Connection Manager, using a custom profile and a post-connect action to run the script.
- A RADIUS (IAS) server to manage authentication.
- A remote access Policy that uses the quarantine attributes, installed with the quarantine agent, to determine whether the connection has been authorized by the script.

NAQC is supported by Windows 98 SE and later clients that support Connection Manager. For details on implementing a quarantine script, consult Microsoft's TechNet site.

Restricting Authentication Methods

You can use the **Authentication-Type** attribute to restrict a policy to certain authentication types. When you add this attribute, you can use the **Authentication-Type** dialog box to add one or more of the possible authentication types, as shown in Figure 7.20.

Figure 7.20 Restricting by Authentication Method





EXAM WARNING

You can also restrict authentication methods in the **Security** tab of the RRAS server's **Properties** dialog box, as described earlier in this chapter. If a method is disabled in the server's properties, it will not be used even if it is enabled for a remote access Policy.

Restricting by Phone Number or MAC Address

You can use the following two attributes to add a phone number condition to a remote access Policy:

- **Called-Station-ID**: The phone number the user called.
- **Calling-Station-ID**: The phone number the call originated from (Caller ID).

Controlling Remote Connections

After a connection is established by matching a remote access Policy, the profile associated with the policy is used to control what the user can do with the connection. Some of the most useful profile settings include the following:

- The amount of time the user is allowed to remain connected or remain idle
- The encryption methods that will be allowed
- Which traffic will be filtered using packet filters
- The client IP address.

Controlling Idle Timeout

The *idle timeout* is the amount of time the RRAS server will keep a session connected when there has not been any traffic to or from the remote access server. You can use this setting to ensure that clients who finish using their remote connection but fail to disconnect are disconnected automatically.

The idle timeout is part of a remote access profile. You can change the timeout on the **Dial-in Constraints** tab of the **Edit Dial-in Profile** dialog box. Exercise 7.11 describes how to change this setting.

Controlling Maximum Session Time

Along with the idle timeout, you can define a maximum amount of time a client can remain connected to the server whether they use the connection or not. When your supply of incoming ports is limited, this is one way to ensure that ports are opened up to enable other users to connect.

The maximum session time is also defined in the **Dial-in Constraints** tab of a profile. Exercise 7.11 demonstrates how to change the idle timeout and session time for a profile.

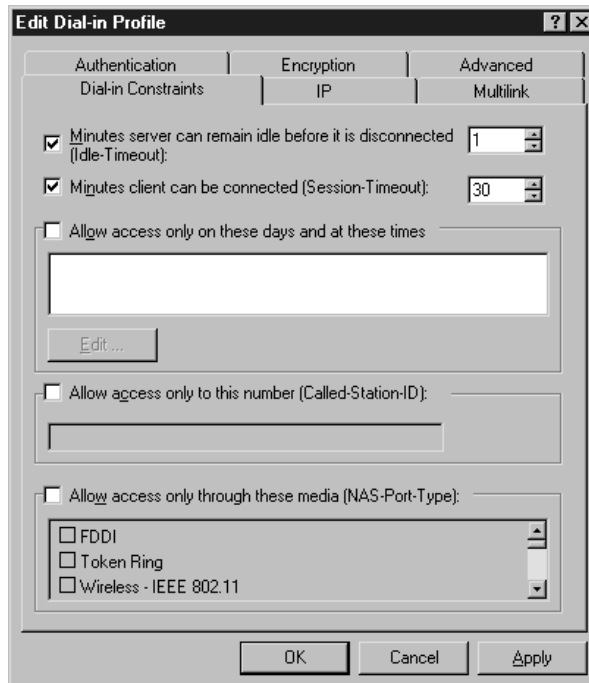
EXERCISE 7.11

CONTROLLING IDLE AND SESSION TIMES

Follow these steps to modify the idle and session times for a remote access policy's profile.

1. From the **Routing and Remote Access** console, select **Remote Access Policies** in the left-hand column. A list of the current policies is displayed in the window.
2. Click one of the policies in the window to highlight it. Select **Action | Properties** from the menu.
3. The **Policy Properties** dialog box is displayed. Click the **Edit Profile** button.
4. The **Edit Dial-in Profile** dialog box is displayed, as shown in Figure 7.21. Check the box next to **Minutes server can remain idle before it is disconnected** and select a number of minutes.

Figure 7.21 Edit Dial-in Profile



5. Check the box next to **Minutes the client can be connected** and select a number of minutes.
 6. Click **OK** to return to the **Policy Properties** dialog box.
 7. Click **OK** to save your changes and return to the RRAS console.
-

Controlling Encryption Strength

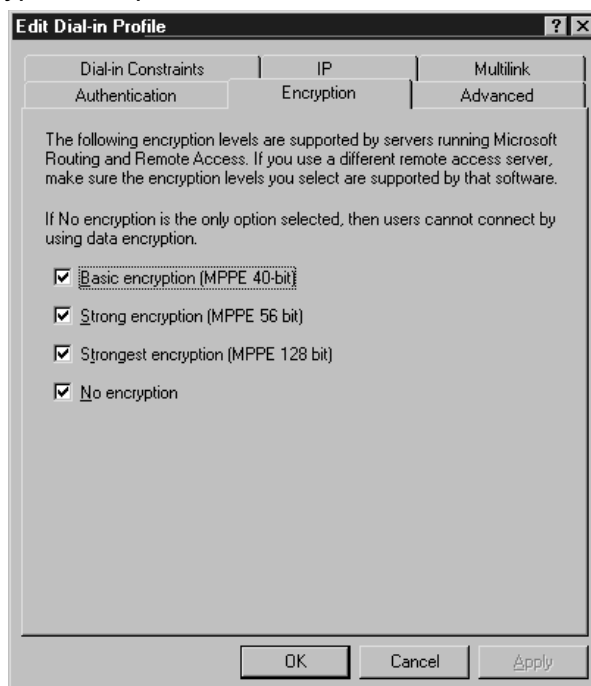
You can use the settings in the **Encryption** tab of a remote access profile's **Properties** dialog box to allow or disallow particular types of encryption for a VPN connection.

Encryption types include the following:

- Basic encryption (MPPE 40-bit)
- Strong encryption (MPPE 56-bit)
- Strongest encryption (MPPE 128-bit)

Which encryption type is used depends on what the server and the client support, but you can use this setting to prevent access with inadequate encryption. The **Encryption** tab of the **Properties** dialog box is shown in Figure 7.22.

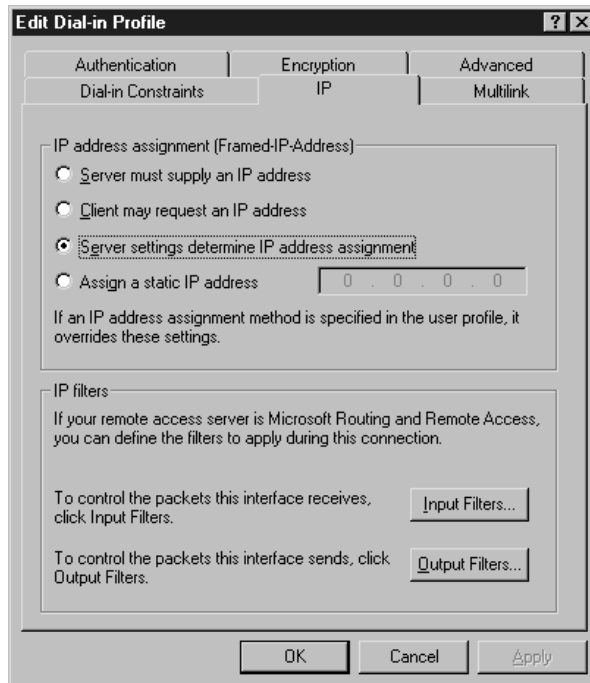
Figure 7.22 Encryption Properties



Controlling IP Packet Filters

You can use IP packet filters to filter incoming or outgoing traffic for connections that match a particular remote access profile. You might find this useful for denying access to a VPN from particular locations, or only allowing access from a particular address. You can manage outgoing and incoming packet filters from the IP settings tab of the **Profile Properties** dialog box, as shown in Figure 7.23.

Figure 7.23 IP Settings



Controlling IP Address for PPP Connections

You can also use the **IP settings** to control IP address assignment for PPP (dial-in) connections. The following options are available:

- Server must supply an IP address
- Client may request an IP address
- Server settings determine IP address assignment
- Assign a static IP address

The last option enables you to specify a single IP address to be assigned to clients that match this profile. If you use this feature, be sure that only one client at a time will match the profile, because the IP address can only be assigned to one client.

Creating a Plan to Offer Remote Assistance to Client Computers

Remote Assistance is a new feature that's designed to allow Windows XP Professional and Windows Server 2003 users to request help from another user. The user requesting help typically sends an request for assistance using Windows Messenger or e-mail via the **Help and Support Center**. The request includes an attachment that contains details of how to connect to the user's PC that the recipient will double-click to begin a Remote Assistance session with the requesting user's PC. Once connected, the helper can view the desktop of the requesting user and chat online with him or her. The helper can also, with the user's permission, take control of his desktop.

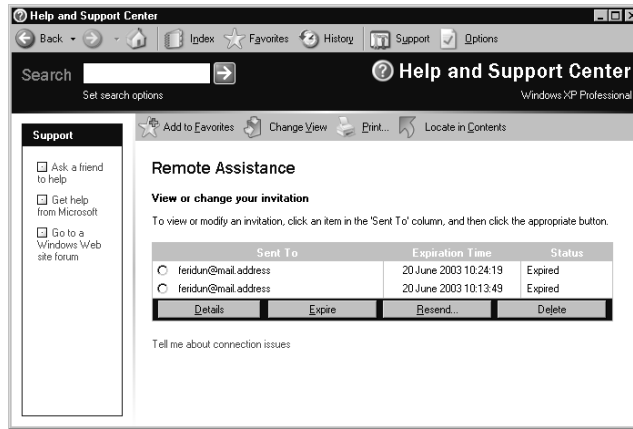
How Remote Assistance Works

Remote Assistance (RA) enables a user at one computer, referred to as the "Novice", to request help from a user at another computer, called the "Expert". The underlying technologies at work with RA are Windows Terminal Services, and the RDP protocol. Although these are the same technologies that were originally developed for thin client computing and that are used for RDA and terminal server, Remote Assistance is *not* designed to be a thin client solution, but rather a support and troubleshooting tool only. Another difference between RA and traditional Terminal Services is that typically a session will be initiated when a Novice sends an invitation to a Expert, soliciting their assistance. The Novice must typically be present at the machine that needs assistance in order to allow the Expert to access his or her system after the Expert receives and accepts the invitation.

A Remote Assistance request can optionally include an "expiry" (expiration) date, after which the Remote Assistance request is no longer valid. This is used to reduce the risk of unauthorized access to the user's computer. The user requesting help can also require the helper to use a password to connect to his or her computer. The user must communicate this password to the helper. Users can review their invitations in the **Help and Support Center**. Figure 7.24 shows a summary of invitations that have been sent from a particular computer. Using RA, the Expert actually views and interacts with the same desktop and applications that the Novice is using, at the same time that the Novice is using it. This is very different from the other forms of Terminal Services, in which a connection is established to a unique session on the Terminal Services computer. During an RA session, both the Novice sitting at the keyboard and the remote assistant (Expert) can control the computer at the same time. With Remote Desktop for Administration or the terminal server role, a user can connect from a wide range of client systems without permission, provided the user has a valid username and password.

Just as with any form of Terminal Services, Remote Assistance uses the RDP protocol so that only screen updates are sent to the client (in this case, the Expert) while keystrokes and mouse movements are sent back to the server (in this case, the Novice). In this way, RA provides remote support and control of client desktops while involving very little use of bandwidth.

Figure 7.24 Summary of Remote Assistance Invitations



Using Remote Assistance

As with Remote Desktop for Administration, the Remote Assistance (RA) components of Windows 2003 are installed with the operating system. And, just as Remote Desktop for Administration needs to be enabled and configured before you can use the feature, the same is true for RA.

Two major components comprise the default RA installation: the Terminal Services service and the Remote Desktop Help Session Manager service. In addition to installing these two components, Microsoft also creates a special user account for connections involving RA, called HelpAssistant_XXXXXX. On your system, the X's will be replaced with a unique alphanumeric code, and the account name will appear as something similar to this: HelpAssistant_e4bb43. This account will be disabled until you enable RA. As we've mentioned, although RA is based on and uses Terminal Services, it works very differently from Remote Desktop for Administration or the terminal server role. Let's take a closer look at how RA works.



TEST DAY TIP

Be sure that you are familiar with Remote Assistance (RA). As a new component in the Windows server family, and one that relates directly to test objectives, it is likely to be featured in one or more exam questions.

Configuring Remote Assistance for Use

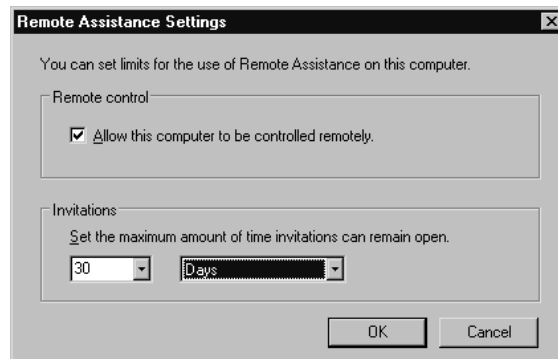
RA is relatively easy to configure; you use the same tab that is used to configure Remote Desktop for Administration. To enable RA, go to **Control Panel** and select the **Remote** tab in the **System** properties. Select the check box next to **Turn on Remote Assistance**

and allow invitations to be sent from this computer, located in the **Remote Assistance** section of the tab.

Invitations do not stay valid indefinitely. They have an expiration time of one hour by default, but the Novice can alter the expiration time of the invitations he or she sends, from 0 minutes to 99 days. The acceptance and opening of a session in response to an invitation does not cause it to expire; it is good until it reaches the specified expiration time. In other words, if you save an invitation to a file with an expiration time of 30 days, that invitation can be used to establish RA connections as many times as desired within that 30-day time-frame. To modify the default expiration time, perform the following steps:

1. Click **Start | Control Panel | System**.
2. Click the **Remote** tab.
3. Click the **Advanced...** button.
4. Choose the desired number (0 to 99) and interval (minutes, hours, or days) under the **Invitations** section in the **Remote Assistance Settings** dialog box, as shown in Figure 7.25.

Figure 7.25 The Remote Assistance Settings Dialog Box



In addition to modifying the expiration time, the **Remote Assistance Settings** dialog box can be used to enable the Expert to control the Novice's desktop and applications during an RA session, or alternately prevent them from doing so. When the **Allow this computer to be controlled remotely** box is checked, the Expert will be allowed to send mouse and keyboard input to the Novice's system and interact directly with his or her desktop and applications. When it is unchecked, the Expert will be able to see the Novice's desktop and any actions the Novice performs, but cannot control the cursor or send keyboard commands.



NOTE

It is important to be aware that, when you enable Remote Assistance (RA), the **Allow this computer to be controlled remotely** checkbox is enabled by default.

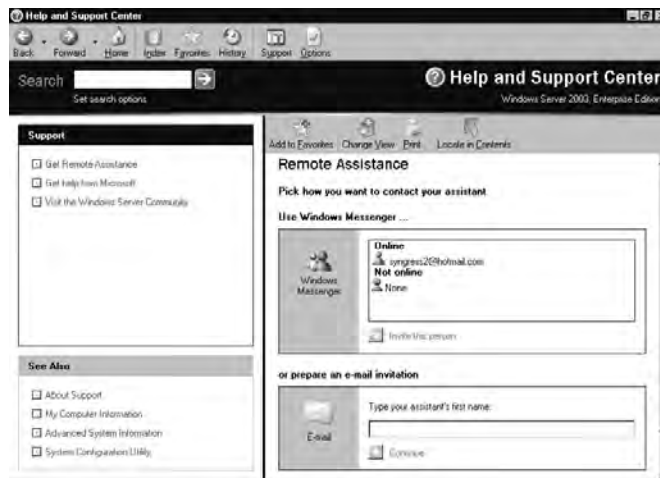
Asking for Assistance

A Novice can use a variety of methods to send an invitation using Remote Assistance:

- The request can be sent using Windows Messenger.
- The request can be sent via e-mail.
- The request can be saved to a file.

To create an invitation, open **Help and Support** from the Windows **Start** menu. On the right side of the **Help and Support Center** utility, click **Remote Assistance** under the Support heading. In the next screen, click the **Invite someone to help you** link. You will then be able to select the method that you want to use in asking for assistance, as shown in Figure 7.26.

Figure 7.26 The “Pick how you want to contact your assistant” Screen in Remote Assistance



EXAM WARNING

Although a Remote Assistance (RA) session can be solicited using an invitation sent in a file or via e-mail, Microsoft emphasizes sending an invitation using Windows Messaging. You should make sure you are familiar with all of the details of this method of solicitation.

Using Windows Messenger to Request Help

Windows Messenger is a chat program available from Microsoft and installed in Windows XP by default that is similar to ICQ and AOL Instant Messenger. (MSN Messenger is a separate but related application; both use the .NET Messenger Service). When you use Windows Messenger for RA, the invitation travels through a messaging server infrastructure that can include the Internet, or can work with Microsoft Exchange Server within the LAN. Expert and Novice "tickets" (data packets) that contain connection information are exchanged through this infrastructure. However, after these have been exchanged, the actual RDP connection attempt and subsequent session take place directly between the Novice and Expert computers.

Windows Server 2003 does not install Windows Messenger by default. If you have not installed it prior to arriving at the Remote Connection screen, you will only see a link notifying you that it is not installed and prompting you to download and install it. If Messenger is installed, the user from whom you wish to solicit help must be on the network and logged on to his or her Windows Messenger client. If this is the case, you can click the name of the contact from whom you want to solicit assistance, followed by the **Invite this person** link. The person you invited can then accept the invitation. A Remote Assistance dialog box will display on your screen until the person accepts, or until you click the cancel button on the dialog box.

You can also request assistance from within the Windows Messenger application, by double-clicking a contact to establish a conversation with him or her and then selecting the **Ask for Remote Assistance** link on the right side of the conversation window. This will add a notification to your conversation window, with a link on which you can click to cancel the request. You will also be notified in the conversation window when the person receives and accepts your request.

Remember that Remote Assistance only works on computers running Windows XP and 2003. If your invitation is sent to a person at a computer running the Windows 2000 or earlier operating system, or a non-Microsoft operating system, it will not be received.

Responding to a Request for Help Using Windows Messenger

If the Expert to whom an invitation is sent has the Windows Messaging application running, a request from a Novice for assistance will be displayed in a Conversation window on the Expert's system. The Expert can click the **Accept** link in the window (or use the key combination **Alt + T**) to initiate the connection, or click the **Decline** link (or use the key combination **Alt + D**) to reject it. If it is neither accepted nor declined before the invitation expires, the Expert will be unable to establish a connection in response to that invitation.

Using E-Mail to Request Help

To use e-mail to send an RA invitation, you must first have a default mail client configured on the Windows Server 2003 computer. This mail client can be Microsoft Outlook Express, which is installed with Windows, Outlook (installed as a separate application or with

Microsoft Office), or a third-party mail application. To create an RA invitation using e-mail, follow these steps:

1. Open the **Help and Support** utility from the Window's Start menu.
2. On the right side of the **Help and Support Center** screen, click **Remote Assistance** under the **Support** heading.
3. On the next screen that is displayed, click the **Invite someone to help you** link.
4. On the next screen, under the **or prepare an e-mail invitation** section, type the first name of the person you want to use as an Expert in the **Type your assistant's first name:** text box and click the **Continue** link.
5. The next screen contains two sections. The first is entitled **Set the invitation to expire** and contains a drop-down box for specifying a number between 0 and 99 and an interval drop-down box with selections for minutes, hours, or days. This means the possible time period during which the invitation is valid ranges from 0 minutes to 99 days.
6. The second section of this screen is entitled **Require the recipient to use a password** and is enabled by a check box. The check box is selected and this section is enabled by default. The intent is that, should the invitation accidentally fall into the wrong hands, a password would still be required to use it. Obviously, you should not include the password in the e-mailed invitation. Instead, you should communicate it to the person in some other manner (for example, by telephone). The password is entered twice, once in the **Type password:** text box and again in the **Confirm password:** text box.
7. After the password had been entered into each box, the **Create Email Invitation** button at the bottom of the screen activates and can be clicked.
8. The final screen is entitled **Was the e-mail invitation successfully sent?** When you clicked the **Create Email Invitation** button on the previous screen, your default e-mail program should have launched, with an e-mail created and ready to be sent to the person whose assistance you are requesting. This final screen alerts you to this and gives you the option to recreate the mail message in case you accidentally closed the window when it popped open. At the bottom of the screen are links to manage your outstanding invitation requests and create additional invitations. After you send the e-mail, you've finished the process of asking for remote assistance using the e-mail method.

Responding to a Request for Help From an E-Mail Request

When e-mail has been used to send you an invitation for remote assistance, a short e-mail message entitled "YOU HAVE RECEIVED A REMOTE ASSISTANCE INVITATION" will show up in your inbox. The message will contain a link to click, which will look something like this:

<https://www.microsoft.com/remotassistance/s.asp#1AjK8A2TD,4H8SQYYfvIpQF5prHYajrReyrAd2j6oHb4Qe/Eo1Ahs=.zb2.0RJ81UIfxb4Xfkp8thzdy8A=Z>.

When you click the link, your browser will open to a page on Microsoft's Web site. The entire process of the two computers finding each other using this method takes place through Microsoft's Web site. In addition, email-based Remote Assistance depends on a downloaded control.

When you visit the site, a **Security Warning** dialog box will appear and you will be prompted to specify whether you wish to install the **Remote Assistance Server Control**. If you select **Yes**, the control will download and the page will load. If you are not accessing the page from a Windows XP or Server 2003 computer, a message will display, informing you that you must be running one of these operating systems to complete the connection. If you are accessing the Web page from a Windows XP or 2003 computer, you will see a button entitled **Start Remote Assistance** in the middle of the Web page. When you click this button, a small Remote Assistance dialog box appears, prompting you to enter the password associated with the invitation (if one was used). After you have typed in the password, click the **Yes** button to begin the connection.

Using a Saved File to Request Help

The third and final way of requesting assistance is to use a saved file. Obviously, if you use this method, you need to somehow transfer the file containing the invitation to the Expert. This can be done in one of several ways:

- You can e-mail the file.
- You can save the file to a share on the network.
- You can create a link to the file on a Web page.
- You can save the file on a floppy diskette and hand it to the person.

To create an RA invitation using a saved file, open the **Help and Support** utility from the Windows **Start** menu. On the right side of the **Help and Support Center** screen, click **Remote Assistance** under the Support heading. In the next screen that is displayed, click the **Invite someone to help you** link.

At the bottom of the next screen, click the **Save invitation as a file (Advanced)** link. This leads to a screen that contains two parts. The first is entitled **Enter your name** and it contains a text box into which you type your name. When you send someone a request using Windows Messenger or e-mail, the recipient can easily see who sent the request. This is not true with a file-based request, so this dialog box is used to embed that information into the request and make it readily available to the Expert.

The second portion of this screen is entitled **Set the invitation to expire** and contains a drop-down box that enables you to specify a number between 0 and 99, and an interval drop-down box with selections for minutes, hours, or days. The possible range for the duration of a valid invitation is from 0 minutes to 99 days.

After you fill in the requested information, click the **Continue >** button at the bottom of the screen. The following page contains a section entitled **Require the recipient to use a password**, which can be enabled by checking a check box. By default, the check box is selected and this requirement is enabled. Again, the intent is that if the invitation accidentally falls into the wrong hands, at least a password will be required to use it. The password must be entered twice, once in the **Type password:** text box and again in the **Confirm password:** text box.

After the password has been entered into each box, the **Save Invitation** button at the bottom of the screen activates and can be clicked. This displays a **Save As** dialog box that enables you to specify a name and location for the file. The file will be saved with an *.msrcincident* extension. After it is saved, the final screen is displayed. It confirms the file name and where it was saved. At the bottom of the screen, there are links to manage your outstanding invitation requests and create additional invitations. Exercise 7.12 walks you through the steps of creating a saved file to use with Remote Assistance.

EXERCISE 7.12

CREATING A SAVED FILE FOR REQUESTING HELP

1. Open the **Help and Support** utility from the Windows **Start** menu.
2. On the right side of the Help and Support Center screen, click **Remote Assistance** under the Support heading.
3. On the next screen that is displayed, click the **Invite someone to help you** link.
4. At the bottom of the next screen, click the **Save invitation as a file (Advanced)** link.
5. In the **Enter your name** text box, type your name
6. In the **Set the invitation to expire** drop-down boxes, specify when the invitation should expire and then click the **Continue >** button.
7. Type the password you would like to use in the **Type password:** and **Confirm password:** text boxes. If you do not wish to use a password, clear the check box next to **Require the recipient to use a password**.
8. Click the **Save Invitation** button at the bottom of the screen.
9. In the **Save As** dialog box, specify a name and location for the file.
10. Review the information on the final screen and close the Help and Support utility.

Responding to a Request for Help that was made using a Saved File

Responding to a remote assistance request that has been saved to a file is a simple matter of double-clicking the file. When you do this, a small Remote Assistance dialog box appears, asking you to enter the password associated with the invitation if one was specified. After you type in the password, click the **Yes** button to initiate the connection. In the following section, we show you how to complete the connection process for each of the methods described, and demonstrate what you can do when the connection has been established.

Completing the Connection

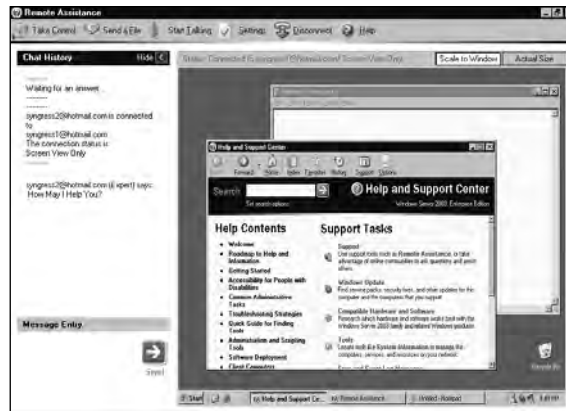
After the Expert user accepts a request for assistance, a small **Remote Assistance** dialog box pops up on the Expert's computer with a message indicating that a connection is being attempted. When the connection is established, the full **Remote Assistance** application opens, displaying a status message that says it is waiting for an answer from the Novice computer. When the connection is accepted by the Novice user, the status of the **Remote Assistance** application changes to **connected**.

During this time, the Novice's system displays a small **Remote Assistance** dialog box that asks the user if he or she wants to allow the Expert to view the computer's screen and chat with him or her. If the Novice clicks the **No** button, the connection is rejected. If the Novice clicks the **Yes** button, the connection is established. If too much time passes after the Expert attempts to establish the connection and before the Novice accepts it, a dialog box opens to inform the Novice that the invitation was accepted but has expired. This dialog box also states that a new invitation needs to be generated and offered. A dialog box is also displayed on the Expert's computer, indicating that the remote connection could not be established. When a connection is successfully established, a **Remote Assistance** application opens on the Novice's system.

Using the Completed Connection as the Expert

The Remote Assistance application on the Expert's computer consists of a tool bar across the top, a chat option on the left side and a replica of the Novice's remote desktop on the right. This is shown in Figure 7.27.

Figure 7.27 The Remote Assistance Utility on the Expert's Computer



The buttons on the tool bar across the top include the following:

- **Take Control** Initiates a request to enable the Expert to remotely control the cursor and keyboard input on the Novice's computer. When this button is clicked, a dialog box pops up on the Novice's computer, asking the Novice to allow or reject control by the Expert. Remote control is only possible if the **Allow this computer to be controlled remotely** box is checked on the **Remote** tab of the **System** properties in **Control Panel**. If remote control is accepted by the Novice, a dialog box appears in the **Remote Assistance** application on the Expert's computer over the display of the Novice's desktop, stating that remote control has been accepted. Either party can end the remote control at any time by using the **ESC** key. After remote control is established, the Remote Control button changes to read **Release Control** and can be clicked to end the remote control of the session without ending the RA session itself. Both the Novice and Expert can control the cursor and keyboard input for the Novice's system, so it is recommended that only one party be using the pointing device or typing at any given time. The Expert can use Remote control by clicking on the Novice desktop that is displayed in his or her **Remote Assistance** application.
- **Send a File** Enables you to transmit a file from the Expert's to the Novice's computer.
- **Start Talking** Establishes an audio connection between the Novice's and Expert's computers for voice and/or video communication. When this button is clicked, the **Audio and Video Tuning Wizard** opens. The wizard enables you to specify and test your microphone, audio card, and other related settings.
- **Settings** Opens the **Remote Assistance Settings** dialog box and enables adjustment of audio quality in accordance with the capacity of the underlying

network. The **Audio and Video Tuning Wizard**, mentioned in the previous bullet point, can also be opened from this dialog box.

- **Disconnect** Terminates the connection between the Novice's and Expert's computers and ends the RA session.
- **Help** Displays the About Remote Assistance help screen.

The left side of the Remote Assistance application on the Expert's computer contains a chat window. This enables the Novice and Expert to exchange text messages. In addition to chat communication, this portion of the application also contains status messages (such as the names of users who are part of the connection, whether remote control is enabled, how to stop remote control, etc).

The right side of the **Remote Assistance** application on the Expert's computer displays the desktop of the Novice's system. When the connection is initially established, the desktop appears in View Only mode. This enables the Expert to view the desktop of the Novice, but the Expert cannot interact with it. The Expert can still exchange text messages or voice communications with the Novice in this mode, and can exchange files. If the Expert and Novice agree to switch from View Only to Remote Control, the Expert can then interact with the remote desktop and applications on the Novice's system. To do this, the Expert uses his or her pointing device and keyboard to select and input data into the desktop that is displayed on the right side of the **Remote Assistance** application.

Using the Completed Connection as the Novice

The Remote Assistance application on the Novice's computer consists of a chat window on the left side and a series of option buttons along the right, shown in Figure 7.28.

Figure 7.28 The Remote Assistance Utility on the Novice's Computer



This application enables the Novice to send messages to and receive messages from the Expert. It also contains the following buttons:

- **Stop Control** Terminates the ability of the Expert to control the cursor and keyboard input on the Novice's computer.
- **Send a File** Enables transmitting a file from the Novice's to the Expert's computer.
- **Start Talking** Establishes an audio connection between the Novice and Expert computers for voice and/or video communication. When clicked, the **Audio and Video Tuning Wizard** opens. The wizard enables you to specify and test your microphone, speaker, and related settings.
- **Settings** Opens a dialog box that enables the adjustment of audio quality in accordance with the capacity of the underlying network. The **Audio and Video Tuning Wizard** can also be opened from this dialog box.
- **Disconnect** Terminates the connection between the Novice's and Expert's computers and ends the RA session.
- **Help** Brings up the **About Remote Assistance** help screen.

The left side of the Remote Assistance application on the Novice's computer contains a chat window. This enables the Novice and Expert to exchange text messages. In addition to chat communication, the left side of the application also displays status messages such as the names of users who are part of the connection, whether remote control is enabled, how to stop remote control, etc.

Managing Open Invitations

Sometimes you might want to know the names of users with whom you have active RA invitations open. You might want to cancel an invitation because you've solved the problem or because you want someone else to help you. Help and Support Center provides a number of options for managing open invitations.

To manage your active invitations, follow these steps:

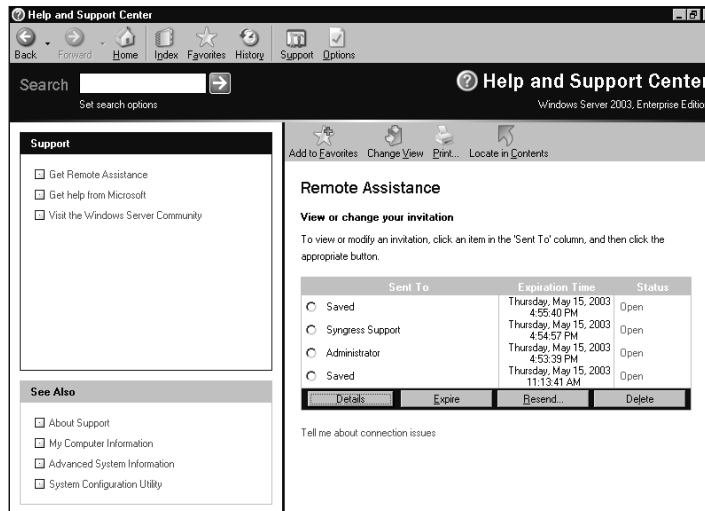
1. Open the **Help and Support** utility from the Windows **Start** menu.
2. On the right side of the Help and Support Center screen, click **Remote Assistance** under the Support heading.
3. On the following screen, click the **View Invitation Status (X)** link. The **(X)** will be replaced on your screen by the number of invitations you have outstanding.
4. The next screen will show you a list of the invitations that are outstanding. The list consists of three columns: Sent To, Expiration Time, and Status. The Sent To column contains the name of the person to whom you sent the Windows

Message or e-mail. If you saved the request to a file, this column will display the word “Saved.” The Expiration Time column will show the date and time that the invitation will expire. The Status column will show whether the invitation’s status is Open or Expired. Now you can view or modify any of these invitations.

Each invitation will have a radio button next to it, as shown in Figure 7.29. You can click a radio button to select one of the invitations, and then choose an action to perform using the buttons under the list box. The buttons include:

- **Details** Enables you to view to whom the invitation was sent, when it was sent, when it expires, its current status, and whether it is password protected.
- **Expire** Enables you to cause an invitation to expire immediately, regardless of the expiration time that was set when the invitation was originally created.
- **Resend...** Can be used only with expired invitations. When selected, this option displays a screen that walks you through the creation process for the invitation all over again. Remember that the request was originally saved to a file or sent via e-mail. Because of this, the screens and options presented are identical to those outlined earlier in the chapter.
- **Delete** Enables you to permanently delete the invitation. If the invitation’s status is Open when you select to delete it, a dialog box will pop up, informing you that the invitation will not be usable for connection. If the invitation’s status is Expired, it is simply deleted and no pop-up box appears.

Figure 7.29 The “View or change your invitation settings” Screen in Remote Assistance



Offering Remote Assistance to your Clients

Although the usual method is for the user requesting help to initiate the Remote Assistance session, it is also possible within a domain for a helper to offer assistance. An administrator can set group policy to prevent users from requesting remote assistance, or to restrict whether users will be able to enable a helper to remotely control their computers or only view them.

To configure your clients to accept Remote Assistance offers, you must ensure that the following three conditions are met:

- The Group Policy on the computer of the novice user must be configured to enable Remote Assistance offers.
- The computers of the novice and expert users must be members of the same domain, or members of trusted domains.
- Both computers must have Windows XP installed (or a newer operating system).

To configure Group Policies for Remote Assistance, you'll need to create a list of "Expert" users from that "Novice" users can accept Remote Assistance offers from. This list must consist of Domain User groups and Domain User accounts. Exercise 7.13 describes how to configure Group Policy to allow your Expert users to offer Remote Assistance to your clients.



NOTE

Experts attempting to offer Remote Assistance will not be able to connect to a Novice computer where Solicited Remote Assistance is disabled.

EXERCISE 7.13

CONFIGURING GROUP POLICY SETTINGS

1. Start the Microsoft Management Console (MMC) Group Policy snap-in: Click **Start**, click **Run**, and then in the **Open** box, type: **gpedit.msc**. Click **OK** to continue.
2. Locate the Offer Remote Assistance policy under **Local Computer Policy | Computer Configuration | Administrative Templates | System | Remote Assistance** folder.
3. Double-click **Offer Remote Assistance**.
4. On the **Offer Remote Assistance Properties** dialog box, click **Enable**.

5. Select whether or not Expert users will have View Only access to the Novice user's computer or View and Control access.
6. Click **Show**. The **Show Contents** dialog box is displayed.
7. Click **Add** to add the groups that Expert users will be able to offer assistance to.
8. Click **OK**, and then click **OK** again to close the **Show Contents** dialog box and the **Offer Remote Assistance Properties** dialog box.
9. Quit the MMC Group Policy snap-in.

These policies are effective immediately. You do not need to restart the client computers for the settings to take effect.

Once you've configured Group Policy to allow you to offer Remote Assistance to your users, you can establish a connection using the following steps:

1. Click **Start**, and then click **Help and Support**.
2. Under Support Tasks, click Tools.
3. Under Tools in the left pane, click Help and Support Center Tools.
4. Under Help and Support Center Tools, click Offer Remote Assistance.
5. Type the name or the IP address of the computer you want to connect to, and then click **Connect**.
6. Follow the directions that appear on the screen.



EXAM WARNING

Although an assistant can offer Remote Assistance without being asked, the user must give permission before the assistant can see the user's computer. In addition, the user must give explicit permission before the assistant can control the user's computer (if that feature is enabled).

Remote Assistance Security Issues

RA is a valuable tool, but it also contains serious security risks that must be planned for and managed. RA makes it easy for any user to ask virtually anyone using a Windows XP or Server 2003 computer to connect to his or her desktop. This person can be inside or a friend that is outside of your company. Although an outside person may be qualified to assist the user, in doing so they will likely receive full control of a client in your network.

This, of course, is unacceptable, because they could place malicious software on the system while in control of it, view sensitive company information that normally isn't allowed outside of the organization, etc. The best way to prevent this is to use your company's firewalls to prevent connection to RA from outside the company's network. RA uses the same port that all Terminal Services components do, 3389. Simply blocking this port on your external firewalls prevents this type of unauthorized access.

Several other key security concerns should be addressed in your company's remote assistance policies. E-mail and file-based invitations enable you to specify passwords. An invitation without password protection can be used by anyone that receives it by accident or intercepts it illegitimately. Because of this, always mandate the use of these passwords.

Your company may also want to protect traffic that contains RA requests. E-mail is normally sent in unencrypted form on the network. This means that the URL that is sent in the e-mail invitation is available for easy interception while it is in transit on the network. Likewise, a simple XML format is used for the invitation file. A simple pattern match could be used when monitoring the network to detect and automatically save this information to an unauthorized system while it is being sent across the network. If the e-mail or file invitations do not have passwords, they can be used immediately when they are captured in this way. Even if a password is specified, there is no limit to the number of times requests like these can be used for connection. A brute force attack could be used to attempt to break the password and successfully establish a session. For this reason, it is important that your remote assistance policy also specify a short expiration time for the invitation. Once expired, no connections are possible with it. A shorter time reduces the chances of success using a brute force attack. And if no password is specified, at least the open window for misuse of the invitation is shorter.

You should also educate your users on when it is appropriate to accept RA requests. As mentioned previously, a request saved to a file is stored in a standard XML file. These can easily be modified to perform malicious actions when run by a user on a local system. The e-mail request contains a URL to click and can also be altered. In this case it may take the user to a page that performs malicious actions on their local system, or requires the download and installation of an unauthorized ActiveX control that is designed to appear legitimate to the user. Even an unsolicited request received through Windows messaging has security worries.

The best option is to maintain a tight policy that asks users to reject RA invitations in all but a few instances. What is acceptable will relate specifically to your company. Some organizations allow acceptance only from immediate co-workers and known help desk staff. Others are more liberal and allow invitations to be accepted from any verifiable employee within the company. The most important rule is to not allow connections from outside of the organization. Again, this can be further prevented by the use of firewall rules.

EXAM
70-293
OBJECTIVE
3
5.4.2

Planning for Remote Administration by Using Terminal Services

Most of what is new in Windows 2003 Terminal Services relates to remote administration. Microsoft has really listened to customer feedback from previous versions of the operating system, and has created some major improvements in this area. The test objectives focus on two major Terminal Services components, Remote Desktop for Administration and Remote Assistance. Although a predecessor to Remote Desktop for Administration (Terminal Services in Remote Administration mode) existed in Windows 2000, it has received many changes in the current release. RA is a new component for Microsoft's server operating systems. It was initially released with Windows XP.

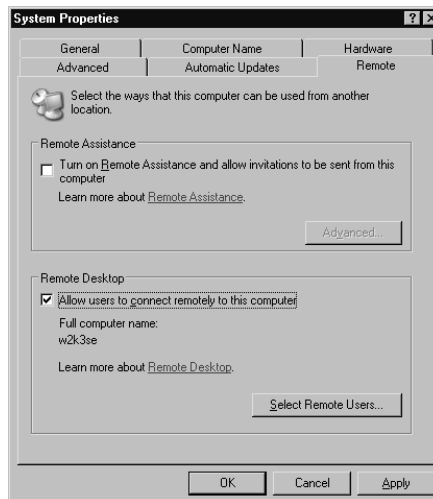
Using Remote Desktop for Administration

As we've mentioned, no installation is necessary for the Remote Desktop for Administration component of Terminal Services. It is installed with the operating system by default. However, for security purposes it is not enabled by default. After it is enabled, members of the administrators group can connect and use it by default. Non-administrators must be specifically granted access. Let's take a look at how to enable and configure this critical component.

Configuring RDA

To configure Remote Desktop for Administration, select **Start | Control Panel | System** and click the **Remote** tab. To enable the feature, simply check the box next to **Allow users to connect remotely to this computer** located in the **Remote Desktop** section of the tab, as shown in Figure 7.30.

Figure 7.30 The System Properties Window



Setting Up Authentication

When RDA is enabled, any user accounts that are members of the Administrators built-in group on the server will be allowed to establish a remote session. However, other accounts must be explicitly approved for access. There are two different ways this can be accomplished. The first is to simply add any accounts that require access to the Remote Desktop Users group on the server. To grant a user access using this method, perform the following steps:

1. Open **Computer Management** and expand the **Systems Tools, Local Users and Groups**, and **Groups** nodes in the console tree on the left side of the utility.
2. Right-click the **Remote Desktop Users** group.
3. Select **Add to Group** from the context menu, and then click the **Add** button.
4. Type (or search for and select) the account name of the user to whom you wish to grant access.
5. Click the **OK** button.

The second, simpler way to access the Remote Desktop Users group and grant access is to use an option provided in the **Remote** tab in the **System** properties located in **Control Panel**. To use this method, perform the following steps:

1. In the Remote Desktop section of the **Remote** tab, click the **Select Remote Users...** button.
2. In the **Remote Desktop Users** dialog box that appears, click the **Add** button.
3. Type (or search for and select) the account name of the user requiring access. (See Chapter 4, *Managing User, Group, and Computer Accounts*, if you need additional information on group management and how to add users to groups in Windows 2003).
4. Click the **OK** button.

Advantages of RDA Over Other Remote Administration Methods

Windows Server 2003 includes many ways to remotely administer your servers. You can install server administration tools (including Active Directory Users and Computers, Active Directory Sites and Services, Active Directory Domains and Trusts, and many others) on a client computer. You can use the Computer Management console on one computer on the network to connect to and manage another. You can use command line tools to connect to and manage computers across the network. What are the advantages of using Remote Desktop for Administration?

Many administrators prefer RDA because you are able to see and use the entire server desktop, exactly as if you were sitting there at the console. You can do things such as pro-

mote or demote a domain controller or defrag the server's disk, install applications, run a backup job, or even upgrade the operating system. You can change configurations, such as Control Panel settings, that are difficult or impossible to configure by other remote methods. You can control the server from a computer on which you would not want to install the administrative tools. With the Remote Desktop Web Connection, you don't even have to have RDC or the Terminal Services client installed on the computer from which you initiate a terminal session; only MSIE v5.0 or later is required. Because of the efficiency of the latest version of RDP, performance over the LAN is almost as fast as if you were physically sitting at the server.

Remote Desktop Security Issues

When enabled, Remote Desktop for Administration opens port 3389 and listens for connection requests. This port is a significant target and is often sought during port scans. Most open ports link to applications that must be attacked in complex ways to permit administrator level access to a computer. This service is designed to actually provide it, which makes it a prime target for attackers. There are several best practices that you should follow to maximize the security of this component.

Remember, with the exception of administrators, users must be authorized to connect using Remote Desktop for Administration. This is accomplished by adding a user's account to the Remote Desktop Users group using one of the methods previously mentioned. If a user does not require this access, his or her account should never be a member of this group. You should control membership in this group through Group Policy or review it manually on a regular basis.

It is important to enforce strong security precautions on all accounts that are enabled to connect using Remote Desktop for Administration. Strong passwords and the use of account lock out are essential to make it difficult for an attacker to successfully use a brute force attack to gain system access. Administrators should be required to log on using a standard user account and perform administrative duties in the session using the **Run as...** feature. This will ensure maximum security of the administrator credentials, minimal damage to the Windows Server 2003 computer if the session is hijacked, and make it more difficult to accidentally install Trojans and other malicious code.

All users should be required to use the most recent client available for their platform. This will ensure that the latest security features are available to them. It should be standard policy to check frequently for software updates to both client and server components, because these may contain critical security fixes. In addition, users should be discouraged from storing their log-on credentials in the properties of the client. This enables anyone with physical access to the user's machine to establish a session. It also stores sensitive information such as the user's username and domain in a clear text file with an RDP extension in the user's My Documents folder.

Finally, denial of service is a significant possibility when using Remote Desktop for Administration because it enables only two sessions to exist on the server. Both active and disconnected sessions count. So, if your company has three administrators and two of them

leave disconnected sessions, the third will not be able to connect using Terminal Services until one of the existing sessions has been terminated. The solution to this may appear to be setting the time out settings so that sessions are reset shortly after they enter the disconnected state. However, this can cause serious problems.

An administrator may establish a session, begin an installation process and then disconnect to enable the installation to finish unmonitored. The previous settings would terminate the session, including the installation routine it was running, with potentially disastrous effects for the server. Special circumstances like these must be taken into account when configuring your policies. Because session timeout settings can be set at the user property level, Microsoft recommends the use of a special shared administrative account for circumstances like this. The strategy applies a timeout for disconnected sessions that are started by every user account except the shared account, which has no timeout settings applied. In this way, there should always be one connection available to a server, even though the second allowed connection is being consumed by a session involving the shared administrative account.

Summary of Exam Objectives

Planning a remote access strategy involves analyzing the needs of the organization, the needs of individual users, and other factors. You should also consider which of the remote access types you need to support:

- Dial-in remote access
- VPN (virtual private network)
- Wireless access

Dial-in access using modems is the traditional type of remote access and is still useful. If you will enable dial-in access, there are a number of factors to plan for. These include whether the RRAS server will assign IP addresses using a static address pool, using DHCP, or using automatic private addressing. You will also need to consider number of incoming ports you will need and whether to manage access by user or using remote access policies.

VPN access uses a client's Internet connection and the server's Internet connection to create a virtual connection, or tunnel, and provide for remote access. A VPN uses one or more VPN protocols to create the tunnels and manage encryption. The VPN tunneling protocols are as follows:

- **PPTP (point-to-point tunneling protocol):** A protocol based on PPP. Uses MPPE for encryption.
- **L2TP (Layer 2 tunneling protocol):** A newer protocol that provides for tunneling and takes advantage of IPSec (IP Security) for encryption. L2TP supports data integrity and sender authentication, unlike PPTP, but requires a public key infrastructure and computer certificates for clients and servers.

Wireless remote access uses the 802.11 standard. A WAP provides access to a number of clients and connects to the LAN. WAPs can use IAS (RADIUS) to provide enhanced security and centralized authentication.

Your plan for the security of a remote access solution should consider the functional levels of domains and the features they support, the authentication methods and encryption levels you will enable, and whether you will support advanced features such as callback security, managed connections, and smart cards.

Remote access policies can be used to grant or deny remote access based on a number of criteria. Each remote access policy includes a profile, which can control what the connection allows after it is established. A profile also includes settings, such as maximum session time and idle timeout, to control the length of remote sessions.

Exam Objectives Fast Track

Planning the Remote Access Strategy

- ☑ Dial-in access requires a modem or ISDN port for each user and is limited in bandwidth, but provides a secure connection without encryption.
- ☑ VPN access can use existing Internet connections but risks sending data (although encrypted) over the public Internet.
- ☑ Wireless remote access uses a wireless access point (WAP) and is usually limited to short ranges.

Addressing Dial-In Access Design Considerations

- ☑ Dial-in clients negotiate with PPP and are issued an IP address. The RRAS server can obtain addresses from a static pool, a DHCP server, or APIPA.
- ☑ You need to determine the number of ports you will need and the bandwidth they will use to plan for dial-in access.
- ☑ Multilink is a system that combines two or more dial-up connections into a single faster connection. It is often used with ISDN.

Addressing VPN Design Considerations

- ☑ PPTP is supported by Windows 95 and later; L2TP is supported by Windows 2000, Windows XP, and Windows Server 2003 only.
- ☑ L2TP supports data integrity and sender authentication; PPTP does not.
- ☑ L2TP requires a public-key infrastructure.
- ☑ L2TP requires machine certificates for each client and VPN server.

Addressing Wireless Remote Access Design Considerations

- ☑ Like other connection types, wireless access can be managed using a remote access policy.
- ☑ A network can support any number of WAPs.
- ☑ RADIUS authentication requires an IAS server configured with the WAPs as clients, and the WAPs configured for RADIUS authentication.

Planning Remote Access Security

- ☑ Windows 2000 mixed domains support Windows NT 4.0 domain controllers and limited security features. Windows 2000 Native and Windows Server 2003 domains support all the Active Directory security features. Windows Server 2003 Interim domains support Windows Server 2003 and Windows NT 4.0 domain controllers.
- ☑ You can raise a domain's functional level, but you cannot lower it.
- ☑ MS-CHAP v2 and EAP are considered the most secure authentication methods.
- ☑ Encryption levels range from no encryption to 168-bit triple DES encryption.

Creating Remote Access Policies

- ☑ Remote Access Policies determine which users can connect remotely and the connection methods they can use.
- ☑ Remote Access Profiles provide further restrictions after the connection is established. Each policy contains exactly one profile.
- ☑ To authorize access by user, use the user's Dial-in properties.
- ☑ To authorize access by group, use the condition in a remote access policy.

Creating a Plan to offer Remote Assistance to Client Computers

- ☑ Remote Assistance is really a tool for end users and you are unlikely to use it for remote server management. You should, however, be aware that Remote Assistance invitations can be sent from a Windows Server 2003 computer, and you should know how to turn off Remote Assistance.
- ☑ End-users can use Remote Assistance to invite another person to view or take control of their desktops.
- ☑ You can use Group Policy to enable your support staff to proactively offer Remote Assistance to end users

Planning for Remote Administration by using Terminal Services

- ☑ Remote Desktop for Administration enables up to two administrators to remotely connect to the server simultaneously, each in his or her own session, to perform administrative tasks.
- ☑ Remote Assistance enables a user, called the Novice, to request help from someone more knowledgeable, called the Expert. The Expert is able to view and interact with the Novice's desktop remotely (if permission is granted by the Novice).
- ☑ Though installed with the operating system, both Remote Desktop for Administration and Remote Assistance must be enabled manually after installation before they can be used.

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: Are dedicated WAN links completely unneeded now that VPN access is common?

A: No, in fact WAN links such as T-1 are now available at lower prices than ever. When you need high bandwidth and don't want to use a public network, these can still be the best option.

Q: Is it better to manage remote access by user or by policy?

A: Either method works. In a small network, managing by user is the simplest approach. For larger networks, this becomes cumbersome and policies are much more convenient.

Q: How many IAS servers should be used for authentication?

A: Although one IAS server provides centralized authentication, it can go down and prevent large numbers of users from accessing the network. For this reason, you may wish to configure a backup server. IAS supports replication between servers.

Q: Can I use L2TP without IPSec?

A: Yes, L2TP does not require IPSec. However, an L2TP connection without IPSec is unencrypted, so it no longer qualifies as a virtual *private* network.

Q: Do I need to choose a single wireless standard?

A: Not necessarily; many products are available that support multiple standards. Most 802.11g equipment supports the 802.11b standard, and many devices add 802.11a support as well. Using WAPs that support multiple standards is the best way to support all clients.

Q: Why is it so difficult to secure a wireless network? WPA is already considered insecure, and it's not even a standard yet.

A: Security by encryption is simply a complex issue and is usually a race between security vendors and hackers. Wired networks have the same issues, but the fact that wireless signals often extend outside your buildings makes exploits much more common.

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 the Remote Access Strategy

1. You are planning a remote access server and need to enable access for several employees. All the employees are in the same city. The company LAN is not currently connected to the Internet, and your security policy specifies that Internet connections should be avoided. Which of the following is the best choice for the remote access solution?
 - A. Dial-in access
 - B. VPN access
 - C. Wireless access
 - D. Dedicated WAN links
2. You are configuring a remote access server on a Windows Server 2003 computer. The same server is acting as a domain controller and DHCP server, assigning IP addresses to clients. Which of the following is the simplest method of assigning IP addresses for remote clients?
 - A. Manually configure each client with an IP address.
 - B. Configure the RRAS server to use DHCP.
 - C. Configure a static address pool.
 - D. Use APIPA.

Addressing Dial-In Access Design Considerations

3. You are configuring a dial-in remote access server on a Windows Server 2003 computer. Employees will use remote access while traveling. You have ten employees with laptops who will require access to the server, but typically only one is traveling at a time. A telecommuting employee will also require access for eight hours a day. How many modems would be the minimum to reliably serve these users?
 - A. 1
 - B. 11
 - C. 2
 - D. 3

4. You have several users who dial in to a remote access server using multilink connections, combining two modems into a single link. Although this provides a higher bandwidth to the users, you find the server runs out of modem lines frequently, and most users are not using their connections to their full potential. Which of the following is a solution to this issue?
 - A. Disable multilink connections.
 - B. Set the maximum number of multilink ports to one.
 - C. Use VPN instead of dial-in access.
 - D. Enable Bandwidth Allocation Protocol (BAP).

Addressing VPN Design Considerations

5. You are configuring a Windows XP client machine to access a VPN server that supports L2TP over IPSec. You need to obtain a computer certificate for the client and wish to do so from the client machine. A CA is present on the local network. Which application can you use to request a certificate?
 - A. A Web browser
 - B. The Certificates MMC snap-in
 - C. The Certification Authority MMC snap-in
 - D. Connection Manager

6. You have configured a VPN server running Windows Server 2003 and RRAS. Most clients are able to access the server, but clients running Windows 98 are reporting that they are unable to connect. Which of the following is most likely the cause of this problem?
 - A. Computer certificates are not installed.
 - B. L2TP is not enabled on the server.

- C. PPTP is not enabled on the server.
- D. Windows 98 does not support VPN client access.

Addressing Wireless Remote Access Design Considerations

- 7. You are setting up wireless access to the network with two WAPs. You want to use a centralized authentication source for both access points. You have an existing IAS server on the network. Which of the following tasks are necessary to support wireless access? (Choose all that apply.)
 - A. Create a remote access policy.
 - B. Configure the WAPs to use RADIUS authentication.
 - C. Install a RADIUS server.
 - D. Add the WAPs as clients in the IAS server's configuration.
- 8. You have configured a WAP using the EAP-TLS protocol. The WAP is connected to a LAN with a Windows Server 2003 server. Which of the following additional tasks may be necessary to ensure that wireless clients can connect? (Choose all that apply.)
 - A. Enable PPP authentication.
 - B. Issue computer certificates to clients.
 - C. Issue user certificates or smart cards to users.
 - D. Install and configure IAS.

Planning Remote Access Security

- 9. You are planning security for your network and have determined that the domain functional level is Windows 2000 Mixed mode. You have a combination of Windows Server 2003 and Windows 2000 Server domain controllers. Which of the following actions may be necessary to enable all of Windows Server 2003's security features? (Choose all that apply.)
 - A. Eliminate or upgrade the Windows 2000 Server domain controllers.
 - B. Eliminate all Windows 2000 clients.
 - C. Raise the functional level to Windows Server 2003.
 - D. Raise the functional level to Windows Server 2003 Interim.

10. You have a network with two Windows Server 2003 servers. You have raised the domain function level to Windows Server 2003. You need to install an additional domain controller and are considering an existing Windows 2000 Server. Which of the following tasks is necessary before using this machine as a domain controller?
- A. Lower the function level to Windows 2000 Mixed mode.
 - B. Lower the function level to Windows Server 2003 Interim.
 - C. Upgrade the Windows 2000 Server to Windows Server 2003.
 - D. Demote the existing domain controller to a member server.

Creating Remote Access Policies

11. You have an RRAS server and have configured two remote access policies. The first policy on the list allows access for all members of the Power Users group. The second policy on the list denies access to clients that connect during evening hours. After testing your configuration, you determine that clients in the Power Users group are able to connect at any time. Which of the following actions would correct this problem?
- A. Delete the first policy in the list.
 - B. Change user account properties to deny remote access.
 - C. Change the order of the policies.
 - D. Install an IAS server.
12. You are operating a remote access server and currently allow VPN access and dial-in access. You have decided to disallow dial-in access after configuring all the clients for VPN access. Which of the following attributes can you check in a remote access policy to deny access to modem users?
- A. Authentication-Type
 - B. NAS-Port-Type
 - C. Framed-Protocol
 - D. NAS-Identifier

Creating a Plan to offer Remote Assistance to Client Computers

13. One of your users is having problems getting a productivity application to work correctly. You suspect that he is performing the steps involved in using the application

incorrectly, but the application interface is complex and it's difficult for you to explain, over the phone, what he needs to do. The user is running Windows XP, and you want to connect to his PC and show him how to perform the task in question so that he can actually see you go through the steps. How would you arrange to do this?

- A. Send the user a Remote Assistance Request.
 - B. Get the user to send a Remote Assistance Invitation.
 - C. Connect to the user's PC using Remote Desktop.
 - D. Connect to the user's PC using the Web Interface for Remote Administration.
14. You are attempting to describe the remote assistance process to a co-worker. The co-worker asks what the correct terms are for the person requesting assistance and the person providing assistance so that he can look them up in Windows Help. Which of the following do you reply with? (Select two.)
- A. Administrator
 - B. Novice
 - C. Expert
 - D. End user

Planning for Remote Administration by using Terminal Services

15. You are attempting to describe the remote assistance process to a co-worker. The co-worker asks what the correct terms are for the person requesting assistance and the person providing assistance so that he can look them up in Windows Help. Which of the following do you reply with? (Select two.)
- A. Administrator
 - B. Novice
 - C. Expert
 - D. End user

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

MCSE 70-293

Planning, Implementing, and Maintaining a High- Availability Strategy

Exam Objectives in this Chapter:

- 4 Planning, Implementing, and Maintaining Server Availability
 - 4.1 Plan services for high availability.
 - 4.2 Identify system bottlenecks, including memory, processor, disk, and network related bottlenecks.
 - 4.2.1 Identify system bottlenecks by using System Monitor.
 - 4.5 Plan a backup and recovery strategy.
 - 4.5.1 Identify appropriate backup types. Methods include full, incremental, and differential.
 - 4.5.2 Plan a backup strategy that uses volume shadow copy.
 - 4.5.3 Plan system recovery that uses Automated System Recovery (ASR).
- 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 for good reason. Ensuring that the network's resources are available to users when they need them is an important part of the network administrator's job. Downtime—whether caused by a disk failure, a performance slowdown, data loss due to an attack, or the loss of an entire server due to a natural disaster such as fire or flood—cuts into worker productivity and impacts the business's bottom line or the organization's ability to accomplish its goals.

In this chapter, we look at the concept of high availability and how it can be attained. We'll provide an overview of performance bottlenecks and what cause them, and show you how to identify such common system bottlenecks as memory, processor, disk, and network components. We'll walk you through the steps of using the System Monitor utility to track server performance and show you how to use Event Viewer and service logs to monitor server issues, as well.

Next, we show you how to plan a backup and recovery strategy. We'll introduce you to the Windows Backup Utility and ensure that you understand the differences between full, incremental, and differential backups. We'll also discuss the use of the Volume Shadow Copy feature as a backup option. You'll learn how to decide what information should be backed up. We'll also show you how to back up user data, System State data, the Dynamic Host Configuration Protocol (DHCP) database, Windows Internet Name Service (WINS) database, Domain Name System (DNS) database, cluster disk signatures, and partition layouts. We'll walk you through the process of using the Windows Backup administrative tool, including the Backup and Restore Wizard feature and the Advanced Mode feature. We'll also discuss the use of command-line tools. Then we'll talk about how to select your backup media, and you'll learn about scheduling backups and how to restore data from backup when necessary.

We'll address how to plan for system recovery using the Automated System Recovery (ASR) feature. You'll learn about system services, how to make an ASR backup, and how to do an ASR restore. We'll explain how ASR works and discuss alternatives to ASR such as the Safe Mode and Last Known Good boot options. Finally, we'll discuss the importance of planning for fault tolerance, including solutions aimed at providing fault tolerance for local network connectivity, Internet connectivity, data on disk, and mission-critical servers.

Understanding Performance Bottlenecks

All system administrators want the systems they install to run perfectly out of the box, all the time. We have all wanted to be able to safely turn off our pagers and cell phones. Our servers should run reliably, quickly, and without interruption, right? Well, if that were the case, we would all be terminally bored or changing careers.

In this chapter, we will examine the conditions and tools for monitoring and ensuring a system's smooth operation. In Chapter 9, we will discuss some of the sophisti-

EXAM
70-293

OBJECTIVE

4
4.1
4.2

cated tools Microsoft has included with Windows Server 2003 for enhancing system uptime and supporting high-volume use.

Identifying System Bottlenecks

For the most part, a Windows Server 2003 system does run well in its default configuration, and, if designed and maintained properly, operates with a minimum of administrative overhead. However, as a general-purpose operating system, it can often be tuned to perform better when used for certain tasks.

The main hardware resources of any computer system are memory, processor, disk storage, and communications or network components. Different applications and circumstances use these resources in different combinations, often taxing one resource more than another. If multiple applications are run on a system, it is often possible to reach the limit of a resource and suffer slow response time, unreliable services, or miss a result. We will take a look at each of these resources, discuss some of the common issues related to them, and consider some of the management options available.

Memory

Memory, random access memory, or simply *RAM* is the working space of the operating system and applications. Its contents are volatile and always in demand. In 1990, a computer system that had a memory capacity of 16MB was very high end. Today, it is possible to purchase systems that support 32GB of memory or more.

RAM is most often the single resource that becomes a bottleneck. A common cause of slow performance is insufficient physical memory. When purchasing new hardware, it is not wise to skimp on memory. The minimum recommended amount of memory for running Windows Server 2003 is 128MB (512MB for Datacenter Edition). These are *very* conservative numbers. Even Microsoft recommends at least 256MB. If you have the ability, double (or more) the amounts to, and you will be happy you did. The short rule with memory is this: more is better.

The Windows operating system controls the access to and allocation of memory and performs “housekeeping tasks” when needed. Applications request memory from the operating system, which allocates memory to the application. When an application no longer needs memory, the application is supposed to release the memory back to the operating system. An application that does not properly release memory can slowly drain a system of available free memory, and overall performance will suffer. This is referred to as a *memory leak*.

Another performance factor related to memory is the use of *virtual memory (VM)* or *paging*. Virtual memory is a method of increasing the amount of memory in a system by using a file on the hard drive called a *page file*. The apparent size of memory is increased without increasing the physical RAM in the system, hence the term *virtual*. Access to hard drives, even on the fastest disk subsystems, is dozens or hundreds of times slower than access to RAM. When the operating system needs more RAM than is available, it copies the least

recently used pages of memory to the page file, and then reassigns those pages of RAM to the application that requested it. The next time a memory request occurs, the operating system may need to reallocate more pages in RAM or retrieve pages from the page file. This paging process can slow even the fastest system.

Page File Tricks

Configuring the Windows page file is not as simple as it may seem. Several factors can affect the performance of the page file and therefore overall system performance. The page file is heavily accessed. Placing it on separate drive and control channel from the operating system and/or applications greatly reduces competition for disk access (known as *contention*).

Spreading the page file across several drives and control channels can boost performance as well, provided that the pieces are not stored on high-contention drives. The use of high-performance drives will also improve performance. In general, SCSI and Fibre Channel provide greater performance than IDE drives.

Keeping the page file defragmented will also help its performance. A single contiguous block of disk space provides better performance by reducing drive read/write head movement. Before adding a new page file to a drive, use the Windows Disk Defragmenter utility to defragment the disk. When you add the new page file to the disk, it will automatically allocate as much contiguous space as it can.

Finally, consider using the **Custom Size** option when configuring virtual memory, and make the **Initial Size** and **Maximum Size** options the same value. This will consume more disk space but will avoid any incidents of *expansion delay*, which occurs when the system must increase the size of the page file. This will also help keep the page file defragmented.

Finally, remember that even the most optimally tuned page file configuration will not make up for insufficient physical RAM.

Tuning memory is often as simple as adding more memory, reducing the number of applications running (including applications that run in the System Tray), or stopping unnecessary services. However, there is an advanced memory-tuning technique that can be applied if the application supports it. Part of the Enterprise Memory Architecture feature of the Enterprise and Datacenter editions of Windows Server 2003 is *4GB tuning* (4GT), also called *application memory tuning*. Using this feature, you can change the amount of RAM addressable by applications from 2GB to 4GB. Your system must have at least 2GB of physical RAM installed, and the application must be written to support the increased memory range. Consult the application documentation or contact your vendor to make this determination.

Processor

If memory is the working space, then the processor is the worker. The central processing unit, or *CPU*, is the “brain” of a computer system and is responsible for, well, processing. It receives data (*input*), performs calculations (*execution*), and reports the result (*output*). The processor is (usually, but not always) responsible for moving data around inside a computer system, transporting it among memory, disk, network, and other devices.

CPUs are commonly described by their type, brand, or model (for example, Pentium 4), and their *clock speed* (for example, 2.0 GHz). In simplest terms, the clock speed is how many times per second the CPU executes an instruction. Generally, the faster the CPU is, the better the computer performs.

The *CPU bus architecture* is another factor when examining performance. Bus architecture is a term used to describe how much data can be moved in and out of the processor at once. It also describes the amount of information that the CPU can process in a single step. A 32-bit CPU (which includes all of Intel’s CPUs from the 80386 through the Pentium 4 and AMD’s CPUs from the Am486 through the Athlon-XP) can use 32 bits wide and access 2^{32} bytes of memory, or 4GB. A 64-bit CPU (Intel’s Itanium series and AMD’s Opteron series) can use 64 bits wide and access (in theory) 2^{64} bytes of memory or 16 exabytes (16 billion gigabytes). No current hardware can support this amount of RAM. Windows Server 2003 supports a maximum of 512GB on Itanium-based hardware with the Datacenter Edition. The point is that 64-bit CPUs can support significantly more memory and run applications that use more of it than 32-bit CPUs, all at a faster speed. Extremely large databases can get a large performance boost on 64-bit systems.

Using multiple CPUs in a computer (called *multiprocessing*) allows a computer system to run more applications at the same time than a single-CPU system, because the workload can be spread among the processors. In effect, this reduces the competition among applications for CPU time. A related programming technology called *multithreading* allows the operating system to run different parts of an application (*threads*) on multiple CPUs at the same time, spreading out the workload. Windows Server 2003 can support up to 64 CPUs, depending on the edition of the operating system in use.

A recent development by Intel is a technology called *hyperthreading*. This feature, introduced in the Xeon and Pentium 4 series of processors, makes a single CPU appear to be two CPUs. Hyperthreading is implemented at the BIOS level and is therefore transparent to the operating system. It typically yields a performance increase of 20 to 30 percent, meaning it is not as efficient as multiple physical CPUs. However, it is included for free on hardware that supports the technology.

One of the downsides of multiple processors is the management of *interrupts*. An interrupt is a hardware or software signal that stops the current flow of instructions in order to handle another event occurring in the system. Disk input/output (I/O), network I/O, keyboard activity, and mouse activity are driven by interrupts. Interrupts are a necessary part of the computer’s operation, but they can impact the performance of a multiprocessor system. The first CPU in a multiprocessor system (processor 0) controls I/O. If an application run-

ning on another CPU requires a lot of I/O, CPU 0 can spend much of its time managing interrupts instead of running an application. Windows Server 2003 manages multiple processors and interrupts quite well, but there is a way to tune the *affinity* of a thread to a specific processor.

Processor affinity is a method of associating an application with a specific CPU. Processor affinity can be used to reduce some of the overhead associated with multiple CPUs and is controlled primarily via the Task Manager utility.

Priority is a mechanism used to prioritize some applications (or threads) over others. Priority can be set when an application is started or can be changed later using the Task Manager utility.

Disk

The *disk* (interchangeably called *hard disk*, *drive*, *storage*, *permanent storage*, *array*, *spindles*, or multiple combinations of these terms) is the permanent storage location of the operating system, applications, and data. A disk is made up of one or more physical *drives*. Disk content remains when power is turned off. Often, a computer system will contain multiple disks configured in *arrays*, which can be useful for increasing performance and/or availability.

Several factors contribute to the performance and reliability of disk:

- The disk controller technology
- The life-expectancy or mean time of the drive(s)
- The way data is arranged on the drive(s)
- The way data is accessed on the drive(s)
- The ratio of drive controllers to the number of drives

Disk Controller Technology

The hard drive itself is a dependent component, meaning it requires some other device to interface with the computer system. The component that provides this interface is the *disk controller*. The disk controller is responsible for converting the request of the operating system into instructions the hard drive can process. The controller also manages the flow of data to and from the drive. A hard drive is always manufactured to work with a specific type of controller.

There are three major types of hard drive interface technologies, each with their own advantages and disadvantages. *ATA* (sometimes called *IDE* or *EIDE*) is the most common interface. Though often found on server-class systems, its primary usage is on workstation-class or midrange systems. Compared to the other technologies, *ATA* drives usually cost the least. The *ATA* interface itself can support a maximum achievable throughput of about 50 Megabytes per second (MBps). *ATA* systems generally support four drives, with two drives on a disk channel. One drive on the channel is configured as *master*, and the other (if pre-

sent) is configured as *slave*. This configuration means that the master drive controls the traffic on the channel as well as the slave drive. This can raise some hardware incompatibility issues, and it is for this reason that you should not mix drives from different manufacturers on the same channel if you can avoid doing so. ATA was designed primarily for use with disk drives, but other ATA devices do exist, such as CD/DVD-ROM drives and tape drives.

The second interface technology commonly used is the *Small Computers System Interface*, commonly called *SCSI* (pronounced “skuzzy”). As a defined technology, SCSI has been around longer than ATA, but it is less prevalent because SCSI devices are generally more expensive than similar ATA devices. SCSI is most commonly found on midrange to high-end server systems. SCSI is a general-purpose interface technology that supports a wide variety of devices: hard drives, tape drives, CD and DVD-ROM drives, scanners, write once/read many (WORM) drives, and more. A SCSI channel can support from 8 to 16 devices, depending on the exact SCSI specification being followed. The current SCSI specification supports 16 devices and a bus speed of 160 MBps. The SCSI bus controller is considered intelligent, meaning that the controller, rather than the system CPU (as is the case with ATA), does the work of managing the channel and the flow of data on the channel. This gives SCSI a performance advantage over ATA but requires SCSI devices to have more advanced circuitry, increasing the cost of SCSI devices.

Fibre Channel (sometimes referred to as *F/C*) is the most recent interface technology to come into widespread use. It is not usually built into a system and requires interface adapters to be installed. Fibre Channel is commonly used to connect servers to large back-end *storage area networks* (SANs). Data rates of 100 MBps, 1000 MBps, and 2000 MBps are possible with Fibre Channel. Fibre Channel is quite expensive, but it supports hundreds of devices per channel. When connected via an SAN, a Fibre Channel configuration can support thousands of devices. Because Fibre Channel can support such a large number of devices, a Fibre Channel configuration can be very complex and difficult to manage. Like SCSI, Fibre Channel is a general-purpose interface technology and supports many devices other than disk drives. Because of its higher cost, however, there are generally fewer devices available for Fibre Channel than for SCSI.

Drive Life Expectancy (MTBF)

Because hard drives contain rapidly rotating moving parts, they are subject to more frequent mechanical failure than most other components in a computer system. Disk drive manufacturers usually predict the rate of failure in terms of mean time between failures (MTBF). This number is often measured in thousands of hours of operation and is an indicator of the statistical reliability of the hard drive device.

It is quite possible to see MTBF ratings of 100,000 hours (11.4 years) or more. But what happens with multiple drives? By understanding that drive failure is a real and predictable event, you can calculate the likelihood of your system experiencing a drive failure. For example, if you were to install two identical hard drives in your system, each with a MTBF of 100,000 hours, when would you (statistically) expect to experience a drive

failure? The answer is sometime within 5.7 years (100,000 MTBF hours/2 drives). This may not seem like too much of an issue until you extend the calculation. It is not uncommon for systems to have six or more drives, depending on the amount of storage needed and the size of the drives involved. Using the same 100,000-hour example with six drives, you could expect to see a failure sometime within 23 months, or just less than two years. If you have more disks, the odds turn against you even further. In short, the more drives you have, the higher the likelihood of experiencing a drive failure.

The mechanism used to control the cumulative effect of this problem is *Redundant Arrays of Independent Disks*, or *RAID*. RAID is a technique of using multiple drives in such a way that data is spread among the drives so that, in the event of failure, the data is available from either a second copy of the data or is re-creatable from a mathematical calculation. A properly configured RAID array counteracts the cumulative effect of MTBF.

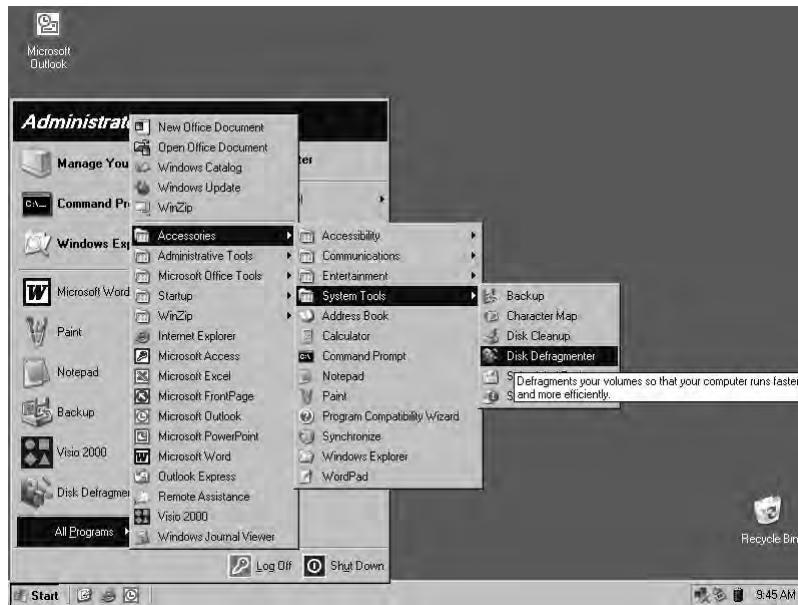
RAID arrays can also be used to increase performance. For example, if you were using five disk drives and spread the data evenly among those drives, you would expect to see it take one-fifth the time to read the data from that array than from a single drive containing the same data. This is referred to as increasing the *spindle count*. (The spindle is the axle that the disks in the disk drive are attached to.) By increasing the number of spindles, you reduce the demand on any single drive for reads or writes.

RAID is configured in one of two ways: in software or in hardware. Software RAID is supported by Windows Server 2003 and can be used to create and manage some types of RAID arrays without the cost of a hardware RAID controller. Hardware RAID requires a hardware RAID controller and usually offers more options for configuring a RAID array, as well as increased performance over software RAID.

Arrangement of Data on Drives

How data is arranged on a drive affects how fast the data is written to a drive or read from a drive. Assume that you are starting with a clean, freshly formatted hard drive. As files are written to the drive, they are assigned to available areas of the disk (called *clusters*) by the file system. If the file is large enough, it will be written to multiple clusters. If a file occupies clusters that are next to each other on disk, the file is said to occupy *contiguous clusters*. Over time, additional files are written to the disk. Later, your first file increases in size. The file system allocates new clusters to the file, some of which it already occupied and some of which are now on a different area of the disk. The file no longer occupies contiguous clusters and is *fragmented*; multiple fragments of the file are spread out around the disk. The effect of fragmentation is to require more time to access the file, thus slowing performance. Fragmentation occurs on every drive, but frequently updated drives are most susceptible to the performance degradation of fragmentation.

Microsoft recognizes that file fragmentation can be a problem and has included a utility with Windows Server 2003 to address this problem. You can use the Disk Defragmenter utility to reduce the total fragmentation on a drive. This utility can be started from the command line or from the Start menu, as shown in Figure 8.1.

Figure 8.1 Starting Disk Defragmenter


Although file fragmentation can have an impact on performance, so can the cure. While running, Disk Defragmenter can have a serious performance impact on a system. It is for this reason that you should use Disk Defragmenter with caution. You can use the Analysis function of Disk Defragmenter to see the fragmentation statistics on a drive without actually performing defragmentation.

New & Noteworthy...

Scheduling Disk Defragmenter

One of the great features missing in Windows NT 4 and limited in Windows 2000 is the Disk Defragmenter utility. It was originally thought that because of the way NTFS functions, file fragmentation would not be an issue. This thinking was proven wrong, however, in the earliest versions of Windows NT, and an appropriate application program interface (API) was built into Windows NT 4 to support defragmentation. Although Microsoft did not include a defragmentation tool in the operating system, a few third-party software publishers produced successful defragmentation products for that operating system.

In Windows 2000, Microsoft included the first version of the Disk Defragmenter utility. This version performed defragmentation adequately, but could run only interactively. This meant that defragmentation needed to be performed by an administrator sitting at the console or via a Terminal Services session. Well, Microsoft has made a nice improvement.

Continued

Disk Defragmenter in Windows Server 2003 can be run from the Start menu or from the command line (Defrag.exe). Although Disk Defragmenter has no built-in scheduling function, if you run it from the command line, you can use the **AT** job-scheduling command to create a scheduled defragmentation run. You no longer need to be logged in to the server to run a defragmentation procedure.

If you've ever needed to manually defragment a lot of servers, you'll appreciate this new capability. If you've never performed a lot of manual defragmentation runs, consider yourself very fortunate to have this feature available now.

The Way Data Is Accessed on Drives

Different applications access data in different ways. For example, Microsoft Exchange Server will write to transaction logs sequentially; the mail store databases may be written to and read from either randomly or sequentially. You should be able to develop a profile of this pattern of reads and writes. This profile will determine the design of the underlying disk system and RAID type used. If the wrong RAID type is used, performance will suffer and/or reliability will be compromised.

If you are running multiple I/O-intensive applications on a system, consider giving each application its own drives and controller. That way, the applications will not conflict with each other for I/O resources.

The Ratio of Drive Controllers to the Number of Drives

The ratio of drive controllers to the number of drives relates to the way data is accessed on drives. As mentioned earlier, a higher number of spindles can yield higher throughput. However, if all the drives are connected to one controller, the controller can potentially become a bottleneck, because all I/O must go through the controller. In systems with multiple drives, it may be beneficial to add more controllers and balance the drives among the controllers. This reduces the load on any one controller and results in improved throughput.

Another potential performance improvement with multiple controllers is the ability of the operating system to perform *split seeks*. If the drives on the controllers are mirrored, the operating system will send read requests to the drive that is able to service the request the quickest.

Network Components

A network is the primary computer-to-computer communications mechanism used in modern computing environments. Windows Server 2003 supports numerous network interface cards (NICs) and multiple network topologies, including Ethernet, Token Ring, Asynchronous Transfer Mode (ATM), and Fiber Data Distributed Interconnect (FDDI), as well as remote-access technologies including dial-up and virtual private Network (VPN) connections.

The primary (and default) communications protocol used by Windows Server 2003 is TCP/IP. The operating system also supports the next generation of the TCP/IP protocol, IP version 6 (IPv6), as well as the AppleTalk, Reliable Multicast, and NWLink protocols. Multiple protocols installed on a system consume additional memory and CPU resources. Reducing the number of protocols in use will improve performance.

On systems that do have multiple protocols loaded, change the binding order of the protocols to the NIC so that the most frequently used protocol is bound first. This will reduce the amount of processing needed for each network packet and improve performance.

When a packet is received by the NIC, it is placed in a memory buffer. The NIC then either generates an interrupt to have the CPU transfer the packet to the main memory of the computer or performs a direct memory access (DMA) transfer and moves the data itself. The method used depends on the hardware involved. The DMA transfer method is much faster and has less impact on the overall system.

The number and/or size of the memory buffers allocated to the NIC can affect performance as well. Some NICs allow you to adjust the number of memory buffers assigned to the card. A larger buffer space allows the NIC to store more packets, reducing the number of interrupts the card must generate by allowing larger, less frequent data transfers. Conversely, reducing the number of buffers can increase the amount of interrupts generated by the NIC, impacting system performance. Some NICs allow you to adjust the buffers for transmitting and receiving packets independently, giving you more flexibility in your configuration. The trade-off of increasing communication buffers may be a reduction in the amount of memory available to applications in the system.

One feature of Windows 2000 and Windows Server 2003 is support for IP Security (IPSec). This feature allows the securing of data transmitted over IP networks through the use of cryptography. IPSec is computer-intensive, meaning that significant amounts of CPU overhead are incurred with its use. Some NICs made in recent years support the offloading of the IPSec calculations to a highly optimized processor on the NIC. This can greatly reduce the amount of CPU time needed, as well as improve communications performance. NICs that support the offloading of IPSec are not much more expensive than regular NICs and should be strongly considered if you are using IPSec.

Another consideration for network performance is the network topology. Although it is not specific to the Windows operating system, network topology can greatly affect communications performance. If the traffic to your server must travel through routers that convert large incoming packets into multiple smaller packets, your server must do more work to reassemble the original packets before your applications can use the data. For example, if a client system is connected to a Token Ring network (which commonly uses a packet size of 4192 bytes or larger) and your server is connected via Ethernet (with a packet size of 1514 bytes), an intermediary router will “chop” the original packet into three or more packets for transmission on the Ethernet network segment. Your server must then reassemble these packets before passing them to applications. Reducing the number of

topologies and/or routers on your network can improve performance by reducing packet conversions and reassembly.

Also, when using Ethernet, consider using switches instead of hubs. Switches are more expensive but allow higher communication rates and also permit more than one device to communicate at the same time. Hubs are cheaper but allow only one computer to be communicating at any given time. Switches are also not susceptible to Ethernet collisions, whereas hubs are at the mercy of collisions.

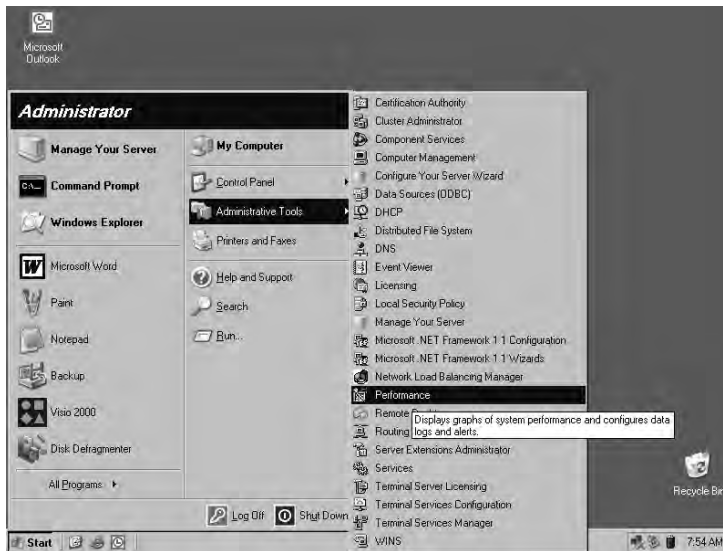
Another topology-related configuration is the duplex setting of the NICs. Primarily an issue with Ethernet, *duplex* describes how data is transmitted. A full-duplex communication link allows the simultaneous transmission and reception of data. Full duplex is the desired setting for servers, because servers normally need to transmit and receive at the same time. Full duplex typically requires switches. Half-duplex communication is the bi-directional communication of data but not at the same time. When transmitting data, receiving data is not possible and vice versa. Half duplex is often acceptable for workstations but should be avoided on servers.

EXAM
70-293
OBJECTIVE
4.2.1

Using the System Monitor Tool to Monitor Servers

Windows Server 2003 includes tools for monitoring the performance of your server. System Monitor is one such tool. System Monitor is an ActiveX control snap-in that is available as part of the Performance administrative tool. You can start it from the Start menu, as shown in Figure 8.2.

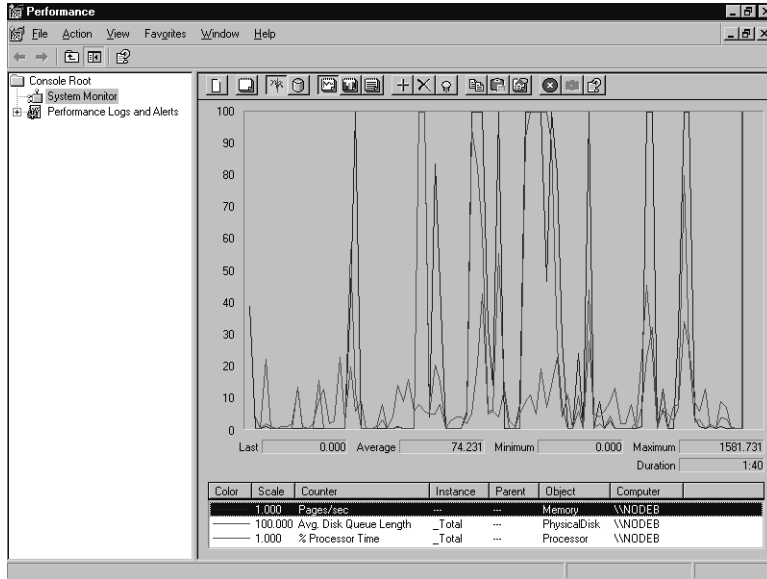
Figure 8.2 Starting the Performance Administrative Tool



System Monitor works by collecting information from *counters* built in to the operating system. Counters are features of the operating system (as well as some utilities and applica-

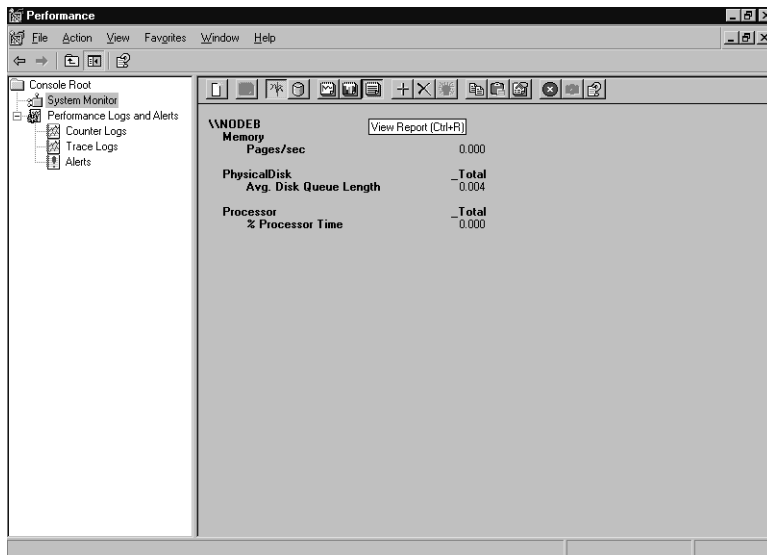
tions) that count specific events occurring in the system (like the number of disk writes each second or the percentage of disk space in use by files). The graphical view, shown in Figure 8.3, graphs the counter statistics.

Figure 8.3 System Monitor, Graphical View with Default Counters



You can also see the counter statistics in System Monitor’s report view, as shown in Figure 8.4.

Figure 8.4 System Monitor, Report View with Default Counters



You can access the data collected by counters via System Monitor, other utilities, or third-party applications. This data provides you with an understanding of what is occurring on your system at the moment or over time. Using this information, you can tune your system's operation to best suit your needs, determine if components are being overutilized, plan for expanding or replacing your system, or perform troubleshooting. Some of the most frequently used counters are shown in Table 8.1.

Table 8.1 Commonly Referenced Performance Counters

Resource	Performance Object:Counter	Recommended Threshold	Comments
Disk	Logical Disk: % Free Space	15%	Percentage of unused disk space. This value may be reduced on larger disks, depending on your preferences.
	Physical Disk: % Disk Time Logical Disk: % Disk Time	90%	If you're using a hardware controller, try increasing the controller cache size to improve read/write performance. The Physical Disk counter is not always available or may be unreliable on clustered disks.
	Physical Disk: Disk Reads/sec Physical Disk: Disk Writes/sec	Varies by disk technology	The rate of read or write operations per second. Ultra SCSI should handle 50 to 70. I/O type (random/sequential) and RAID structure will affect this greatly.
	Physical Disk: Avg. Disk Queue Length	Total # of spindles plus 2	The number of disk requests waiting to occur. It is used to determine if your disk system can keep up with I/O requests.
Memory	Memory: Available Bytes	Varies	The amount of free physical RAM available for allocation to a process or the system.
	Memory: Pages/sec	Varies	Number of pages read from or written to disk per second. Includes cache requests and swapped executable code requests.
Paging File	Paging File: % Usage	Greater than 70%	The percentage of the page file currently in use. This counter and the two Memory counters are linked. Low Available Bytes and high Pages/sec indicate a need for more physical memory.

Continued

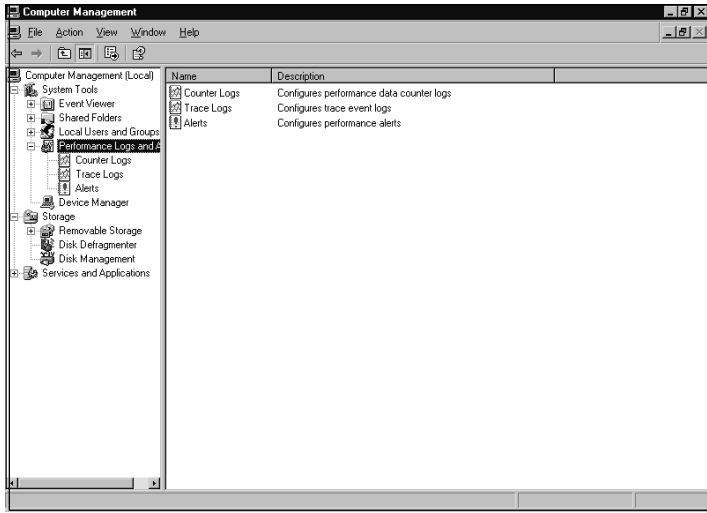
Table 8.1 Commonly Referenced Performance Counters

Resource	Performance Object:Counter	Recommended Threshold	Comments
Processor	Processor: % Processor Time	85%	Indicates what percentage of time the processor was not idle. If high, determine if a single process is consuming the CPU. Consider upgrading the CPU speed or adding processors.
Server Work Queues	Processor: Interrupts/sec	Start at 1000 on single CPU; 5000 multiple	Indicates the number of hardware interrupts generated by the components of the system (network cards, disk controllers, CPUs and so on) per second. A sudden increase can indicate conflicts in hardware. Multiprocessor systems normally experience higher interrupt rates.
	Server Work Queues: Queue Length	4	Indicates the number of requests waiting for service by the processor. A number higher than the threshold may indicate a processor bottleneck. Observe this value over time.
System	System: Processor Queue Length	Less than 10 per processor	The number of process threads awaiting execution. This counter is mainly relevant on multiprocessor systems. A high value may indicate a processor bottleneck. Observe this value over time.

Before you attempt to troubleshoot performance issues, you should perform a process called *baselining*. Baselining is the process of determining what the normal operating parameters are for your system. You develop a baseline by collecting data on your system in its initial state and at regular intervals over time. Save this collected data and store it in a database. You can then compare this historical data to current performance statistics to determine if your system's behavior is changing slowly over time. Momentary spikes on the charts are normal, and you should not be overly concerned about them. Sustained highs can also be normal, depending on the activity occurring in a system, or they can signal a problem. Proper baselining will allow you to know when a sustained high is detrimental.

The Heisenberg Principal of physics (greatly paraphrased) states that you cannot observe the activity of something without altering its behavior. The same is true with performance monitoring. The collection of statistics requires computer resources. Collect only the counters you specifically need. If the system you wish to examine is extremely busy, consider using the noninteractive Performance Logs and Alerts function, shown in Figure 8.5, rather than System Monitor. Collecting the counters will still incur overhead, but not as much.

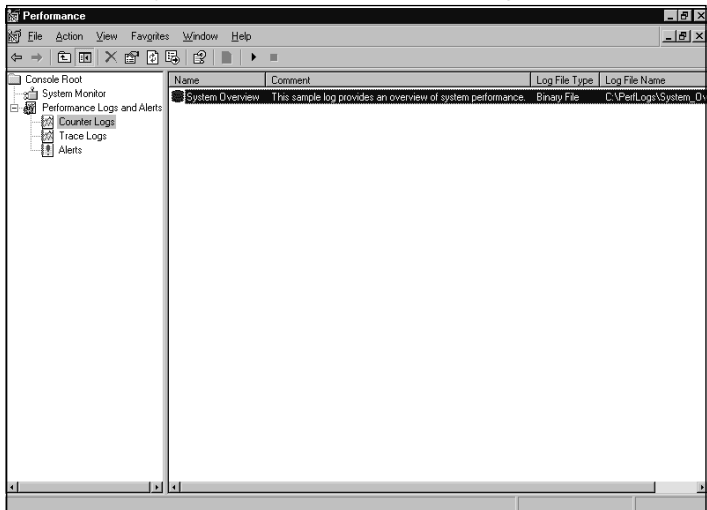
Figure 8.5 Performance Logs and Alerts, Accessed from Computer Management



You might also consider lowering the update interval or collecting data from the system on an hourly or daily basis, rather than continuously. If you're monitoring disk counters, store the logs on a different disk (and, if possible, on a different controller channel) than the one you are monitoring. This will help avoid skewing the data. You can also save the collected statistics into a file or database, and then load and review them later (configured on the **Source** tab of the **System Monitors Properties** dialog box, shown later in Figure 8.12).

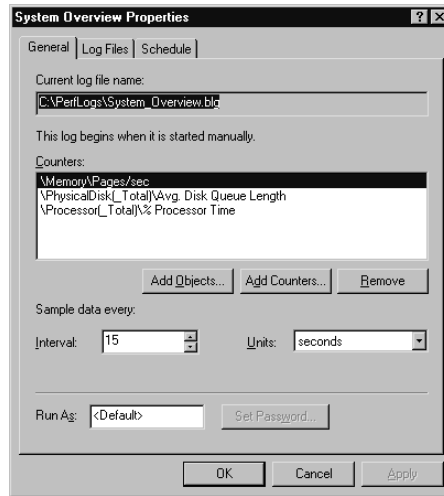
A counter log called System Overview, shown in Figure 8.6, is provided by Microsoft as an example. This log is configured to collect the same default counters as System Monitor at 15-second intervals.

Figure 8.6 The Sample System Overview Counter Log



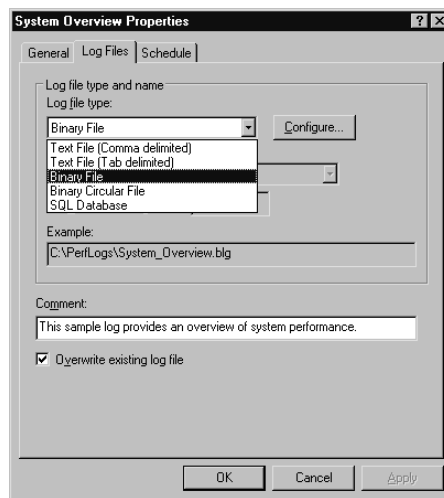
You can view the properties of the System Overview log by right-clicking it and selecting **Properties** from the context menu. Figure 8.7 shows the dialog box that appears.

Figure 8.7 Properties of the System Overview Sample Log



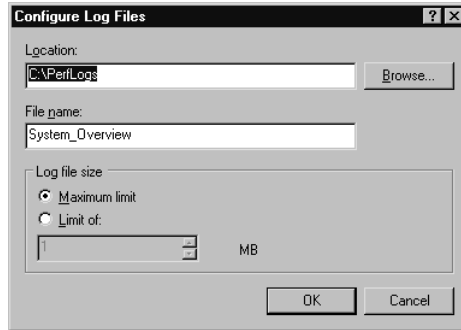
The **Log Files** tab, shown in Figure 8.8, specifies the type (format) of the log file, its location, and the comment assigned to the log. The drop-down menu for the **Log File Type** option (shown in Figure 8.8) lists the formats available for log files. The binary format is compact and efficient. The text formats require more space to store the log, but they can be read by other applications like Microsoft Word and Excel. The binary circular file format overwrites itself when it reaches its maximum size, potentially saving disk space.

Figure 8.8 Properties of the System Overview Sample Log, Log Files Tab



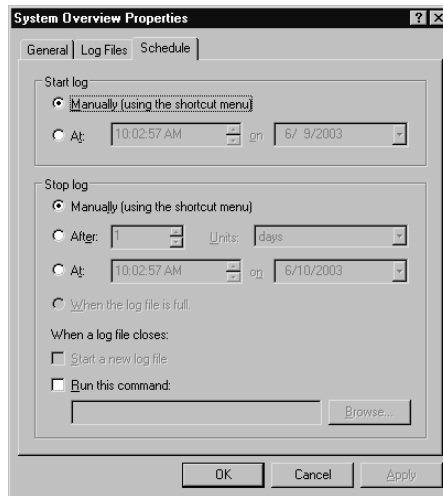
If you click **Configure...** in the **Log Files** tab of the **System Overview Properties** dialog box, you will be presented with the **Configure Log Files** dialog box, as shown in Figure 8.9. Here, you can specify the location, name, and size limit of the log file.

Figure 8.9 Configuring Log Files



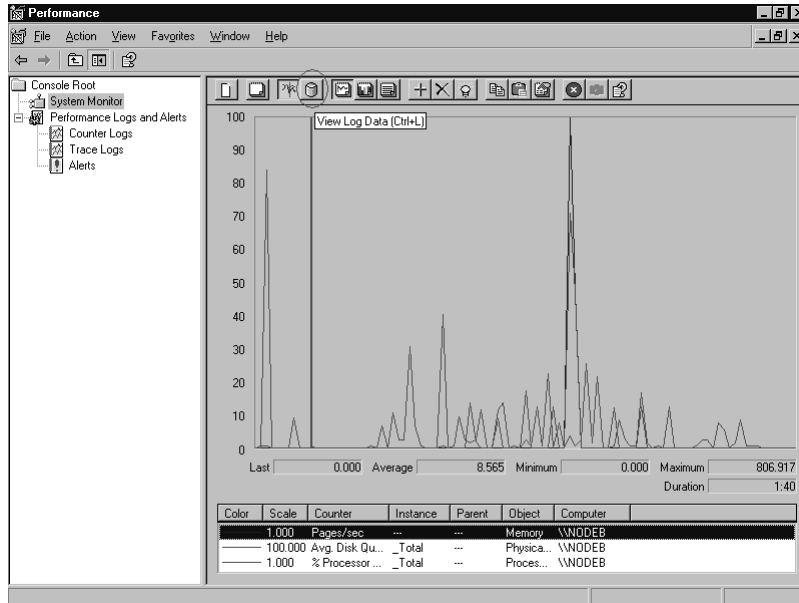
The **Schedule** tab, shown in Figure 8.10, lets you specify the schedule for collecting data from the selected counters. The log can be manually controlled or can be scheduled to start at a specific date and time. You can stop the log manually or after a specific duration, a certain date and time, or when the log reaches its maximum size.

Figure 8.10 Properties of the System Overview Sample Log, Schedule Tab



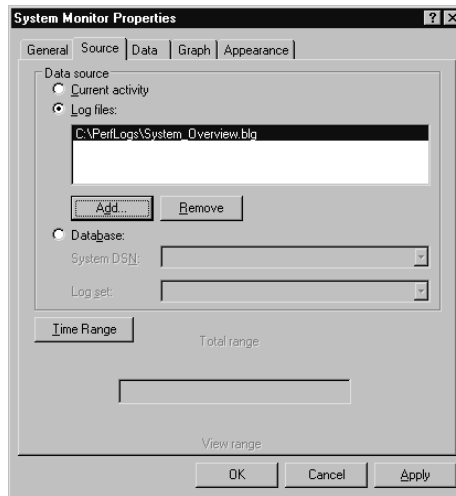
By default, System Monitor tracks real-time data, but you can also have it display data from log files. To view log file data, click the **View Log Data** button in the main System Monitor window, as shown in Figure 8.11 (the icon that looks like a disk, the fourth from the left; circled for clarity in the figure).

Figure 8.11 Selecting the View Log Data Button



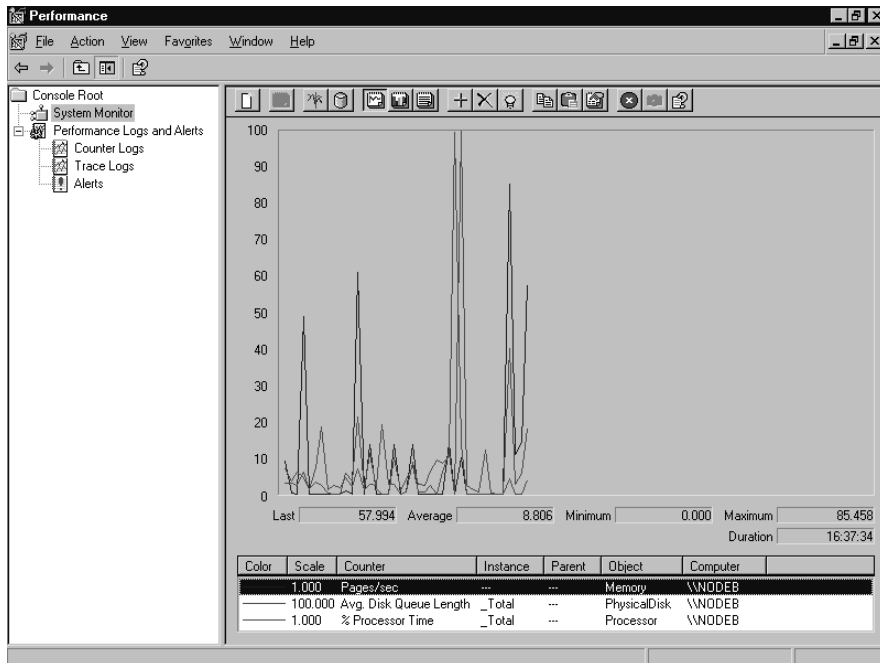
On the **Source** tab, select the **Log files** option button and click **Add**. Browse to the log file you wish to view, select it, and click **Open**. With a log file added, the **Source** tab should look similar to Figure 8.12.

Figure 8.12 System Monitor Properties, Source Tab



Click **OK**, and you will be viewing the data collected in the log file. Figure 8.13 shows an example of viewing log file data.

Figure 8.13 System Monitor, Viewing Log File Data



Determining if performance is acceptable can be highly subjective. It varies depending on the system, role, and environment. There are several general counters and specific thresholds for these counters that you can use to monitor performance. You should examine these counters as ratios over a period of regular intervals, rather than as the average of specific instances. This will provide a more realistic picture of the actual activity occurring on your system. In addition, watch for consistent occurrences of the threshold values being exceeded. It is not uncommon for momentary activity in a system to cause one or more counters to exceed threshold values, which may or may not be acceptable in your environment.

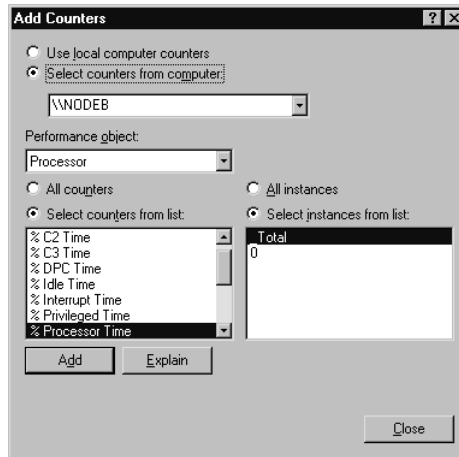
You can use System Monitor and Performance Logs and Alerts to monitor the local system or another computer on the network, as shown in Figure 8.14. It can be useful to compare the performance of the same resource on multiple systems. Be cautious when you do this, though. Ensure that you are comparing appropriately similar objects. Watch out for the “apples and oranges” mismatch. Also, consider that a server being monitored locally may have less monitoring overhead than one that is monitored remotely. This is particularly true regarding the network- and server-related counters, which can be skewed by the transmission of the performance data to your monitoring system. Be sure to account for this difference when developing your statistics.



NOTE


You can also use System Monitor to track statistics from several different computers in the same log file or System Monitor graph. It can be very helpful to have all critical server data in one window or log, so that you can instantly spot problems on one or more of your servers by looking at a single screen.

Figure 8.14 Selecting Counters from Another Computer



If the **Explain** button in the **Add Counters** dialog box is not grayed out, a supporting explanation of the counter is available. The explanation for the `Memory:Pages/sec` counter is shown in Figure 8.15.

Figure 8.15 Viewing a Counter Explanation



Once you have become comfortable and proficient with reading counters, developing baselines, unobtrusively monitoring system activity, and comparing performance, you are ready for the final performance task: determining when your system will no longer be capable of performing the tasks that you want it to perform. Eventually, every computer will be outdated or outgrown. If you have developed the skills for monitoring your system, you should be able to determine in advance when your system will be outgrown. This will allow you to plan for the eventual expansion, enhancement, or replacement of the system. By taking this proactive approach, you can further reduce unplanned downtime by being prepared.



NOTE

Often overlooked, Task Manager is a useful tool for managing the system. Task Manager can assist you in getting an immediate picture of the activities occurring on your system. Stalled applications can be identified on the **Applications** tab. The amount of CPU and memory in use by each active process in the system is available on the **Processes** tab. The **Performance** tab provides a wealth of information on overall system activity, including a real-time bar graph for each processor in the system, showing the amount of time each is in use. The new **Networking** tab shows a real-time graph of the percentage of network bandwidth the system is using.

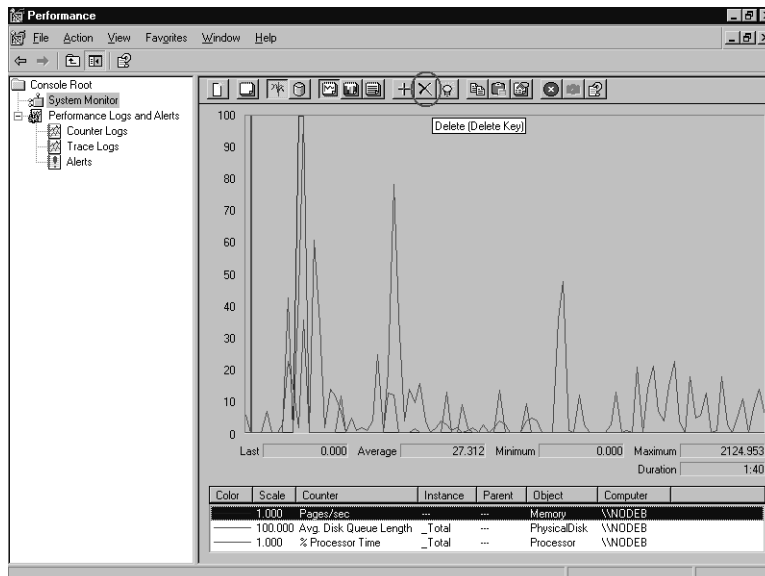
Exercise 8.1 will help you to become proficient in using System Monitor in the Performance console. You can complete the exercise from any Windows Server 2003 computer. Refer to Table 8.1, earlier in the chapter, for information about common counters.

EXERCISE 8.01

CREATING A SYSTEM MONITOR CONSOLE

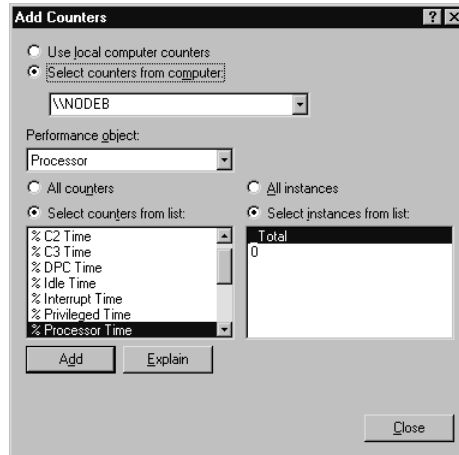
1. Select **Start | All Programs | Administrative Tools | Performance**. Click **System Monitor**. If any counters are already present, click the **Delete (X)** button on the toolbar, circled in Figure 8.16, until the System Monitor window is empty.

Figure 8.16 Empty System Monitor

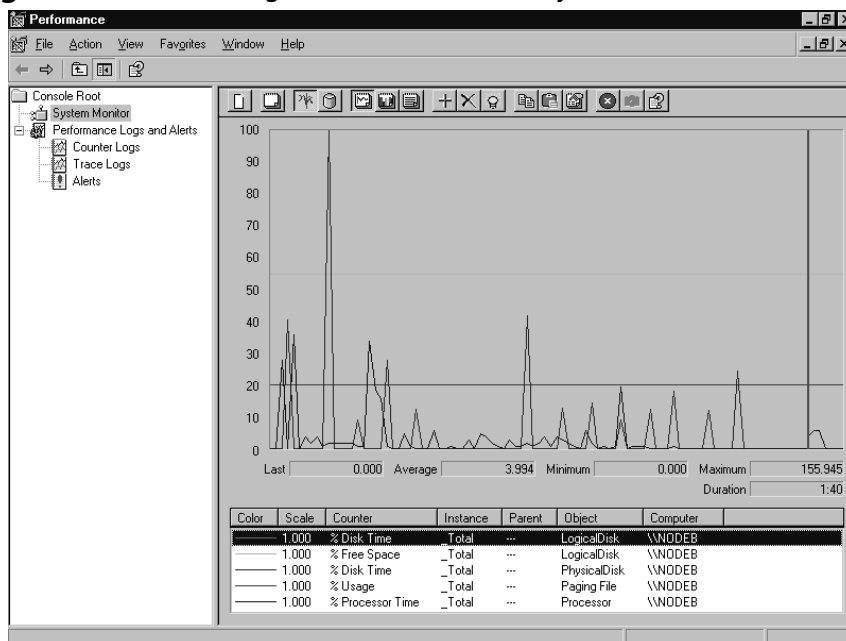


- Click the **Add (+)** button on the toolbar. The **Add Counters** dialog box appears, as shown in Figure 8.17.

Figure 8.17 Add Counters



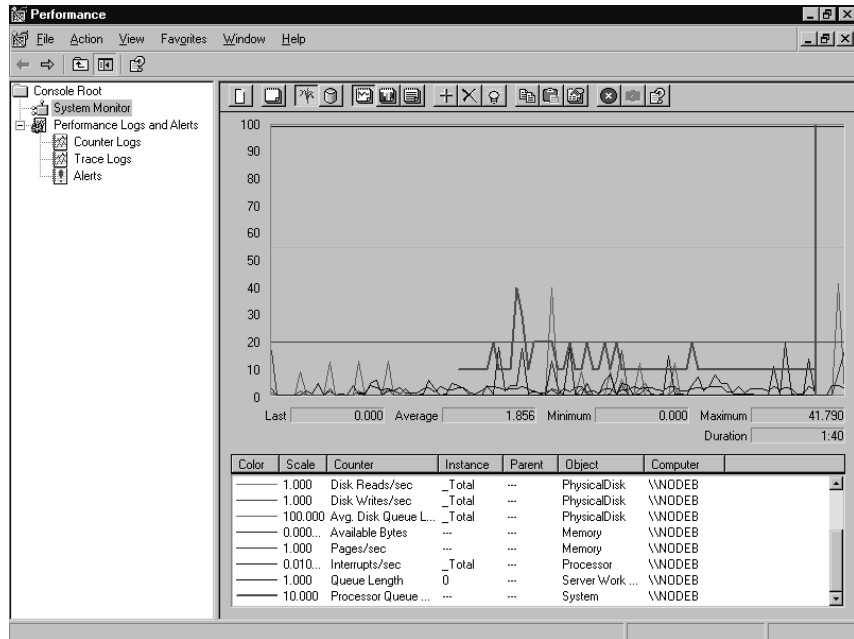
- In the **Performance object** drop-down list, select the **Logical Disk** object.
- In the **Select counters from list** box, select the **%Free Space** counter.
- In the **Select instances from list** box, select **_Total**, and then click **Add**.
- Repeat steps 3, 4 and 5, but select the following performance objects and counters (listed in the form *performance object:counter*):
 - **Physical Disk:%Disk Time**
 - **Logical Disk:%Disk Time**
 - **Paging File:%Usage**
 - **Processor:%Processor Time**
- Click **Close**.
- You should now see a System Monitor window similar to Figure 8.18. Observe the graph as it progresses. Compare the scale of the counters to each other.

Figure 8.18 Percentage-based Counters in System Monitor

9. Click the **Add (+)** button.
10. In the **Performance object** drop-down list, select **Physical Disk:Disk Reads/sec**.
11. In the **Select instances from list box**, select **_Total**, and then click **Add**.
12. Repeat steps 9, 10, and 11 to add the following counters:
 - **Physical Disk:Disk Writes/sec**
 - **Physical Disk:Avg. Disk Queue Length**
 - **Memory:Available Bytes**
 - **Memory:Pages/sec**
 - **Processor:Interrupts/sec**
 - **Server Work Queues:Queue Length**
 - **System:Processor Queue Length**
13. Click **Close**.
14. Your System Monitor window should look similar to Figure 8.19. Notice how busy the chart is beginning to look. Again, compare the scale of the counters to each other. Notice how several counters seem to stay at the bottom of the chart even though they are active. This illustrates

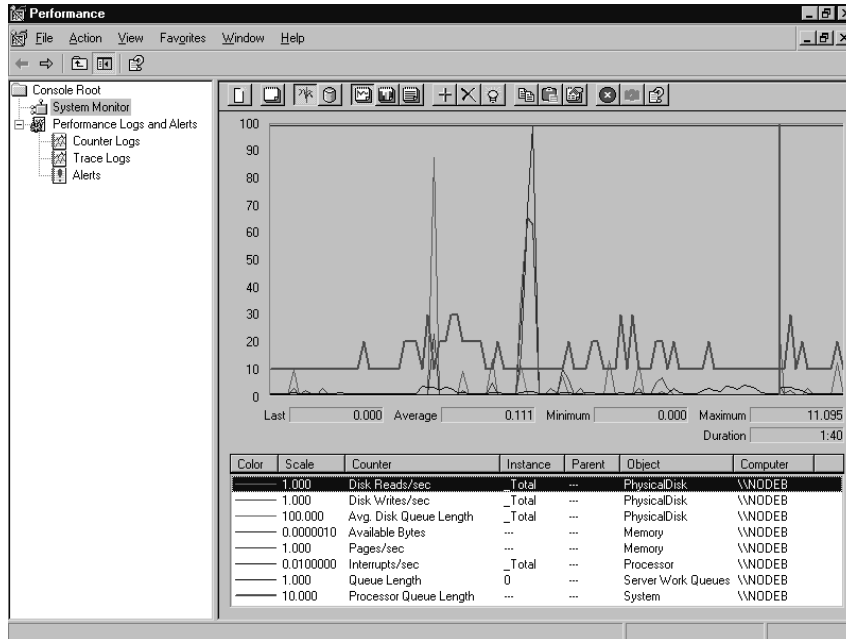
that you should consider scale and try not to mix percentage-based counters with nonpercentage-based counters on the same graph.

Figure 8.19 All Common Counters in System Monitor



15. Click **Logical Disk:%Free Space**.
16. Click the **Delete (X)** button.
17. Repeat steps 9, 10, and 11 to add the following counters:
 - **Physical Disk:%Disk Time**
 - **Logical Disk:%Disk Time**
 - **Paging File:%Usage**
 - **Processor:%Processor Time**
18. You have removed all of the percentage-based counters from the graph. Your System Monitor window should appear similar to Figure 8.20. Compare the scale of the nonpercentage counters.

Figure 8.20 Common Nonpercentage Counters



19. Close the **Performance** console.

Using Event Viewer to Monitor Servers

Windows Server 2003 includes several log files that collect information on events that occur in the system. Using these log files, you can view your system's history of events. A standard Windows Server 2003 system has three event logs that record specific categories of events:

- **Application** Contains events generated by server-based applications, such as Microsoft Exchange and WINS. The specific events logged by each application are determined by the application itself and may be configurable by an administrator within the application.
- **Security** Contains events relating to system security, including successful and failed logon attempts, file creation or deletion, and user and group account activity. The contents of this file will vary depending on the auditing settings selected by the system administrator.
- **System** Contains events relating to the activity of the operating system. Startups and shutdowns, device driver events, and system service events are recorded in the

System log. The configuration and installed options of the operating system determine the events recorded in this log. Because of the nature of its entries, this log is the most important for maintaining system health.

In addition to these three basic logs, a Windows Server 2003 system configured as a domain controller will also have the following two logs:

- **Directory Service** Contains events related to the operation of Active Directory (AD). AD database health, replication events, and Global Catalog activities are recorded in this log.
- **File Replication Service** Contains events related the File Replication Service (FRS), which is responsible for the replication of the file system-based portion of Group Policy Objects (GPOs) between domain controllers.


Finally, a server configured to run the DNS Server service will have the DNS Server log, which contains events related to the operations of that service. Client DNS messages are recorded in the System log.

Events entered into these log files occur as one of five different event types. The type of an event defines its level of severity. The five types of events are as follows:

- **Error** Indicates the most severe or dangerous type of event. The failure of a device driver or service to start or a failed procedure call to a dynamic link library (DLL) can generate this type of event. These events indicate problems that could lead to downtime and need to be resolved. The icon of an error event appears as a red circle with a white X in the middle.
- **Warning** Indicates a problem that is not necessarily an immediate issue but has the potential to become one. Low disk space is an example of a Warning event. This event type icon is a yellow triangle with a white exclamation point (!) in it.
- **Information** Usually indicates success. Proper loading of a driver or startup of a service will generate an Information event. This icon is a white message balloon with a blue, lowercase letter *i* in it.
- **Success Audit** In the Security log, indicates the successful completion of an event configured for security auditing. A successful logon will generate this event. This icon is a gold key.
- **Failure Audit** In the Security log, indicates the unsuccessful completion of an event configured for security auditing. An attempted logon with an incorrect password or an attempt to access a file without sufficient permissions will generate this type of event. The event's icon is a locked padlock.


The event logs are very helpful for collecting data, but we need a tool to present, filter, search, and help us interpret the data. That tool is Event Viewer, shown in Figure 8.21, which can be accessed by selecting **Start | All Programs | Administrative Tools | Event Viewer**.

Figure 8.21 The Event Viewer Window



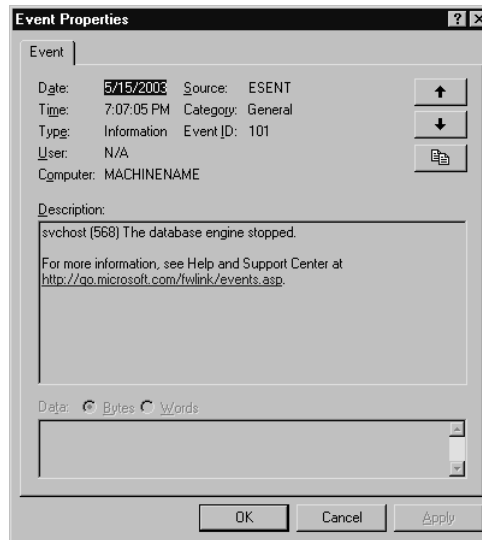
Event Viewer can also be accessed as a component of the **System Tools** snap-in within the **Computer Management** utility, as shown in Figure 8.22.

Figure 8.22 Event Viewer, as Viewed from Computer Management



When viewing an event log, the events appear in the order they occurred. Double-clicking an event will bring up the properties of that event, as shown in Figure 8.23. Click the arrows to navigate to either the next or previous event.

Figure 8.23 Viewing Event Properties




Each event captured follows the same format and contains the same set of data points. Those data points form the *event header* and are as follows:

- **Date** The date the event occurred.
- **Time** The time the event occurred.
- **Type** The applicable type of event (Error, Warning, and so on).
- **User** The user or account context that generated the event.
- **Computer** The name of the computer where the event occurred.
- **Source** The application or system component that generated the event.
- **Category** The classification of the event from the event source's perspective.
- **Event ID** A number identifying the specific event from the source's perspective.
- **Description** A textual description of the event. This may be in any readable structure.
- **Data** A hexadecimal representation of any data recorded for the event by the source.

An event log can contain thousands or even millions of events. Because the event header follows the same structure regardless of the event source, you can use the filter func-

tion to focus in on specific patterns of events. The filter function is available from the log's Properties dialog box. Right-click a log in the left tree view and select **Properties** from the context menu, as shown in Figure 8.24, to view its properties.

Figure 8.24 Accessing the Properties of an Event Log



Click the **Filter** tab to display the filter options, as shown in Figure 8.25. (You can also access the filter function by clicking **View | Filter**.)

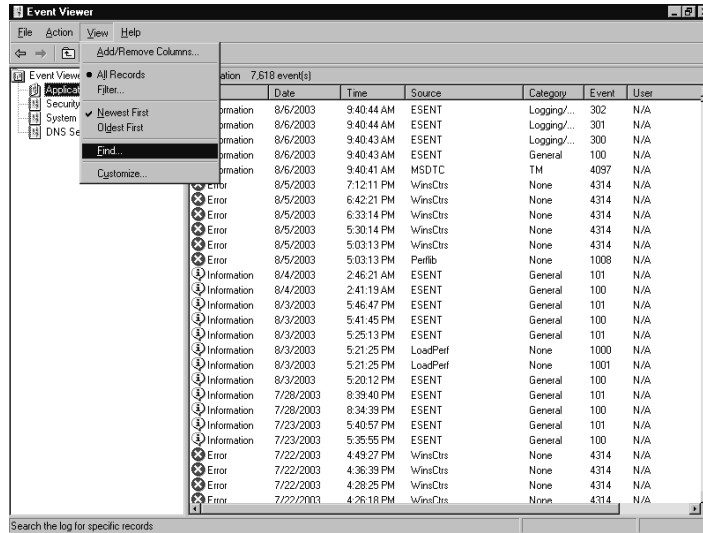
Figure 8.25 Filtering Event Log Data



By changing the selections on the **Filter** tab, you can exclude from view those events that do not fit the filter selections. Put another way, events that do not match the filter

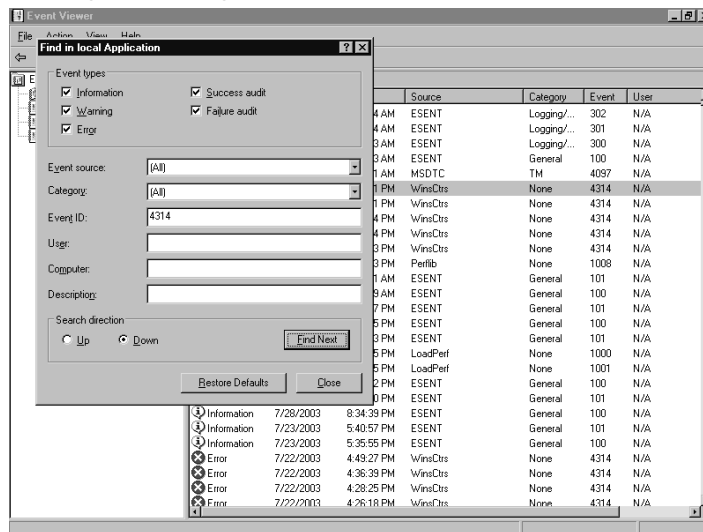
selections are filtered out. The events are still in the log; they are just not displayed as long as the filter is active. In addition to using the filter function, you can search event logs for specific events. From the Event Viewer main window, select **View | Find...**, as shown in Figure 8.26.

Figure 8.26 Using Find in an Event Log



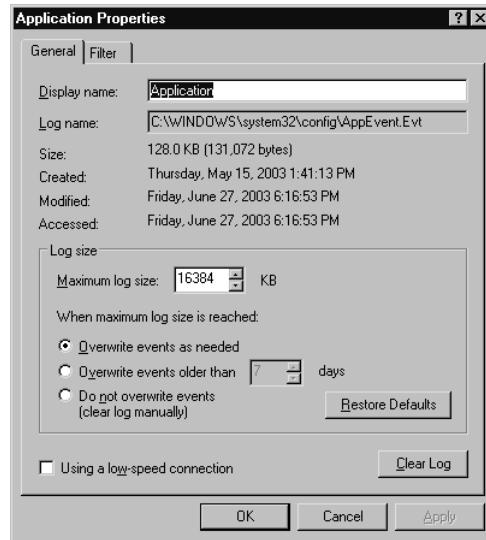
In the **Find** dialog box, enter your criteria for the search and click the **Find Next** button. The next event that matches your criteria will be highlighted in the Event Viewer main window, as shown in the example in Figure 8.27.

Figure 8.27 Finding Event Log Data



The event log files themselves are stored in a compact binary format in the `%systemroot%\System32\Config` directory. You can configure the maximum size of these files and what action is taken when this size is reached on the **General** tab of the log **Properties** dialog box, as shown in Figure 8.28.

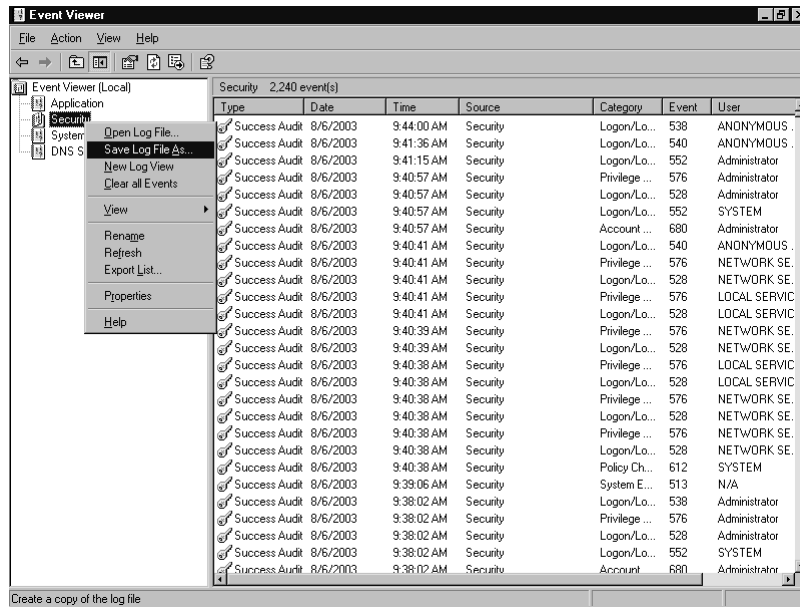
Figure 8.28 Event Log General Properties



Accessing the Properties dialog box of an event log gives you access to information about the log itself, and allows you to change certain characteristics of the log. Referring to Figure 8.28, you can see the log's name, location, and size. The **Maximum log size** option allows you to limit the amount of space the log consumes. The three radio buttons below this option allow you to specify what will happen when this maximum size is reached.

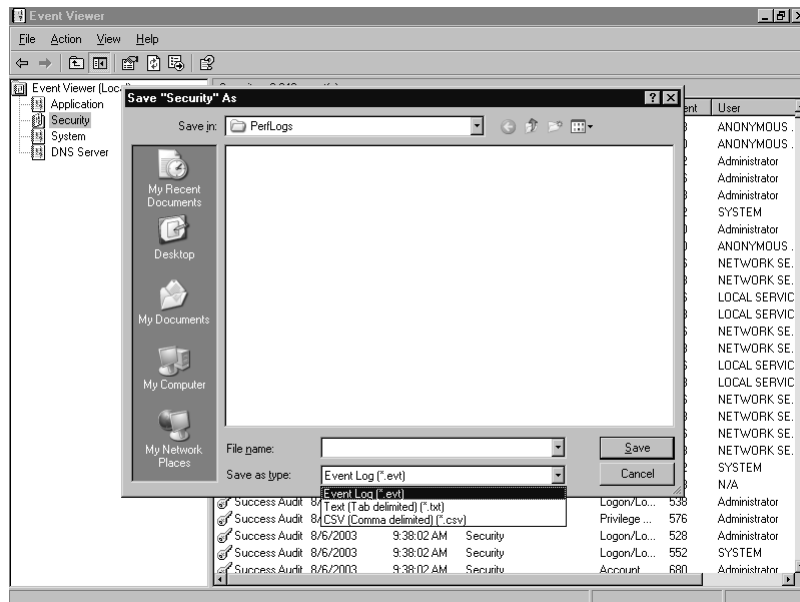
Event logs can be archived on the computer on which they occur for long-term storage and analysis. This can be accomplished in two ways. The first is through the use of the **Clear Log** button on the event log Properties dialog box. You can click this button to delete all entries from a log file, but this process will also prompt you to save the events prior to deletion. The second method is through the use of the **Save Log File As...** option on the context menu for a log file, as shown in Figure 8.29.

Figure 8.29 Saving a Log File, Selection Menu




The **Save Log File As...** selection brings up a **Save AS** dialog box, as shown in Figure 8.30, which allows you to choose the name, location, and format of the archive.

Figure 8.30 Saving a Log File




You can save events in a binary .evt, comma-delimited, or tab-delimited text file. You can use the .evt format to retain the log file in a compact format, which you can reopen in Event Viewer by selecting **Action | Open Log File...**, as shown in Figure 8.31. The delimited archive file formats consume more disk space than the .evt format, but they can be imported into a database or an application like Microsoft Excel for further analysis.

Figure 8.31 Opening an Archived Log File



Choosing to open a log file brings up an Open dialog box, shown in Figure 8.32. In this dialog box, you can locate and choose the archived file.

Figure 8.32 Selecting an Archived Event Log



Using Service Logs to Monitor Servers

As mentioned previously, there are additional event logs for servers in certain roles. A server running the DNS server will have a DNS Server log. A server acting as a domain controller will have logs for the Directory Service and File Replication Service. It is also possible that other services or applications may create their own log files, but most do not.

These server log files follow the same format as the other event logs. They can be filtered, searched, and archived using the methods described in the previous sections. These logs exist mainly to collect the events from these services in one place other than the System log. These services generate a greater number of events than do other services.



Planning a Backup and Recovery Strategy

Backups and documentation are usually of critical importance to the continuing operation of an organization. Organizations often account for the value of their computer and communications equipment, but they overlook their data, which can be difficult to value. Equipment can be replaced. Staff can be hired. But if data cannot be restored, it is lost forever.

You should consider good backups as a form of insurance. Hard drives fail. Cooling fan bearings wear out, and systems overheat. Lightning strikes buildings. Viruses contaminate or destroy data. Buildings get flooded. Bugs in applications do things no one intended. Burglars steal equipment. People enter the wrong information, delete the wrong files, or get emotional and destroy data. Untrained employees accidentally damage data and hardware. And since September 11, 2001, the threat of high-tech terrorism has become an important consideration. For our purposes, the cause of loss is irrelevant. The most important point is the ability to recover from any loss that occurs.

When considering the factors that make backups necessary, you should also consider the human side of the situation. The software won't operate itself, and someone must change tapes. Do not forget to develop good procedures for the people (or person) responsible for your backups. It is important to ensure that there is more than one person who is capable of restoring data, which makes good written procedures essential.

When developing your backup and restore procedures, consider the following guidelines:

- **Develop a log** This gives you a hardcopy record of your backup activities.
- **Test your procedures, devices, and media frequently** A failure in any one of these areas can make data impossible to restore.
- **Keep multiple copies** Media can and does go bad. Shelf life, manufacturing defects, and environmental or physical damage can render media impossible to read.
- **Rotate copies offsite** Keep the backups in a different location. That way, a local disaster won't destroy all of your backups.

- **Back up the system** The operating system and your applications are a form of data, too, and should be protected accordingly.
- **Use the new Automated System Recovery (ASR) feature** This feature saves time in the event of a disaster and can also act as a “last-ditch” effort before a complete rebuild. Perform an ASR backup after each major system change and also on a regular basis.
- **Secure your backups** Secure your backup media in the same way that you would secure any other valuable item. Keep your media locked up in a safe, if possible.
- **Know your data** Data that changes frequently may need more frequent backups. Databases require different strategies than documents and spreadsheets. Encrypted File System (EFS) files and folders should have the recovery agent’s EFS private key backed up as well; otherwise, recovering EFS files and folders may not be possible. The DHCP, WINS, DNS, and AD services have specific backup or restore requirements. It is important to understand these requirements when you plan your backup strategy.

EXAM
70-293
OBJECTIVE
4.5.1

Understanding Windows Backup

Windows Server 2003 includes the Backup Utility for performing backups, restores, and running the ASR Wizard. The utility can back up data to and restore data from almost any removable media device identified by the operating systems—tape drives, hard drives, and even file shares on the network. You cannot, however, back up to recordable CD or DVD drives.

In order to perform a backup or restore operation, you must have the appropriate user rights. The Administrators and Backup Operators groups are assigned the necessary rights to perform both functions, so using an account that is a member of either of these groups will suffice.

The specific user rights required to perform a backup or restore can be individually assigned by using the Local Security Policy utility, shown in Figure 8.33, or a GPO if the user is a member of an AD domain. The following are the user rights required to perform backup and restore operations:

- **Back up files and directories** Allows a user to bypass (if necessary) established permissions on files, directories, and Registry keys and values. Be cautious when assigning this right, because this can be a security risk. A user with this right could easily back up all of your company’s most sensitive information and carry it out the door.



NOTE

Any user with Read permission to a file can back it up, without needing the Back up files and directories right.

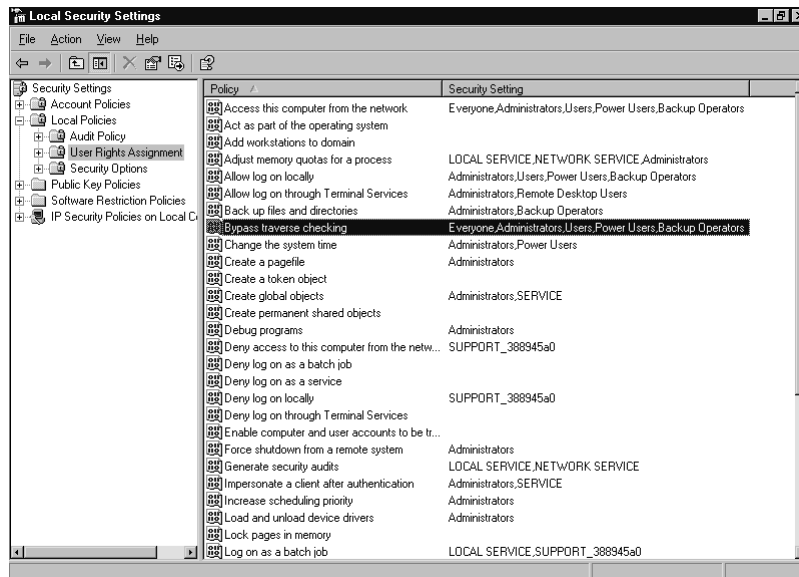
- **Bypass traverse checking** Gives a user the ability to cross directories, whether or not that user has permissions to those directories.
- **Restore files and directories** The corollary user right to Back up files and directories. Allows a user to bypass (if necessary) the established permissions on files, directories, and Registry keys and values. This effectively gives a user the ability to restore objects, regardless of the objects' assigned permissions. You should be cautious with this right due to the potential security risk and possibility of destroying or corrupting data.



NOTE

A user with the Restore files and directories right can strip objects of their permissions during the restore operation. This means that users with this right will be able to restore and access any file or Registry object they wish. To prevent this, consider using a form of encryption such as EFS. When a file is encrypted, it is backed up as encrypted. When restored, it is still encrypted, regardless of the permissions it has assigned to it.

Figure 8.33 Detailed User Rights, Accessed from Local Security Policy



Types of Backups

Most good backup strategies adopt a method of backing up different amounts of data at different times and for different purposes. The length of time required to back up data on a server increases as the amount of data on a server grows. On many systems, a large amount of data is static or changes infrequently. Finally, the costs associated with consumable media (such as tape cartridges) mean that economics force the issue of using the media in cycles. The basic backup cycle includes a complete or *full backup* and several *incremental* or *differential backups*. Each type of backup serves a specific need.

Full Backups

The Windows Backup Utility calls a full backup a *normal* backup. The full backup, as its name implies, backs up everything specified by the user performing the backup operation. A full backup can include the operating system, system state data, applications, and any other data. With a full backup, everything that is backed up has the file system archive bit reset (cleared). This allows the incremental and differential backup types to determine if the file needs to be backed up. If the bit is still clear, the other backup types know that the data has not changed. If the bit is set, the data has changed, and the file needs to be backed up.



NOTE


There is a variation on the full backup type called a *copy backup*. This works exactly like the full backup type with one important exception: the archive bit on backed up files is not cleared. This means the subsequent full, incremental, or differential backups will not be aware that these files have been backed up. This is a useful feature if you want to perform an extra backup between other scheduled backups. Because the differential and incremental backups do not know this backup occurred, they are operationally unaffected by a copy backup. This feature is also useful when you need to get a backup of files but preserve the state of the file system. This is sometimes necessary when installing some software applications. Check the documentation for the application.

The full backup is usually the first backup performed on a server. It takes the longest of all the backup types to complete, because it backs up all specified files regardless of the state of the archive attribute. A full backup consumes the largest amount of backup media of any backup type. Depending on the amount of information chosen to back up and the underlying backup technology involved, it may require multiple backup media to complete.

The main advantage of the full backup type is the ability to rapidly restore the data. All of the information is contained in a single backup set when this type of backup is used. The disadvantages of full backups are high media consumption and long backup times.

Figure 8.34 illustrates a series of full backups. The values listed are relative.

Figure 8.34 Full (Normal) Backup Pattern




Incremental Backups

During an incremental backup operation, all specified files have their archive bit examined. If the bit is set, the file is backed up, and then the bit is cleared. This backup type is used to back up data that has changed or been created since the last full (normal) or incremental backup. It can also be used after a copy or differential backup, but because these do not reset the archive attribute, there is no way for the incremental backup to tell which files have changed since one of those backups last ran. As a result, every file with the archive attribute set is backed up.

The incremental backup type is used between full backups. It is quick to perform, collects the least amount of data, and consumes the smallest amount of media. A complete restore, however, requires the last full backup and every incremental backup (in sequence) since the full backup was performed.

The primary benefits of using the full/incremental backup combination, as illustrated in Figure 8.35, are time and media savings. The main drawback of this combination is longer and more complex restore operations if there are long periods between full backups.

Figure 8.35 Full (Normal) Backup/Incremental Backup Pattern




Differential Backups

The differential backup type is sometimes used as a substitute for the incremental type. A differential backup collects data that has changed or been created since the last full (normal) or incremental backup, but it does not clear the archive bit on the file. It can also be used after a copy or differential backup, but as with an incremental backup, every file with the archive attribute set is backed up.

The differential backup is advantageous when you want to minimize the restoration time. A complete system restore with a full/differential backup combination, as illustrated in Figure 8.36, requires only the most recent full backup and the most recent differential backup. Differential backups start with small volumes of data after a recent full or incremental backup, but often grow in size each time, because the volume of changed data grows. This means that the time to perform a differential backup starts small but increases over time as well. In theory, if full or incremental backups are infrequent, a differential backup could end up taking as long and reaching the same volume as a full backup.

Figure 8.36 Full (Normal) Backup/Differential Backup Pattern





NOTE

You may also want to use combinations of full (normal), incremental, and differential backups. For instance, if you begin with a full backup over the weekend, it might make sense to perform differential backups on Monday and Tuesday. By later in the week, the quantity of changes may be such that a differential backup cannot be performed overnight. An incremental backup on Wednesday will likely solve the problem, with differential backups continuing after that. Using this system, the restore times are still minimized, because the maximum restoration would involve tapes from the full, incremental, and one differential backup. If a failure occurred before Wednesday, it may take tapes from only the full and, possibly, a differential backup to restore the system.

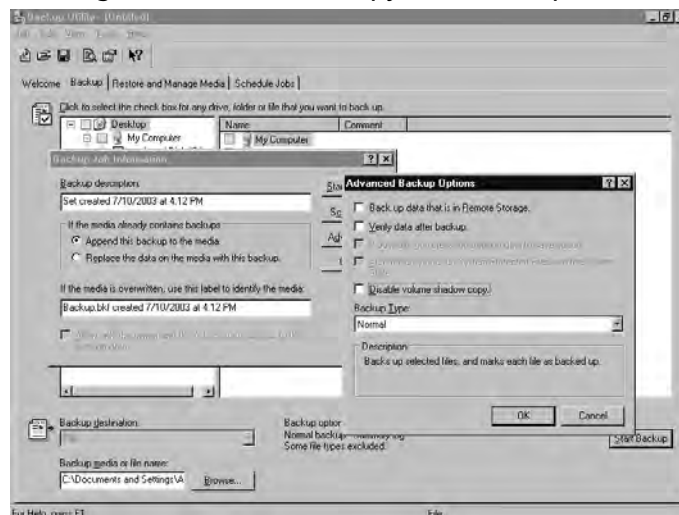
EXAM
70-293
OBJECTIVE
4.5.2

Volume Shadow Copy

More of a new feature than a backup type, Volume Shadow Copy allows you to back up all files on the system, including files that are open by applications or processes. In previous versions of Windows, the applications would need to be stopped or users logged out to allow these files to be closed and backed up using the Windows Backup Utility. With Volume Shadow Copy, these files can continue to remain in use without affecting the integrity of the backup.

This feature is enabled by default, but it may need to be disabled if data managed by some critical applications would be affected by the use of Volume Shadow Copy. The feature can be temporarily disabled by clicking the **Advanced** button in the Backup Utility's **Backup Job Information** dialog box, as shown in Figure 8.37. Unless specified by vendor documentation, leave this feature turned on.

Figure 8.37 Disabling Volume Shadow Copy for a Backup



Determining What to Back Up

Because the data on your servers may be largely static, frequent backups of such data may be redundant. The corollary of this is that more dynamic data needs more frequent backups. Some types of data are structured as multiple files but must be backed up and restored as a single unit to maintain integrity. These factors and more combine to make the development of an efficient backup strategy challenging.

One of the basic techniques you can use to assist you in developing an effective backup and restore strategy is to place your data into basic categories and structure the system around them. For example, on a server that is used for file and print sharing as well as hosting a database, a good structure would be to have separate logical drives for the operating system, the shared files, the application software packages, and the databases. This allows you to easily treat each set of data differently for backup purposes, meeting the specific requirements of each type of data.

Data Backup

If you separate your data into categories, the time required to perform backups can be greatly reduced. For example, once a month, the static parts of the system (operating system and software volumes) could be backed up to tape. For the rest of the month, you can perform either incremental or differential backups. The shared file volume can follow a different schedule, depending on the rate and volume of change in the data. The volume that contains the database files may need full backups nightly in order to expedite restore procedures, and also due to the nature of the database application. It, too, can be easily backed up on a separate schedule from the rest of the system. Tailoring the behavior of backups to each type of data will speed backup and restore operations and minimize the ongoing costs associated with consumable media.

System State Data

The system state data is a special collection of key system and service information. The System State data is present on all Windows Server 2003 systems and includes the following:

- The Registry
- The COM+ Registration database
- Critical boot and system files
- Files protected by Windows File Protection
- The AD database and logs, and the SYSVOL directory (on domain controllers)
- The Certificate Services database (on Certificate Services servers)
- The Cluster Services data (on cluster member servers)
- The Internet Information Server (IIS) Metadirectory (when IIS is installed)

The System State components are designed to allow a system's full identity to be restored, and therefore they are backed up as an entire unit. You can back up System State only locally (unless you're using a third-party application) and restore it only to the system from which it originated.

The Restore to Alternate Location feature is available with a System State restore, but only the Registry, SYSVOL, cluster data, and boot files will be restored. The other components of System State cannot be put in an alternate location and will not be restored. The normal (and arguably best) practice is to back up the System State, boot, and system volumes together. Also, use the ASR feature, which is covered in the "Planning System Recovery with ASR" section later in this chapter.

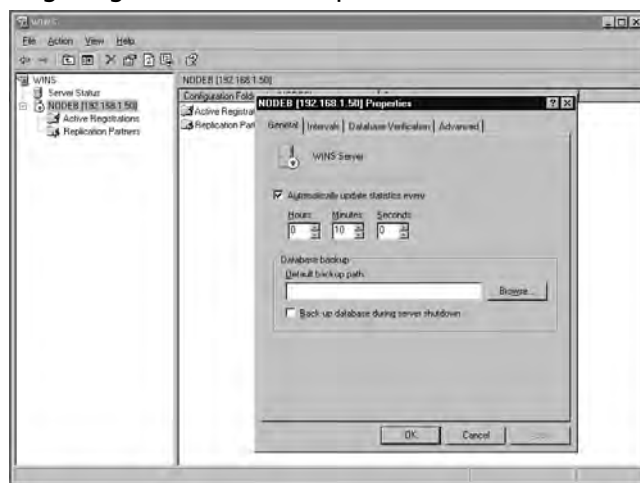
DHCP, WINS, and DNS Databases

DHCP, WINS, and DNS are services that can be hosted by Windows Server 2003. However, each requires some amount of special treatment.

DHCP allows the automatic assignment of IP addresses to systems on the network. When installed, DHCP operates continuously and creates an automatic backup of the DHCP database in %systemroot%\System32\Dhcp\Backup. To manually back up the DHCP database, use the **Action | Backup** command in the DHCP utility. You should then use the Windows Backup Utility to copy this file to your backup media. To restore a DHCP database, first restore the database backup from your backup media, and then use the **Action | Restore** command in the DHCP utility. The DHCP service will be temporarily stopped during the restore operation.

WINS is a service that provides a method of mapping NetBIOS names to IP addresses. WINS is commonly (but not exclusively) used with older versions of Windows. WINS has a built-in backup function, but the function is not activated until you first specify a backup path for the database in the WINS administrative tool by selecting the WINS server and selecting **Action | Properties**, as shown in Figure 8.38.

Figure 8.38 Configuring the WINS Backup Path



Once you have specified a backup directory path, WINS automatically performs a backup of the local WINS database every 24 hours. You should use the Windows Backup Utility to back up this directory to your backup media.

To restore the WINS database, you must first restore the WINS backup directory path from your backup media. Then stop the WINS service, remove all files from the WINS database path, start the WINS utility, select **Action | Restore Database**, and select the file from which to restore the database.

DNS is the name resolution protocol and service used to convert host names to IP addresses. AD is designed to use DNS, and Windows Server 2003 can be used as a DNS server. How DNS data is backed up and restored depends on how DNS is configured. If DNS is configured as an Active Directory-integrated zone, the DNS information is stored in the AD database. This means it is backed up and restored as part of the System State data.

If DNS is not configured as an Active Directory-integrated zone, the individual zone files are automatically backed up by the DNS service, and these files should be used for backup and restore operations. These files are stored in `%systemroot%\DNS\Backup`.

Cluster Disk Signatures and Partition Layouts

Some special care must be taken when backing up and restoring clustered computers. If a clustered server needs to be restored, the original disk signatures and partition structure must also be restored. This is best accomplished by using the ASR feature (covered in the “Planning System Recovery with ASR” section later in this chapter). All cluster nodes should have an ASR backup performed on them, making sure that one node has ownership of the cluster’s quorum resource when the ASR Wizard is running. In the event that clustered disks need recovery, you can use the ASR backup to restore the clustered disk partitions and disk signatures.

Using Backup Tools

The Windows Backup Utility is included in Windows Server 2003 for backing up and restoring your servers. The Backup Utility uses all of the new backup- and restore-related features of Windows Server 2003, including ASR and Volume Shadow Copy. If you are currently using a third-party backup and restore application, you may be surprised by all of the features that the Backup Utility offers in Windows Server 2003.

Using the Windows Backup Utility

The Windows Backup Utility supports three modes of operation: the Backup or Restore Wizard, Advanced Mode, and command-line operation. Each mode is meant to fit different circumstances. The Backup Utility is accessed from **Start | All Programs | Accessories | System Tools | Backup**. It can also be started from a command-line by typing `NTBackup.exe`.

Backup or Restore Wizard

The first time you start the Backup Utility, you are presented with the Backup or Restore Wizard, as shown in Figure 8.39. The purpose of the wizard is to simplify the backup or restore process by stepping you through the process, making the most common options available. The Wizard is best used for initial or manual backups on standardized hardware configurations.

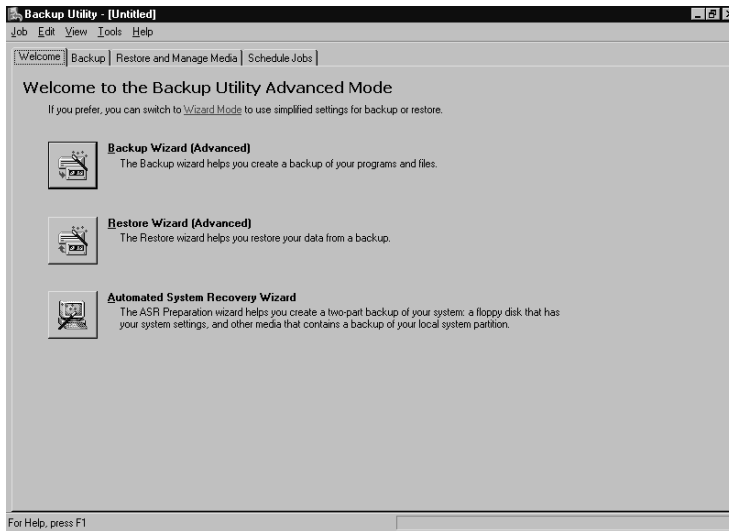
Figure 8.39 The Backup or Restore Wizard



The Wizard does allow you to take advantage of some of the more advanced options, like scheduling, but these options are best configured and controlled by using Advanced Mode.

Advanced Mode

Advanced Mode is accessed by clicking the **Advanced Mode** link in the opening Backup or Restore Wizard window (see Figure 8.39). Advanced Mode gives you direct access to the ASR Wizard, customization options, and reporting and media management functions. If you click Advanced Mode, you are presented with the window shown in Figure 8.40.

Figure 8.40 The Windows Backup Utility, Advanced Mode

Using Advanced Mode, you can predefine different backup jobs and save their settings. You can then set up schedules for these backup jobs to accommodate the needs of your organization and your data. You can also access the **Report** option (available from **Tools | Report**) to get detailed information on the backup and restore activity that has occurred on your system.

Using the Command-Line Tools

You can also run the Windows Backup Utility as part of a batch file or directly from a command prompt. Using this capability, you can integrate the Windows Backup Utility into sophisticated batch files or scripts. Most of the options available in Advanced Mode are available when using the command-line mode. However, you cannot do a restore from the command-line. Restores must be performed with the Wizard or Advanced Mode.

Selecting Backup Media

An important part of your backup and restore strategy is your choice of backup media. Many different types of media are usable by Windows Backup, and each has advantages and disadvantages. You must consider factors such as backup and restore speed, media capacity, media cost, device cost, media shelf life, and the reliability of the technology.

When analyzing these factors, take a long-term view. Technology changes rapidly, but data stays around for a long time. Examine the necessary life of your data. Accounting data usually needs to be recoverable for seven years. Data relating to legal proceedings may need to be retained for decades. Medical research data may need to be retained for *centuries*. No single media or technology will meet all of these requirements, but with proper planning, you can ensure that you and your successors can manage the retention of data.

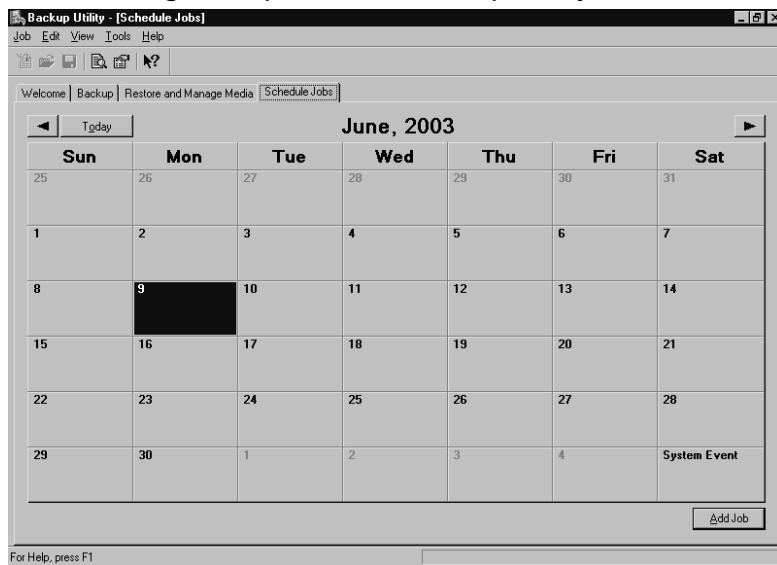
Tape technology has been around for a long time. Tape lends itself very well to high-volume, long-term storage of data. Tape is the most common type of backup media used and is almost always the eventual endpoint of saved data.

The Windows Backup Utility can use any type of tape drive and tape technology supported by Windows Server 2003. When purchasing a tape drive, make sure that the operating system supports it. Choosing the type of tape drive and media can be difficult, since tape technology is widely varied and available in several different formats, capacities, and speeds. Extensive research may be required to choose a technology that matches your requirements for data volume, backup speed, and restore speed.

Scheduling Backups

You can use the **Schedule Jobs** tab in the Windows Backup Utility, shown in Figure 8.41, to create an automated schedule of backup jobs. You can define different types of jobs and different schedules. For example, you can define and schedule normal (full) backups every Friday starting at 6:00 P.M., and differential backups every weeknight starting at 10:00 P.M. The jobs will automatically execute when their scheduled times occur.

Figure 8.41 Scheduling Backups with the Backup Utility



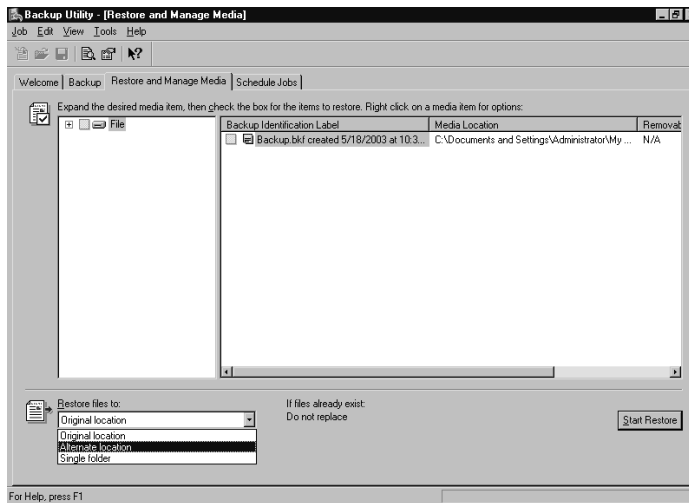
Restoring from Backup

Backing up data is important, but the objective of any backup and restore application is the successful restoration of data after it is lost or corrupted. A backup process without a restore process is useless. As with backups, knowing your data is important when attempting a restore operation. Some types of data must be restored as a unit, some data may require

additional preparation or utilities for a successful restore (AD), and some data may require noting more than a place to put it (normal shared files).

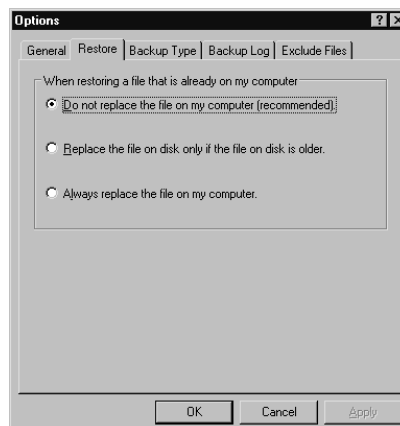
As mentioned previously, you can restore files using either the Backup Utility's Backup or Restore Wizard or Advanced Mode. The first step is to select the backup media to restore from. When using Advanced Mode, you can click the **Restore and Manage Media** tab to select the media, as shown in Figure 8.42. You can expand the media listing on this tab until you find the items you wish to restore, and then select those items by clicking the check box next to each item.

Figure 8.42 Choosing the Restore Source Media



By default, files restored from media will not overwrite existing files of the same name. You can alter this behavior by changing the restore options available on the **Restore** tab of the **Options** dialog box (accessed by selecting **Tools | Options**), as shown in Figure 8.43.

Figure 8.43 The Restore Options



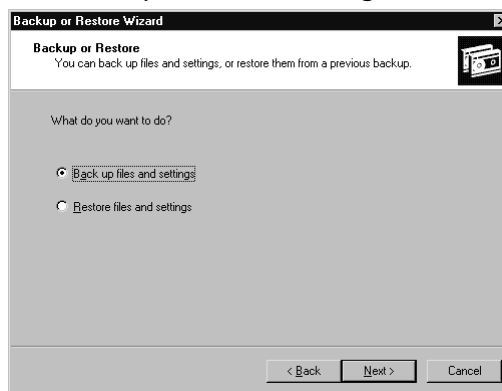
When you have selected the items you wish to restore, you must determine if you want to restore them to their original locations or to an alternate location. This is determined by the setting you select in the **Restore files to** drop-down list on the **Restore and Manage Media** tab (see Figure 8.42). Once you have selected the restore options desired, click the **Start Restore** button to begin the restore process.

EXERCISE 8.02

CREATING A BACKUP SCHEDULE

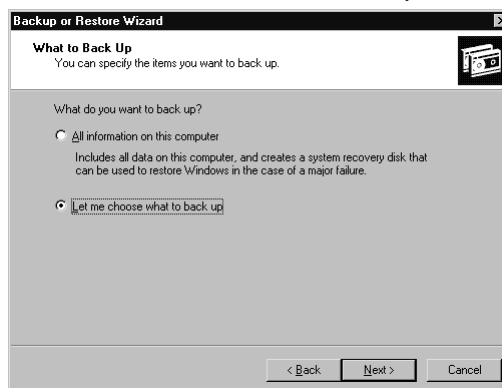
1. Select **Start | All Programs | Accessories | System Tools | Backup**. In the **Backup or Restore Wizard**, click **Next**. Select **Back up files and settings**, as shown in Figure 8.44, and click **Next**.

Figure 8.44 Select Backup Files and Settings



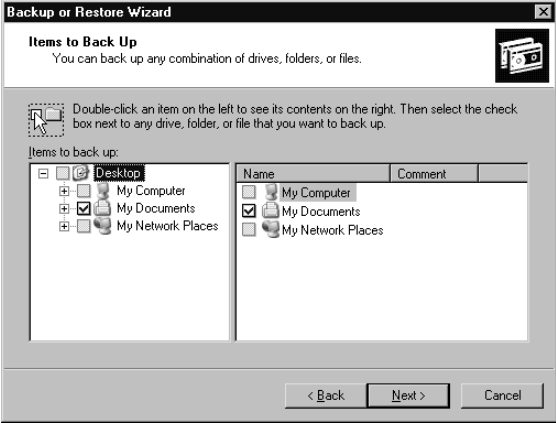
2. Select **Let me choose what to back up**, as shown in Figure 8.45, and click **Next**.

Figure 8.45 Select to Choose What to Back Up



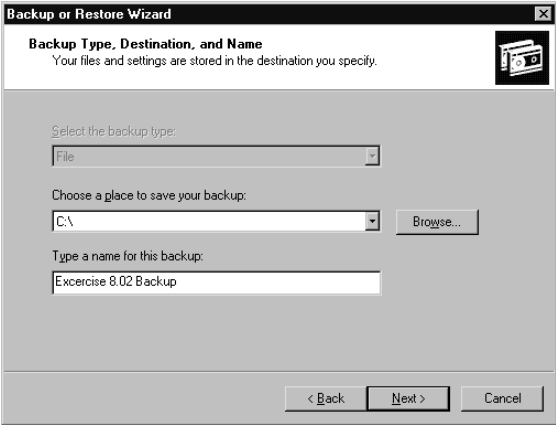
- 3. For this exercise, select **My Documents** in the **Items to Back Up** window, shown in Figure 8.46, and click **Next**.

Figure 8.46 Choose Items to Back Up



- 4. Select the destination, location, and name for your backup, and then click **Next**. Note that for this exercise, a file has been chosen for the destination, as shown in Figure 8.47.

Figure 8.47 Selecting a Destination for the Backup



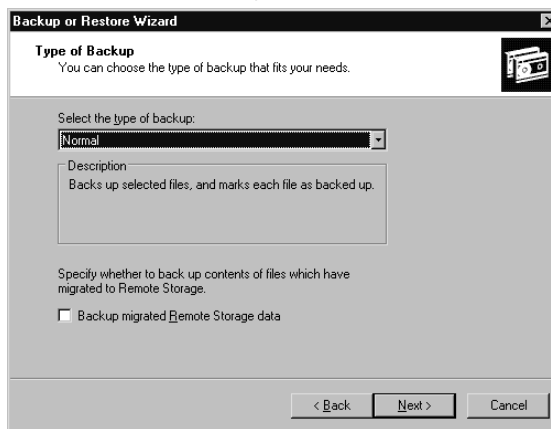
- 5. In the **Completing the Backup or Restore Wizard** window, shown in Figure 8.48, click **Advanced**.

Figure 8.48 Choose Advanced to Specify Backup Options



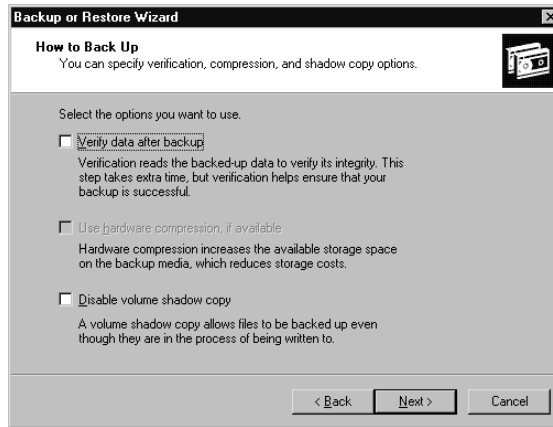
6. Select **Normal** as the backup type, as shown in Figure 8.49, and click **Next**.

Figure 8.49 Select the Backup Type



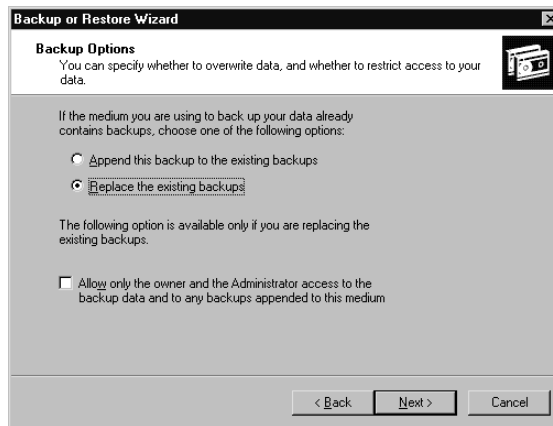
7. Make no changes in the **How to Back Up** window, shown in Figure 8.50, and click **Next**.

Figure 8.50 How to Back Up Options



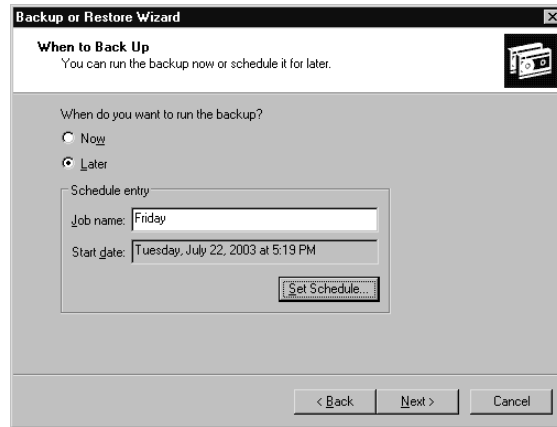
8. In the **Backup Options** window, select **Replace the existing backups**, as shown in Figure 8.51, and click **Next**.

Figure 8.51 Select Backup Options




9. In the **When to Back Up** window, select **Later** and enter **Friday** in the **Job name** text box, as shown in Figure 8.52. Then click the **Set Schedule** button.

Figure 8.52 Specify When to Back Up

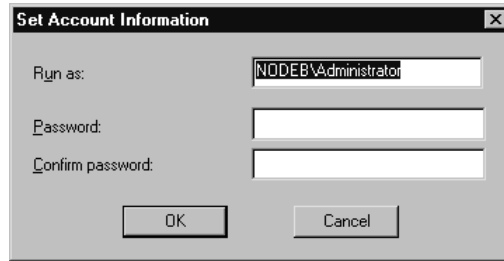


10. In the **Schedule Job** dialog box, change the backup to run **Weekly** on **Friday** at **8:00PM**, as shown in Figure 8.53, and click **OK**.

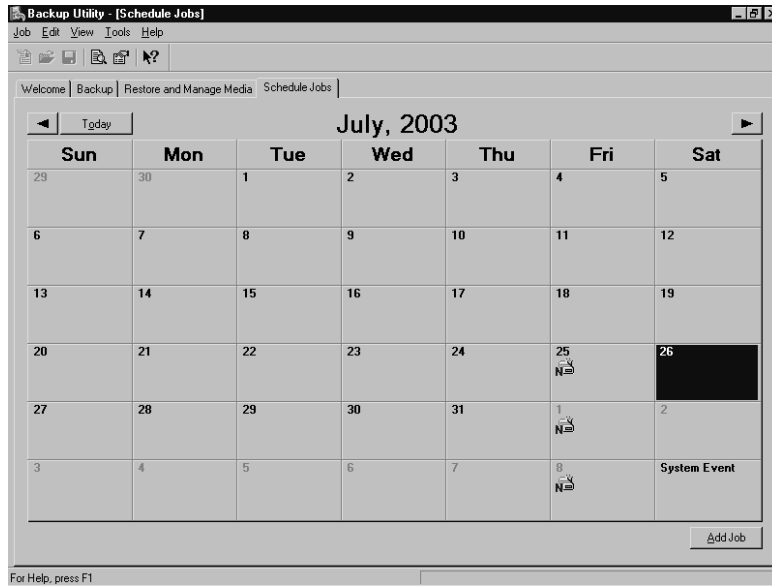
Figure 8.53 Schedule a Weekly Backup Job



11. In the **Set Account Information** dialog box, shown in Figure 8.54, enter an account and password with sufficient permissions to perform the backup and click **OK**. You may be prompted for this information more than once.

Figure 8.54 Set Account Information.

12. Scheduling of the backup is now complete. Close and reopen the Backup Utility.
13. Click **Advanced Mode** and select the **Schedule Jobs** tab. You will see your scheduled backups ready to go, as shown in Figure 8.55.

Figure 8.55 View Scheduled Backups in Advanced Mode

EXAM
70-293
OBJECTIVE
4.5.3
4
4.1

Planning System Recovery with ASR

The ASR feature of Windows Server 2003 is new and replaces the older emergency Repair Disk (ERD) concept. You might have heard the saying, “Outages take seconds; recoveries take days.” ASR was designed to specifically address this issue.

In the past 10 years, the state-of-the-art operating system has gone from DOS to Windows Server 2003. The operating system's complexity has increased, along with the difficulty in troubleshooting and repair. If an operating system component of a DOS system became corrupted, often a single command (SYS) could be used to re-create the operating system on the affected disk within minutes. Starting with Windows NT, however, repair was not that simple. The operating system no longer consisted of a few basic files, but hundreds of files that were linked together by the Registry. Troubleshooting became extremely difficult. A complete reinstallation and recovery was often necessary, followed by hours of tweaking and reconfiguration in an attempt to return to the previous operational state.

Now, with ASR, you can re-create *and* restore the entire operating system *exactly as it was* in one simple and quick process. It's important to note that ASR is not meant as a substitute for regular backups. ASR protects only the operating system and any other data that is on the same partitions or volumes as the operating system files. Typically, applications and data must continue to be backed up on a regular basis outside of your adopted ASR procedures. However, a proper ASR routine can mean the difference between spending a weekend or a couple of hours on recovery.



TEST DAY TIP

ASR is primarily meant as a disaster-recovery tool but has another extremely useful function. It can be used to migrate to different hardware. The ASR utility has enough rudimentary intelligence built in to allow for differences in hardware when a restore is taking place. In effect, this means that ASR can also be used for server migrations.

What Is ASR?

ASR is a two-part, last-resort, system recovery feature for all components of the operating system, including the system state, system services, disk signatures, and partition layouts. It is similar to some third-party disaster-recovery tools, but it is more specific in purpose.

Unlike an operating system reinstallation, an ASR restore will re-create the *exact state* of the operating system at the time the ASR backup was performed. This means that for ASR to be effective, you should make sure that an ASR backup is performed after each change in the operating system.

How ASR Works

ASR involves two main processes:

- **ASR backup** The process of creating an *ASR set*, which consists of a 1.44MB floppy diskette and a linked backup media containing ASR-created backup data.

These two components are necessary for performing an ASR restore and must be kept together.

- **ASR restore** The process of re-creating the operating system and system-related disk partitions/volumes from an ASR set. In addition to the ASR set, you will need to have the original media used to install Windows Server 2003 on your server.

An ASR backup creates a set of all of the information necessary to re-create the operating system at the time the ASR backup is performed. When an ASR restore is performed, the operating system is reinstalled using the original Windows Server 2003 media. However, instead of generating new disk signatures, security identifiers, and Registry content, these items are restored from the ASR set.



NOTE

When operating on a nonclustered server, members of the Backup Operators group can perform ASR backups. This is not the case on clustered servers. Either a member of the Administrators group must perform the ASR backup or the Backup Operators group must be added to the security descriptor for the cluster service.

Alternatives to ASR

Before resorting to an ASR restore, there are a few alternatives that you should attempt for expediency and simplicity. Sometimes, these alternatives resolve the issue, so that an ASR restore is unnecessary.

Safe Mode Boot

A Safe Mode startup starts the system with the minimum number of drivers enabled. Only keyboard, mouse, base video, monitor, disk, and default services are loaded. No network is available. This startup option can sometimes be used to get around a failed software application, service, or device driver that is causing system problems. If the system boots successfully, you can then disable or uninstall the problem driver, service, or application.

Last Known Good Boot Mode

The Last Known Good option starts the system normally but uses the Registry settings from the last successful logon to the system. This is useful to get past misconfiguration issues, especially regarding drivers that can cause system instability. A successful boot with this option will wipe out any setting or configuration changes that have occurred since the last successful logon. Once a logon occurs, these settings will then become the new Last Known Good configuration.

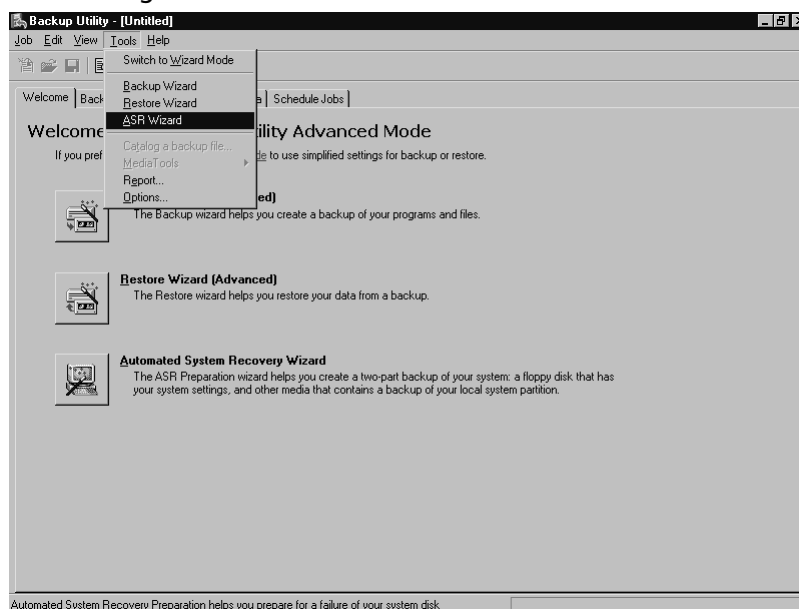
ASR As a Last Resort

If none of the above alternatives work, then an ASR restore may be necessary. Remember that ASR restores and re-creates the system as it was when the ASR set was created. Because of this, it is important to keep your ASR set up-to-date. At a minimum, an ASR backup should be performed after each operating system or system change.

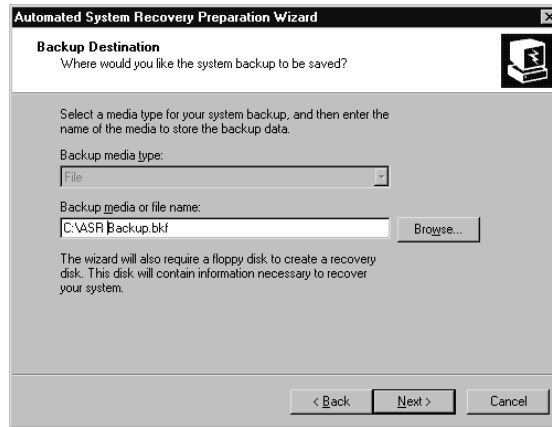
Using the ASR Wizard

The ASR Wizard is accessed from the Windows Backup Utility in Advanced Mode. To start the Wizard, click its icon on the **Welcome** tab or select it from the **Tools** menu, as shown in Figure 8.56.

Figure 8.56 Starting the ASR Wizard



The ASR Wizard will start, prompt you for a destination for the backup, as shown in Figure 8.57, and proceed to create the backup.

Figure 8.57 The ASR Preparation Wizard, Choose a Destination

When the partitions or volumes that contain operating system components have been backed up, you will be prompted to insert a blank 1.44MB diskette, as shown in Figure 8.58. Insert the diskette into the floppy drive and click **OK**.

Figure 8.58 Creating the ASR Diskette

You are not required to have a diskette drive installed to perform an ASR backup, but you *are* required to have a diskette drive installed to perform an ASR restore. You can create the ASR diskette after the Wizard completes by copying the files `asr.sif` and `asrpnf.sif` (located in the `%systemroot%\Repair` directory) to a diskette. If you do not have a floppy disk drive installed in your system, you will see the warning in Figure 8.59. This does not mean that the ASR process will fail; it just means that you will need to create the diskette manually later. Click **OK** to close the warning dialog box.

Figure 8.59 No Floppy Drive Warning

If you are performing an ASR backup without using a diskette, you will see the warning shown in Figure 8.60. Click **OK** to close the dialog box.


Figure 8.60 ASR Diskette Warning Message



Performing an ASR Restore

An ASR restore is a fairly straightforward process. Boot from your original Windows Server 2003 CD-ROM. If a third-party storage driver needs to be loaded, press **F6** when prompted to load the driver. To begin the ASR recovery process, press **F2** when prompted, as shown in Figure 8.61.

Figure 8.61 Text-Mode ASR Prompt



Next, you will be prompted to insert the ASR diskette into the floppy drive, as shown in Figure 8.62.

Figure 8.62 Insert the ASR Diskette Prompt



The ASR process will begin partitioning and formatting your server's boot and system partitions/volumes, as well as any other partitions or volumes that contained operating system files. This process will automatically re-create the operating system as it existed at the time the ASR set was created. If the backup media that is part of the ASR set cannot be located, you will be prompted for its location. Once the ASR restore is complete, the system will reboot.

Planning for Fault Tolerance

Fault tolerance is the ability to encounter failures and still continue to function. Fault tolerance is created by using a combination of *redundancy* (the duplication of components or resources), efficient distribution of workload, proper planning, proper procedures, and training. When all of these are done correctly and in the right proportions, high availability is the result.

To properly plan for fault tolerance, examine all of the possible areas a failure could occur that would affect continuous operation. The following are the most common areas of failure:

- Hardware (disk, RAM, CPU, power supply, cooling fans, and network)
- Infrastructure (power feeds, environmental, and wide-area communications)
- Operational (documentation, change of media, and procedures)
- Functional (placing too many critical processes into a failure-susceptible area).

One fault-tolerant-related phrase you may have heard before is *five nines*, which is a reference to the larger *scale of nines* measure of computer system availability first developed by Jim Gray. The scale of nines refers to the percentage of downtime allowed per year, described by the number of nines in the availability statistic. Five nines refers to an achievable level of reliability in the middle scale. Table 8.2 illustrates the amount of downtime each level of “nines” means per year.

Table 8.2 The Scale of Nines and What Five-Nines Means

Name	Percentage of Uptime per Year	Effective Downtime per Year
One nine	90%	36 days, 12 hours
Two nines	99%	3 days, 15 hours, 36 minutes
Three nines	99.9%	8 hours, 45 minutes, 36 seconds
Four nines	99.99%	52 minutes, 34 seconds
Five nines	99.999%	5 minutes, 15 seconds
Six nines	99.9999%	31.5 seconds
Seven nines	99.99999%	3.2 seconds
Eight nines	99.999999%	0.32 second
Nine nines	99.9999999%	0.03 second

Five nines reliability is commonly discussed because it is possible to achieve given current technology. The primary factor with the scale of nines is cost. Higher levels of availability are becoming possible to achieve, but they usually come at a steep price.

Network Fault-Tolerance Solutions

One area of component failure is the network interface. If a system has one interface to a network, and a component of that interface fails (the switch, the cable, or the NIC), the whole interface fails. As a result, it is a good idea to build redundancy into your network interfaces.

Several manufacturers sell NICs that have two or more ports. Using the appropriate drivers, these cards usually support either a failover configuration or a load-balanced configuration, which work as follows:

- **Failover** Keeps one port idle and waiting, while the other port(s) handle communications. If a component of that interface fails, the idle port comes online and takes over for the failed port. A failover configuration can be used with switches or nonswitched network hubs.
- **Load-balanced configuration** Uses multiple ports simultaneously and spreads the communication load among the ports. In the event of an interface failure, the communications load is reassigned to the remaining active ports. A load-balanced configuration yields higher availability and performance but can be used only in conjunction with higher-end intelligent switches.

Some network topology issues can affect network availability as well. When designing a network, keep in mind all of the potential failure points, including routers, switches, bridges, and wide area network (WAN) components.

In all but the smallest networks, it is a good idea to have redundant functionality for critical services. If you are using AD, make sure that you have more than one domain controller and DNS server. If you are using WINS, create a secondary WINS server and have it replicate with the primary WINS server. If you are using DHCP, create a secondary DHCP server on each subnet and configure each with the appropriate scopes. Following these guidelines will ensure continued operation of these services in the event of failures.

Internet Fault-Tolerance Solutions

Many of the Internet fault-tolerance solutions are the same as general network fault-tolerance solutions, but there are a few extra considerations. Network Load Balancing (NLB) is a set of features included with all versions of Windows Server 2003 that can increase the redundancy, performance, and availability of Web sites. NLB allows multiple Windows Server 2003 Web servers to share and distribute the load of serving Web pages and other network-based applications. NLB is discussed in Chapter 9.

Most medium and large networks access the Internet through a *proxy server*. A proxy server is a server that accesses the Internet on behalf of a requesting client, caches the results, and then passes the requested pages back to the client. Subsequent requests are then served from cache, and actual traffic to the Internet is reduced. This can increase the performance for clients when accessing Web pages. Some proxy servers also add security and other features. If your environment includes a proxy server, consider building redundancy into it. A secondary proxy server may be in order.

The actual communication circuits and Internet Service Providers (ISPs) are other potential points of failure. It is common for large companies and organizations to have multiple WAN circuits and even multiple circuits to more than one ISP. This increases cost but also reduces the likelihood of a communications failure usually outside your control.

Disk Fault-Tolerance Solutions

The most common hardware component that fails is the hard drive. Even though modern disk drives commonly operate for months or years without incident, failure is a given. As a result, disk fault-tolerance solutions are some of the most well-developed and reliable technologies, and they employ some of the oldest and simplest techniques. It is not uncommon to see other areas of technology borrow concepts that were first seen in the development of disk fault tolerance.

The disk controller is the first component to consider. Although the controller itself is generally considered more reliable than the disk drive, a controller can (depending on the technology in use) support multiple drives. This can make the failure of a controller have a stronger impact than if a single drive fails. A configuration that uses multiple controllers connected to a set of drives is often called *duplexing*. Multiple controllers increase the cost and complexity of the configuration, but they can help to eliminate the controller as a point of failure.

RAID

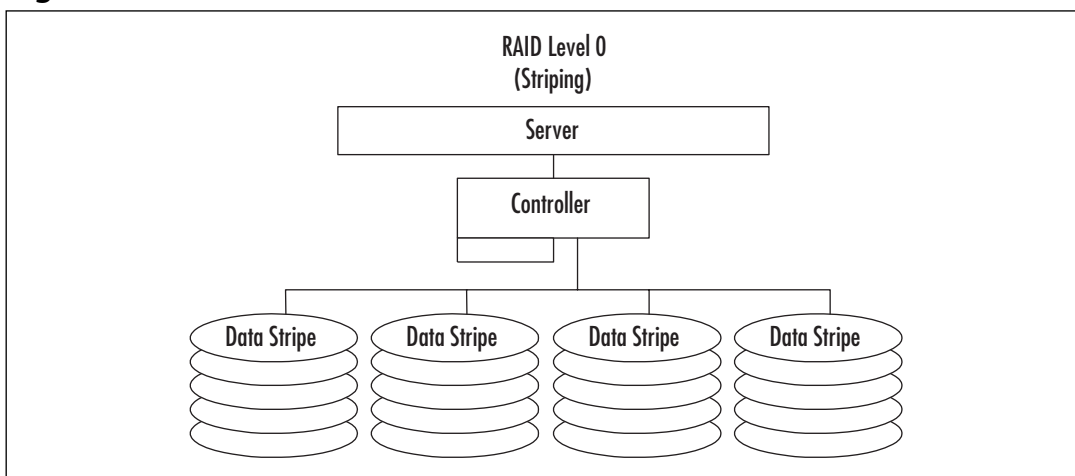
Redundant drives are more common than redundant controllers. Using a technique called RAID, data is duplicated on multiple drives, prevents the failure of any single drive from causing downtime. Like many fault-tolerant solutions, RAID can also have a performance impact. RAID can be done via specialized hardware controllers or via software. Hardware solutions are generally faster and more portable; software solutions are cheaper.

A group of RAID-configured drives is called a *drive array* or simply an *array*. The structure of an array is usually described as its *level*. A level refers to a specific type or combination of redundancy in use. There are several defined RAID levels. The most common RAID levels are described in the following sections.

RAID 0

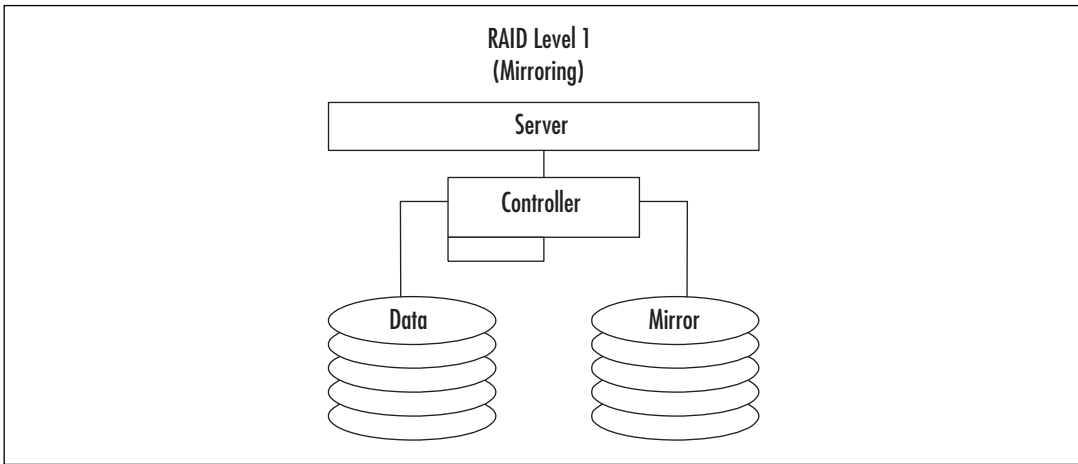
RAID 0 (also called *striping*) provides increased read and write performance and does not provide redundancy. Data is split into smaller uniform-sized blocks (or *stripes*), as illustrated in Figure 8.63, and is written to read from multiple drives at the same time. Multiple drives with smaller blocks of data increase the performance, but the failure of any drive destroys all data. The level is used mainly for ultra high-performance databases. RAID 0 requires a minimum of two drives, although more are commonly used. The maximum number of drives allowed using a Windows Server 2003 striped volume is 32.

Figure 8.63 RAID 0

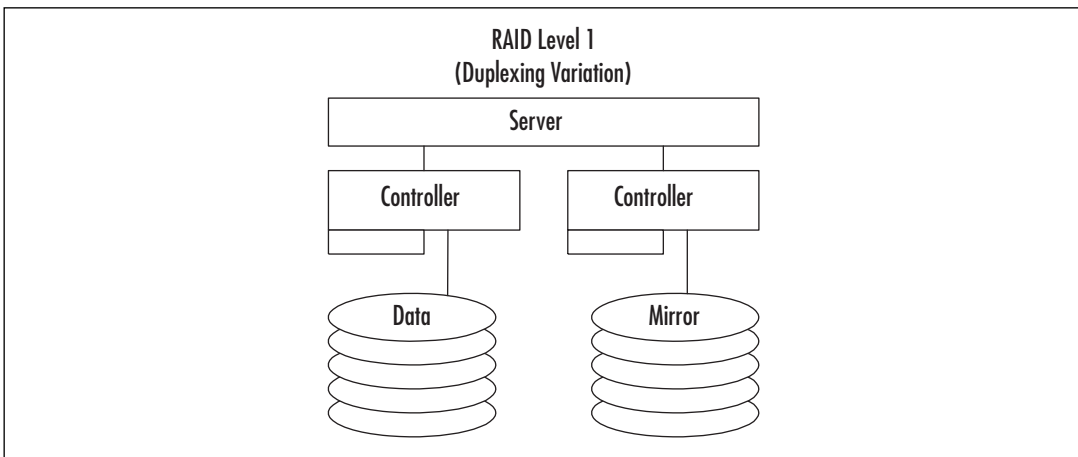


RAID 1

RAID 1 (also called *mirroring*) provides redundancy by duplicating the exact contents of one drive onto another, as illustrated in Figure 8.64. If one drive fails, the other has a complete copy of the data. This is a very reliable method of protecting data. It has a small write performance impact, because data must be written twice (once to each drive). It also has a read impact, because often information can be read from either or both drives at the same time.

Figure 8.64 RAID 1

A variation of this RAID level called *duplexing* uses a controller for each mirrored disk, as illustrated in Figure 8.65. Duplexing can improve performance while increasing fault tolerance. RAID 1 is considered expensive because one half of the total disk space is used for providing redundancy. This level is commonly used for high-value, moderate performance data like log files.

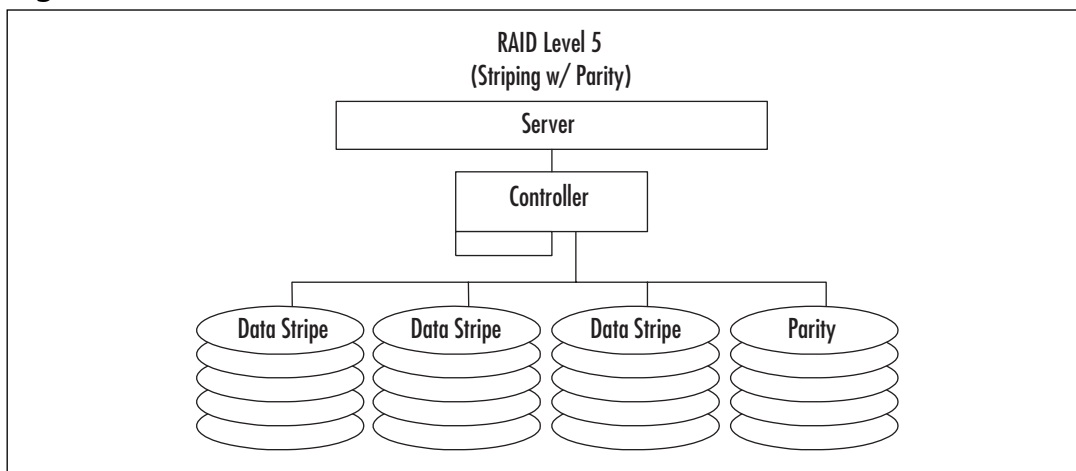
Figure 8.65 RAID 1 Duplexing Variation

RAID 5

RAID 5 (also called *striping with parity*) is perhaps the most common level of RAID. As shown in Figure 8.66, this level combines the concept of *parity* with the technique of striping. *Parity* is a mathematical value that can be used to re-create missing data. When using Windows Server 2003 software-based RAID 5, a minimum of three drives is required

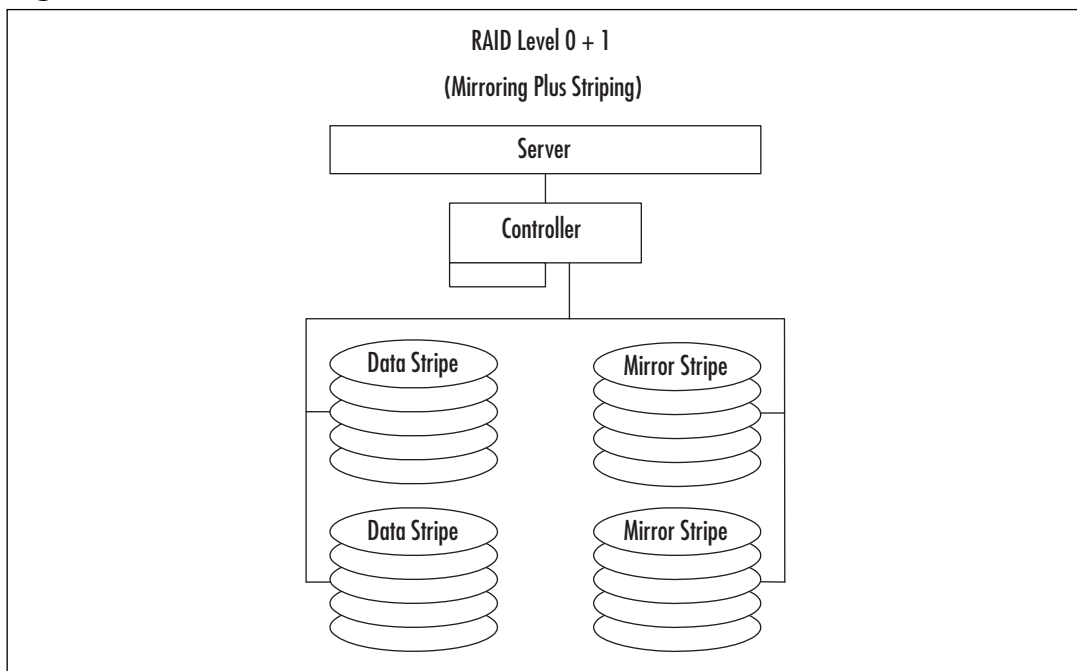
and a maximum of 32 are allowed. When a block of data is written to disk, it is split into smaller, uniform-sized blocks. An additional block of data (the parity block) is created from a mathematical calculation based on the other blocks of data. Finally, all the blocks of data are written in a stripe to the disks. If a drive fails, the original data can be re-created by doing the reverse of the parity calculation. RAID 5 exacts a performance impact for writes (the parity calculation) but yields a performance increase on reads (from striping). RAID 5 is commonly used to provide redundancy at a lower effective cost than RAID 1.

Figure 8.66 RAID 5



RAID 0+1

RAID 0+1 (also called *mirroring plus striping*), as its name implies, combines the performance and redundancy benefits of RAID levels 0 and 1. As shown in Figure 8.67, each disk in a RAID 0+1 array is mirrored, preventing the failure of any single drive (and sometime multiple drive failures) from causing downtime. In addition, each drive in a RAID 0+1 array receives only a portion of the total data load, yielding a tremendous performance increase. RAID 0+1 is the most expensive and complex RAID configuration, but it is also the highest performing. RAID 0+1 is used where the value of the data is high, the performance demand is high, and the cost considerations are secondary. Databases commonly reside on RAID 0+1 arrays.

Figure 8.67 RAID 0+1

Hot Spare Drives

Another important feature to consider is the use of *hot spare* drives. Most modern, fault-tolerant disk controllers support the use of additional hard drives that wait for an existing drive in an array to fail. When the controller determines that a drive has failed, it activates the hot spare drive and uses it to replace the failed drive in the array. The data that was present on the failed drive is re-created on the hot spare drive by the fault-tolerant controller, not Windows. In this way, an array may operate for a brief time in a decreased availability state, but will not require attention from an administrator to recover from that state.

Server Fault-Tolerance Solutions

The server is our final point of consideration for fault-tolerance. Our focus in this section is on the server system itself, not on the workload or the services it provides. There are two basic methods for introducing fault-tolerance on a server: hardware redundancy and virtualization (called *clustering*).

Modern server hardware is designed around increasing performance and reliability. Higher-end (more complicated and expensive) servers often include many built-in redundancy features. It is possible to find servers that support spare RAM and CPUs, redundant power supplies and cooling fans, built-in hardware RAID support, and many other features integrated into the basic system. In addition, many components in modern higher-end

servers are *hot-swappable*, meaning the power does not need to be turned off in order to remove or change the component.

Another hardware component that is often overlooked but is easily acquired and implemented is a redundant power source. Ideally, you want duplicate power sources all the way back to duplicate utility companies, but that is usually not possible. What is possible is the installation of an Uninterruptible Power Supply (UPS) and the software to communicate with it. A UPS is basically a large battery, although this term is sometimes also used to refer to a generator. Your equipment plugs into the UPS, and the UPS plugs into utility power. If utility power is cut, the UPS continues to power your equipment. Most often, a UPS is used to provide power long enough for a proper system shutdown. Size a UPS by the amount of power it must provide and the length of time needed to run when on battery. The more equipment on a UPS or the longer the required runtime, the “larger” the UPS must be. In very large environments, consider multiple UPSs operating in parallel (never “daisy-chain” UPSs) and possibly a backup generator.



EXAM WARNING

In addition to your servers, you may also need to have your switches and hubs, a monitor, and other equipment plugged into a UPS. Many UPS systems come with software that uses the network to notify servers that power has been lost. If the network hardware does not have power, servers will not receive these messages. As a result, they will not know they need to shut down. In addition, you may want to leave hardware plugged in that allows you to interact with the server during power outages. That will be hard to do without a monitor.

Server virtualization refers to a method used to reduce the dependence of the services provided by a server on the hardware it runs on. Server *clusters* are used for this purpose. Server clusters are discussed in Chapter 9.

Summary of Exam Objectives

Windows Server 2003 often performs reasonably well in its default configuration, but insufficient memory, CPU, disk, or network resources can reduce performance to an unacceptable level. Proper tuning and allocation of these resources will ensure adequate performance. Proper configuration of the server's page files can improve performance. Regular use of the Disk Defragmenter utility will ensure that your file systems do not become a bottleneck for read and write operations. Use efficient and intelligent network adapters to handle some of the processing load and reduce the overall impact communications have on the system.

The System Monitor utility can be used to monitor various counters present in the system. These counters display real performance information about what is occurring in the system. Some counters can display statistics as percentages, others as cumulative counts of events, and others as immediate absolute values. System Monitor can be used to view current activity in the system or to view data from log files.

A properly developed baseline can help in planning for increased growth and in identifying resources that are being overutilized. A baseline provides a mechanism for identifying what normal operating conditions are for a server. The baseline acts as a reference for troubleshooting performance issues.

The operating system and some applications record events in numerous event log files. The events in these files are always in the same format and can be viewed, searched, and monitored to determine if a system is functioning properly. Entries in the event logs indicate the severity or nature of the events. Security auditing can be enabled and security-related events captured in the event logs. The logs themselves can be archived to create a historical record of a server's activities.

Backing up data is a must to ensure system availability. Only user accounts with elevated user rights can perform backups or restores. Different methods (normal, differential, and incremental) for performing backups are available to accomplish different objectives. Backups can be performed to tape drives, network shares, or local disks, but not writable or read-writable CD-ROMs or DVD-ROMs.

Some services like DHCP, WINS, and DNS may have special considerations or configuration issues that need to be addressed before backups are performed. Clustered server disks also require special consideration for backups, and the new Volume Shadow Copy feature assists in creating backups of open files.

The Windows Backup Utility can be run either as a Wizard or in Advanced Mode. The Wizard works in most situations and steps you through the process of creating or restoring a backup. The Advanced Mode gives you access to the more powerful options of the utility and lets you fine-tune your backups. The Backup Utility also lets you schedule backup sessions, so that you can create a relatively simple and regular backup process.

The new ASR feature of Windows Server 2003 simplifies the process of re-creating a failed server installation. The ASR process replaces the older ERD process used in previous versions of Windows. Proper planning and preparation must be completed before ASR can

be used to restore a system, and performing an ASR restore should be the last resort. An ASR restore requires a floppy diskette drive to be present in the server, but one is not required to create an ASR backup.

The proper use of fault tolerance will mean that services will continue to be provided even when something breaks down. Redundancy in hardware, software, and communications ensures a reliable environment. The use of redundant network interfaces and proxy servers will ensure reliable communications. Using disk RAID arrays for the storage of applications and data will help prevent downtime due to a hard drive failure and may also be used as a performance enhancer. Using redundant components to help cool your server and provide power when the utility power fails will ensure your server operates in adverse conditions.

Exam Objectives Fast Track

Understanding Performance Bottlenecks

- ☑ RAM is the one resource that most often becomes a performance bottleneck.
- ☑ A good rule of thumb is that more RAM is better.
- ☑ Virtual memory uses hard disk space to expand the apparent memory available in the system. Performance decreases as virtual memory on the disk is heavily used.
- ☑ The processor is the brain of the computer. A computer can have multiple processors.
- ☑ Disk controller technology and the use of RAID determines how fast data can be read from or written to disk.
- ☑ Defragmenting a file system can improve read and write performance.
- ☑ Running multiple network protocols decreases overall network performance.
- ☑ Modern NICs can offload some of the communication processing overhead from the CPU to the NIC, which can increase system performance.
- ☑ The use of IPSec can greatly increase the security of information as it travels on the network. Using appropriate NICs to offload IPSec processing can improve system performance.
- ☑ Full-duplex communication is desired for servers. Switches are typically required to support full-duplex communications.
- ☑ System Monitor displays information collected by counters that let you examine the performance of your system. Counters are installed by default by the operating system and some applications.

- ☑ Baselineing is used to determine the average operating parameters of your system so that variations can be detected.
- ☑ Monitoring a large number of counters can impact system performance. Monitor only the necessary counters.
- ☑ Information about various events that occur in the system is collected in a number of event log files, which can be viewed using the Event Viewer utility.
- ☑ Event Viewer can be used to search the logs and filter out events you do not wish to examine.

Planning a Backup and Recovery Strategy

- ☑ Data backup is an essential part of a high-availability strategy.
- ☑ Many things can cause loss of data—from hardware and software problems to human factors.
- ☑ Good procedures are an indispensable part of a backup and restore strategy.
- ☑ The Windows Backup Utility is used to perform backups and restores, as well as to create ASR sets.
- ☑ Specific user rights and permissions are required to perform a backup or restore.
- ☑ Several different backup types exist, including normal (full), copy, differential, and incremental.
- ☑ Backup types are most effective when used in combination.
- ☑ Volume Shadow Copy is a new feature in Windows Server 2003 that allows the Windows Backup Utility to back up open files.
- ☑ Different applications and components of the operating system may have specific needs for either backup or restore.

Planning System Recovery with ASR

- ☑ ASR is a new feature of Windows Server 2003 that assists in the rapid re-creation of a server after a major failure.
- ☑ ASR is a last-resort option. Booting into Safe Mode and the Last Known Good mode should be tried before attempting an ASR restore.
- ☑ ASR sets are created from the Windows Backup Utility.
- ☑ ASR sets consist of a floppy disk and media containing the data on every partition or volume that contained system components.

- ☑ A floppy diskette is required to perform an ASR restore. The diskette contains files that describe the disk identities and structure of the system being restored.
- ☑ An ASR restore requires the ASR backup media, the ASR diskette, and the original Windows Server 2003 media.

Planning for Fault Tolerance

- ☑ Fault tolerance allows for components of a system to fail while the system continues to function.
- ☑ Fault tolerance is achieved through a combination of redundancy, efficient load distribution, proper planning, proper procedures, and training.
- ☑ *Five nines* refers to a system that is available 99.999 percent of the time.
- ☑ Network interfaces can be made fault tolerant by configuring multiple NICs for failover or load-balanced operation.
- ☑ Using multiple Web and proxy servers increases availability.
- ☑ Use of RAID technology can reduce or eliminate downtime caused by disk drive failure.
- ☑ Several RAID levels exist. Each RAID level is suitable for a specific type of use.
- ☑ Modern servers often have built-in redundancy, increasing their reliability.
- ☑ A UPS can prevent or reduce the downtime caused by a power failure.

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 best way to ensure a server has enough memory to operate acceptably?

A: Install more than the minimum amount of RAM.

Q: What is using part of a server's hard disk as an expansion of memory called?

A: Virtual memory.

Q: What is the file that is used for virtual memory called?

A: The paging file.

Q: What does a multiprocessing system contain?

A: More than one CPU.

Q: What are the three most common disk interfaces?

A: ATA, SCSI, and Fibre Channel.

Q: What happens when a data packet crosses a router from a network with a large packet size to a network with a smaller packet size?

A: The router re-creates the original packet as multiple smaller packets and forwards them to their destination.

Q: What devices can the Windows Backup Utility use for performing backups?

A: Any device supported by the Windows Server 2003 operating system as having removable and writable media. Tape devices are the most common.

Q: Can ASR be used to completely restore a failed system?

A: No, the purpose of ASR is to re-create the operating system. Additional restores of data and applications are required for a full system restore.

Q: Why is RAID 5 not recommended for write-intensive situations?

A: Every time data is written to a RAID 5 array, the parity block must be recalculated.

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 Performance Bottlenecks

1. You have been tasked with the implementation of enhancing the security of your network and have been allocated a modest budget to accomplish the task. You decide to implement IP Security (IPSec) between your three Windows Server 2003 servers and your Windows 2000 Professional and Windows XP Professional workstations. As the implementation proceeds, you begin hearing reports that the network does not seem as responsive. You confirm that performance has decreased. What can you do to return performance to the previous level and still accomplish your objectives?
 - A. Remove IPSec from the workstations, leaving the servers configured with IPSec.
 - B. Remove IPSec from the servers, leaving the workstations configured with IPSec.
 - C. Add NICs to your servers and configure the cards for load balancing.
 - D. Purchase new NICs that support IPSec on the NIC.
2. You have inherited the responsibility of supporting a server from a previous administrator. The system has dual 1 GHz CPUs, 2048MB of RAM, and a dual-channel caching hardware RAID controller with sixteen 18GB hard drives configured as a RAID 5 array. The system has been running an important SQL database for some time, but over the last few weeks, responsiveness has decreased as more people have been accessing the SQL databases. Your part-time SQL administrator has told you that recent database growth is not the culprit. The databases have been consistently using between 40 and 45 percent of the available disk space. You have been asked to resolve this problem. What can you do to increase the responsiveness of the SQL database?
 - A. Install more RAM in the server.
 - B. Change the RAID array to a RAID 0+1 configuration.
 - C. Change the RAID array to a RAID 0 configuration.
 - D. Increase the cache size on the array controller.

3. You have recently purchased a new single-CPU, Intel Xeon-based server. This hardware will be used to run a multithreaded CPU-intensive application. How can you ensure that the application performs at its best on the hardware provided?
 - A. Turn on hyperthreading.
 - B. Add a second CPU.
 - C. Boost the processing priority of the applications threads.
 - D. Disable hyperthreading.

4. Your server seems slow to respond to file requests from drive D: at times. You have examined the system with Performance Monitor, and the counter LogicalDisk:Current Disk Queue Length for the D: instance consistently varies between 8 and 20 during these periods of slow response. Drive D: resides on an external, 14-slot disk array with 4 slots populated with hard drives. How should you resolve this problem?
 - A. Defragment drive D:.
 - B. Add more memory to the system to increase file-caching efficiency.
 - C. Add more physical drives to the external array; either expand drive D: across the new drives or create another drive and move some heavily accessed files from drive D: to the new logical drive.
 - D. Add processors or turn on hyperthreading.

5. You have recently purchased and installed two new name-brand servers. The servers are identical in all respects, except that one server has a single CPU and the other has two. The single-CPU system will be used for basic file and print services, and the dual-CPU system will be used for running Microsoft Exchange Server. Both systems respond adequately. While developing a performance baseline, you notice that the dual-CPU system seems to be experiencing more interrupts per second than the other server. What should you do to resolve this increased level of interrupts?
 - A. Do nothing. This is a peculiarity of Microsoft Exchange Server.
 - B. Increase the communication buffers on the multiple-CPU server's NIC.
 - C. Remove the second CPU from the dual-CPU system.
 - D. Do nothing. This is normal for a multi-CPU system.

Planning a Backup and Recovery Strategy

6. You have been asked to develop a backup strategy for your company's three Windows Server 2003 servers. You have been told that the primary objective is to have the sys-

tems up and running again as quickly as possible should a disaster occur. To accomplish this goal, initial funds have been allocated and, if necessary, ongoing funds will be made available. What backup strategy should you adopt?

- A. Full backups nightly to a tape drive installed in each server
 - B. Full backups nightly to a single, centralized tape drive
 - C. Full backups weekly, with daily differential backups to a tape drive installed in each server
 - D. Full ASR backups nightly
7. You have been asked to develop a backup strategy for your company's three Windows Server 2003 servers. You have been told that the primary objective is to minimize the ongoing cost of performing backups. To accomplish this goal, you have been given a modest budget. What backup strategy should you adopt?
- A. Full backups monthly, differential backups on the weekends, and incremental backups daily to a tape drive installed in each server
 - B. Full backups monthly, differential backups on the weekends, and incremental backups daily to a single, centralized tape drive
 - C. Incremental backups daily to a single, centralized tape drive
 - D. Periodic full backups and daily incremental backups to a single, centralized tape drive
8. You have been asked to develop a backup strategy for your company's three Windows Server 2003 servers. You have been told that the primary objective is to minimize the time required for performing backups on regular business days. You do not have the use of any advanced storage technology, and an older application on the server requires you to shut down the application and disable Volume Shadow Copy to get a successful backup. To accomplish this goal, you have been given a sufficient budget. What backup strategy should you adopt?
- A. Full backups on the weekends and incremental backups daily to a tape drive installed in each server
 - B. Full backups monthly and differential backups daily to a single, centralized tape drive
 - C. Incremental backups daily to a single, centralized tape drive
 - D. Periodic full backups and daily incremental backups to a single, centralized tape drive

9. Your company uses a well-known and respected third-party backup utility for all of its servers. You are adopting Windows Server 2003 early after its release and have upgraded a number of servers to the operating system. You have high hopes about improving backup performance on some of your higher volume file servers (including the ability to back up open files) and have installed the third-party client agent software on your servers. After a few days, you notice that the speed of backups has not increased. What is the most likely reason that backup performance has not increased?
- A. Volume Shadow Copy has not been turned on for the appropriate volumes.
 - B. The third-party backup software does not use the new features present in Windows Server 2003.
 - C. An ASR backup needs to be performed before the third-party utility will show increased performance.
 - D. The drives hosting the files need to be defragmented for performance to improve.

Planning System Recovery with ASR

10. You have inherited the responsibility for supporting an important server recently upgraded from Windows NT 4 to Windows Server 2003. When the server was upgraded, it met the hardware requirements, but not by much. Increasing demand on the system has led to lower than desirable performance. Company management has authorized the purchase of new server hardware and would like you to upgrade the server as quickly as possible with the least amount of risk and additional expense. What is the best way to accomplish the upgrade in the fastest possible time, with the lowest risk, and no additional cost?
- A. Use a third-party product to duplicate the server onto the new hardware.
 - B. Create an ASR backup of the existing server. Use the ASR backup on the new hardware. Back up the existing server. Restore the backup to the new hardware.
 - C. Install Windows Server 2003 onto the new hardware. Back up the existing server. Restore the backup to the new hardware.
 - D. Shut down the existing server and move the existing hard drives to the new server. Boot the new server with the old hard drives.
11. A few weeks ago, you installed a new server. You have been performing regular full and incremental backups for all files on the system. You did not perform an initial ASR backup. When you arrived this morning, you discovered that the hard drive failed sometime last night after the backup completed, and the server will no longer boot. You replaced the failed hard drive with an identical one you had on hand. What is the quickest way to get the server back to its previous operational state?

- A. Start an ASR restore. Since the hard drive is new and identical to the failed drive, ASR will automatically re-create the previous configuration.
 - B. You cannot restore the server. It is permanently lost.
 - C. Reinstall Windows Server 2003 in a minimal configuration, restore the most recent full backup, and then restore all of the incremental backups in sequence.
 - D. Reinstall Windows Server 2003 in a minimal configuration, perform an ASR backup, perform an ASR restore, restore the most recent full backup, and then restore all of the incremental backups in sequence.
12. You are working on an existing server. The NIC manufacturer has notified you of an updated driver for your card that will greatly improve performance. You download and install the new driver. Before you reboot the system, you perform an ASR backup. When you reboot the system, it reaches the graphical portion of the boot process and presents a STOP message. What is the proper process for recovering from this problem?
- A. Perform an ASR restore from the ASR backup set you created before the reboot.
 - B. Reboot the system, press F8 when prompted during the boot process, select Last Known Good Configuration, and press Enter.
 - C. Reinstall the operating system and do a restore of the system from tape backup.
 - D. Reboot the system, press F8 when prompted during the boot process, select Safe Mode, and press Enter.

Planning for Fault Tolerance

13. You are responsible for administering a Windows Server 2003 system. The system has a Pentium III 800 MHz CPU, 1024MB of RAM, and four hard drives configured in a RAID 5 array that reside in an external seven-slot chassis. The array is controlled by a modern, high-performance hardware RAID controller and presents the array to the operating system as a single drive. You arrive on a Monday morning to find your server has crashed. On investigation, you find that two of the hard drives failed. The server has a built-in display that tells you that one drive failed late Friday night and the second drive failed Sunday afternoon. What should you have done to prevent the second drive failure from causing the server to crash?
 - A. Ensure that backups complete during business hours.
 - B. Use Volume Shadow Copy to automatically create a backup on the remaining good drive.
 - C. Install a second hardware RAID controller and distribute the drives evenly on the controllers.
 - D. Purchase another hard drive and configure it as a hot spare drive.

14. You are replacing a single-port NIC in your server with a new four-port NIC. Your switches support 100 Mbps full-duplex operation. Your switches also support either load-balancing or failover configurations. Which configuration choice is best for increased performance and availability?
 - A. Configure the card for two-way load balancing with failover to the remaining two ports.
 - B. Configure the card for four separate links to the switch. Windows Server 2003 automatically determines that the ports connect to the same switch and enables failover.
 - C. Configure the card for four-way load balancing.
 - D. Leave the old NIC in the server and add the new four-port card into an empty slot on the server. Configure the new card as a failover backup for the existing card.

15. Your data center recently experienced a utility power failure that took down all of the computer systems. Some systems experienced major problems (hard drive and fan failures) when the power was restored. Because of the failures, company management decides to install an Uninterruptible Power Supply (UPS) for the data center to protect the systems from another power failure. A few months later, another power failure hits the data center and the systems run for a time, then go down when the UPS runs out of power. This time, hard drive failures occur and data is lost. What was missed

during the implementation of the new UPS that would have prevented the second power failure from impacting the servers?

- A. Neither the proper procedures nor the automated software controls were implemented to enable a controlled shutdown.
- B. The UPS that was purchased did not have a high enough power runtime rating to handle the load of the equipment in the data center.
- C. Windows Server 2003 does not support the use of a UPS.
- D. Windows Server 2003 does not support the use of a nondedicated UPS. Each server must have a dedicated UPS.

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**

2. **B**

3. **A**

4. **C**

5. **D**

6. **A**

7. **D**

8. **A**

9. **B**

10. **B**

11. **C**

12. **B**

13. **D**

14. **C**

15. **A**

MCSE 70-293

Implementing Windows Cluster Services and Network Load Balancing

Exam Objectives in this chapter:

- 4.1.1 Plan a high availability solution that uses clustering services.
 - 4.3 Implement a cluster server.
 - 4.3.1 Recover from cluster node failure.
 - 4.1.2 Plan a high availability solution that uses Network Load Balancing.
 - 4.4 Manage Network Load Balancing. Tools might include the Network Load Balancing Monitor Microsoft Management Console (MMC) snap-in and the WLBS cluster control utility.
-
- Summary of Exam Objectives
 - Exam Objectives Fast Track
 - Exam Objectives Frequently Asked Questions
 - Self Test
 - Self Test Quick Answer Key

Introduction

Fault tolerance generally involves redundancy; for example, in the case of disk fault tolerance, multiple disks are used. The ultimate in fault tolerance is the use of multiple servers, configured to take over for one another in case of failure or to share the processing load. Windows Server 2003 provides network administrators with two powerful tools to enhance fault tolerance and high availability: server clustering (only available in the Enterprise and Datacenter Editions), and Network Load Balancing included in all editions.

This chapter looks first at server clustering and shows you how to make clustering services part of your enterprise-level organization's high-availability plan. We'll start by introducing you to the terminology and concepts involved in understanding clustering. You'll learn about cluster nodes, cluster groups, failover and failback, name resolution as it pertains to cluster services, and how server clustering works. We'll discuss three cluster models: single-node, single quorum device, and majority node set. Then we'll talk about cluster deployment options, including N-node failover pairs, hot standby server/N+1, failover ring, and random. You'll learn about cluster administration, and we'll show you how to use the Cluster Administrator tool as well as command-line tools.

Next, we'll discuss best practices for deploying server clusters. You'll learn about hardware issues, especially those related to network interface controllers, storage devices, power-saving features, and general compatibility issues. We'll discuss cluster network configuration and you'll learn about multiple interconnections and node-to-node communications. We'll talk about the importance of binding order, adapter settings, and TCP/IP settings. We'll also discuss the default cluster group. Next, we'll move onto the subject of security for server clusters. This includes physical security, public/mixed networks, private networks, secure remote administration of cluster nodes, security issues involving the cluster service account, and how to limit client access. We'll also talk about how to secure data in a cluster, how to secure disk resources, and how to secure cluster configuration log files.

The next section addresses how to make Network Load Balancing (NLB) part of your high-availability plan. We'll introduce you to NLB concepts such as hosts/default host, load weight, traffic distribution, convergence, and heartbeats. You'll learn about how NLB works and the relationship of NLB to clustering. We'll show you how to manage NLB clusters using the NLB Manager tool, remote-management tools, and command-line tools. We'll also discuss NLB error detection and handling. Next, we'll move onto monitoring NLB using the NLB Monitor Microsoft Management Console (MMC) snap-in or the Windows Load Balancing Service (WLBS) cluster control utility. We discuss best practices for implementing and managing NLB, including issues such as multiple network adapters, protocols and IP addressing, and NLB Manager logging. Finally, we'll address NLB security.

EXAM 70-293
OBJECTIVE 4.1.1

Making Server Clustering Part of Your High-Availability Plan

Certain circumstances require an application to be operational more consistently than standard hardware would allow. Databases and mail servers often have this need. What if it were possible to have more than one server ready to run the critical application? What if there were a software component that automatically managed the operation of the application so that, if one server experienced a failure, another server would automatically take over and keep the application running? Such a technology exists, and it's called *server clustering*.

The basic idea of server clustering has been around for many years on other computing platforms. Microsoft initially released its server cluster technology as part of Windows NT 4.0 Enterprise Edition. It supported two nodes and a limited number of applications. Server clustering was further refined with the release of Windows 2000 Advanced and Datacenter Server Editions. Server clusters were simpler to create, and more applications were available. In addition, some publishers began to make their applications “cluster-aware,” so that their applications installed and operated more easily on a server cluster. Now with the release of Windows Server 2003, we see another level of improvement on the server clustering technology. Server clusters now support much larger clusters and more robust configurations. Server clusters are easier to create and manage. Features that were available only in the Datacenter Edition of Windows 2000 have now been made available in the Enterprise Edition of Windows Server 2003.

Terminology and Concepts

Although it has been used previously, a more formal definition of a server cluster is needed. For our purposes, a *server cluster* is a group of independent servers that work together to increase application availability to client systems and appear to clients under one common name. The independent servers that make up a server cluster are individually called *nodes*. Nodes in a server cluster monitor each other's status through a communication mechanism called a *heartbeat*. The heartbeat is a series of messages that allow the server cluster nodes to detect communication failures and, if necessary, perform a *failover* operation. A failover is the process by which resources are stopped on one node and started on another.

Cluster Nodes

A server cluster node is an independent server. This server must be running Windows 2000 Advanced Server, Windows 2000 Datacenter Server, Windows Server 2003 Enterprise Edition, or Windows Server 2003 Datacenter Edition. The two editions of Windows Server 2003 cannot be used in the same server cluster, but either can exist in a server cluster with a Windows 2000 Advanced Server node. Since Windows Server 2003 Datacenter Edition is available only through original equipment manufacturers (OEMs), this chapter deals with server clusters constructed with the Advanced Server Edition of Windows Server 2003 unless specifically stated otherwise.

A server cluster node should be a robust system. When designing your server cluster, do not overlook applying fault-tolerant concepts to the individual nodes. Using individual fault-tolerant components to build fault-tolerant nodes to build fault-tolerant server clusters can be described as “fault tolerance in depth.” This approach will increase overall reliability and make your life easier.

A server cluster consists of anywhere between one and eight nodes. These nodes do not necessarily need to have identical configurations, although that is a frequent design element. Each node in a server cluster can be configured to have a primary role that is different from the other nodes in the server cluster. This allows you to have better overall utilization of the server cluster if each node is actively providing services. A node is connected to one or more storage devices, which contain disks that house information about the server cluster. Each node also contains one or more separate network interfaces that provide client communications and support heartbeat communications.

Cluster Groups

The smallest unit of service that a server cluster can provide is a *resource*. A resource is a physical or logical component that can be managed on an individual basis and can be independently activated or deactivated (called bringing the resource *online* or *offline*). A resource can be owned by only one node at a time.

There are several predefined (called “standard”) types of resources known to Windows Server 2003. Each type is used for a specific purpose. The following are some of the most common standard resource types:

- **Physical Disk** Represents and manages disks present on a shared cluster storage device. Can be partitioned like a regular disk. Can be assigned a drive letter or used as an NTFS mounted drive.
- **IP Address** Manages an IP address.
- **Network Name** Manages a unique NetBIOS name on the network, separate from the NetBIOS name of the node on which the resource is running.
- **Generic Service** Manages a Windows operating system service as a cluster resource. Helps ensure that the service operates in one place at one time.
- **Generic Script** Manages a script as a cluster resource (new to Windows Server 2003).
- **File Share** Creates and manages a Windows file share as a cluster resource.

Other standard resource types allow you to manage clustered print servers, Dynamic Host Configuration Protocol (DHCP) servers, Windows Internet Name Service (WINS) servers, and generic noncluster-aware applications. (It is also possible to create new resource types through the use of dynamic link library files.)

Individual resources are combined to form *cluster groups*. A cluster group is a collection of server resources that defines the relationships of resource within the group to each other

and defines the unit of failover, so that if one resource moves between nodes, all resources in the group also move. As with individual resources, a cluster group can be owned by only one node at a time. To use an analogy from chemistry, resources are atoms and groups are compounds. The cluster group is the primary unit of administration in a server cluster. Similar or interdependent resources are combined into the same group. A resource cannot be dependent on another resource that is not in the same cluster group. Most cluster groups are designed around either an application or a storage unit. It is in this way that individual applications or disks in a server cluster are controlled independently of other applications or disks.

Failover and Failback

If a resource on a node fails, the cluster service will first attempt to reactivate the resource on the same node. If unable to do so, the cluster service will move the cluster group to another node in the server cluster. This process is called a *failover*. A failover can be triggered manually by the administrator or automatically by a node failure. A failover can involve multiple nodes if the server cluster is configured this way, and each group can have different failover policies defined.

A *failback* is the corollary of a failover. When the original node that hosted the failed-over resource(s) comes back online, the cluster service can return the cluster group to operation on the original node. This failback policy can be defined individually for a cluster group or disabled entirely. Failback is usually performed at times of low utilization to avoid impacting clients, and it can be set to follow specific schedules.

Cluster Services and Name Resolution

A server cluster appears to clients as one common name, regardless of the number of nodes in the server cluster. It is for this reason that the server cluster name must be unique on your network. Ensure that the server cluster name is different from the names of other server clusters, domain names, servers, and workstations on your network. The server cluster will register its name with the WINS and DNS servers configured on the node running the default cluster group.

Individual applications that run on a server cluster can (and should) be configured to run in separate cluster groups. The applications must also have unique names on the network and will also automatically register with WINS and DNS. Do not use static WINS entries for your resources. Doing so will prevent an update to the WINS registered address in the event of a failover.

How Clustering Works

Each node in a server cluster is connected to one or more storage devices. These storage devices contain one or more disks. If the server cluster contains two nodes, you can use either a SCSI interface to the storage devices or a Fibre Channel interface. For three or more node server clusters, Fibre Channel is recommended. If you are using a 64-bit edition

of Windows Server 2003, Fibre Channel is the required interface, regardless of the number of nodes.

Fibre Channel has many benefits over SCSI. Fibre Channel is faster and easily expands beyond two nodes. Fibre Channel cabling is simpler, and Fibre Channel automatically configures itself. However, Fibre Channel is also more expensive than SCSI, requires more components, and can be more complicated to design and manage.

On any server cluster, there is something called the *quorum resource*. The quorum resource is used to determine the state of the server cluster. The node that controls the quorum resource controls the server cluster, and only one node at a time can own the quorum resource. This prevents a situation called *split-brain*, which occurs when more than one node believes it controls the server cluster and behaves accordingly. Split-brain was a problem that occurred in the early development of server cluster technologies. The introduction of the quorum resource solved this problem.

Cluster Models

There are three basic server cluster design models available to choose from: single node, single quorum, and majority node set. Each is designed to fit a specific set of circumstances. Before you begin designing your server cluster, make sure you have a thorough understanding of these models.




EXAM WARNING

Make sure you understand the differences between the server cluster models and the circumstances in which each is normally used.

Single Node

A single-node server cluster model is primarily used for development and testing purposes. As its name implies, it consists of one node. An external disk resource may or may not be present. If an external disk resource is not present, the local disk is configured as the cluster storage device, and the server cluster configuration is kept there.

Failover is not possible with this server cluster model, because there is only one node. However, as with any server cluster model, it is possible to create multiple *virtual servers*. (A virtual server is a cluster group that contains its own dedicated IP address, network name, and services and is indistinguishable from other servers from a client's perspective.) Figure 9.1 illustrates the structure of a single-node server cluster.

Figure 9.1 Single Node Server Cluster

If a resource fails, the cluster service will attempt to automatically restart any applications and dependent resources. This can be useful when applied to applications that do not have built-in restart capabilities but would benefit from that capability.

Some applications that are designed for use on server clusters will not work on a single-node cluster model. Microsoft SQL Server and Microsoft Exchange Server are two examples. Applications like these require the use of one of the other two server cluster models.


Single Quorum Device

The single quorum device server cluster model is the most common and will likely continue to be the most heavily used. It has been around since Microsoft first introduced its server clustering technology.

This type of server cluster contains two or more nodes, and each node is connected to the cluster storage devices. There is a single quorum device (a physical disk) that resides on the cluster storage device. There is a single copy of the cluster configuration and operational state, which is stored on the quorum resource.

Each node in the server cluster can be configured to run different applications or to act simply as a hot-standby device waiting for a failover to occur. Figure 9.2 illustrates the structure of a single quorum device server cluster with two nodes.


Figure 9.2 Single Quorum Device Server Cluster



Majority Node Set

The majority node set (MNS) model is new in Windows Server 2003. Each node in the server cluster may or may not be connected to a shared cluster storage device. Each node maintains its own copy of the server cluster configuration data, and the cluster service is responsible for ensuring that this configuration data remains consistent across all nodes. Synchronization of quorum data occurs over Server Message Block (SMB) file shares. This communication is unencrypted. Figure 9.3 illustrates the structure of the MNS model.

Figure 9.3 A Majority Node Set Server Cluster




This model is normally used as part of an OEM predesigned or preconfigured configuration. It has the ability to support geographically distributed server clusters. When used in geographically dispersed configurations, network latency becomes an issue. You must ensure that the round-trip network latency is a *maximum* of 500 milliseconds (ms), or you will experience availability problems.

The behavior of an MNS server cluster differs from that of a single quorum device server cluster. In a single quorum device server cluster, one node can fail and the server cluster can still function. This is not necessarily the case in an MNS cluster. To avoid split-brain, a majority of the nodes must be active and available for the server cluster to function. In essence, this means that 50 percent plus 1 of the nodes must be operational at all times for the server cluster to remain operational. Table 9.1 illustrates this relationship.

Table 9.1 Majority Node Set Server Cluster Failure Tolerance

Number of Nodes in MNS Server Cluster	Maximum Node Failures before Complete Cluster Failure	Nodes Required to Continue Cluster Operations
1	0	1
2	0	2
3	1	2
4	1	3
5	2	3
6	2	4
7	3	4
8	3	5



Server Cluster Deployment Options

When you use either the single quorum device model or MNS model, there are a variety of ways that you can configure your clustered applications to act during a failover operation. The choices vary with the number of nodes in your server cluster, and each has advantages and disadvantages.

These deployment options are not always mutually exclusive. In a server cluster with several nodes and multiple cluster groups, it is possible that some groups will use one deployment option while other groups use a different one. Consider these options carefully when you design larger server clusters.



EXAM WARNING

Expect questions related to the cluster deployment options. A good understanding of each deployment option, how the options are configured, and the advantages/disadvantages of each will help you on the exam.

N-Node Failover Pairs

The N-node failover pairs deployment option specifies that two nodes, and only two nodes, may run the application. This is the simplest option and is, in essence, the only option available in a two-node server cluster. If configured in a larger server cluster with three or more nodes, the application will not be able to function if both nodes are not operational. In larger server clusters made up of nodes with different processing capabilities or capacities, you can use this option to limit an application to running on only the nodes capable of adequately servicing the application.

An N-node failover pair is configured by specifying the two nodes in the Possible Owners property for the cluster resource, as shown in Figure 9.4. You can set the Possible Owners property using the server cluster administrative tools described in the “Server Cluster Administration” section later in this chapter. Every cluster resource has a Possible Owners property that can be configured or left blank.

Figure 9.4 Setting the Possible Owners Property

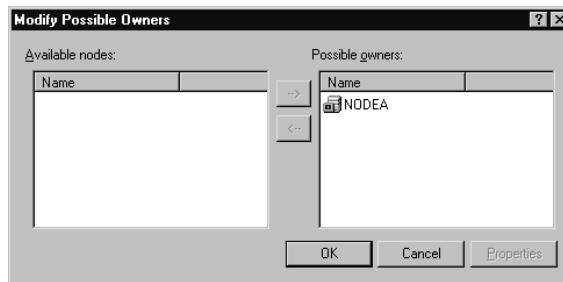


Figure 9.5 illustrates an N-node failover configuration in a server cluster with four nodes—A, B, C and D—in its normal operational state. Nodes A and B are configured as a failover pair, and nodes C and D are also a failover pair. Assorted virtual servers are active and are spread among the nodes.

Figure 9.5 N-Node Failover, Initial State

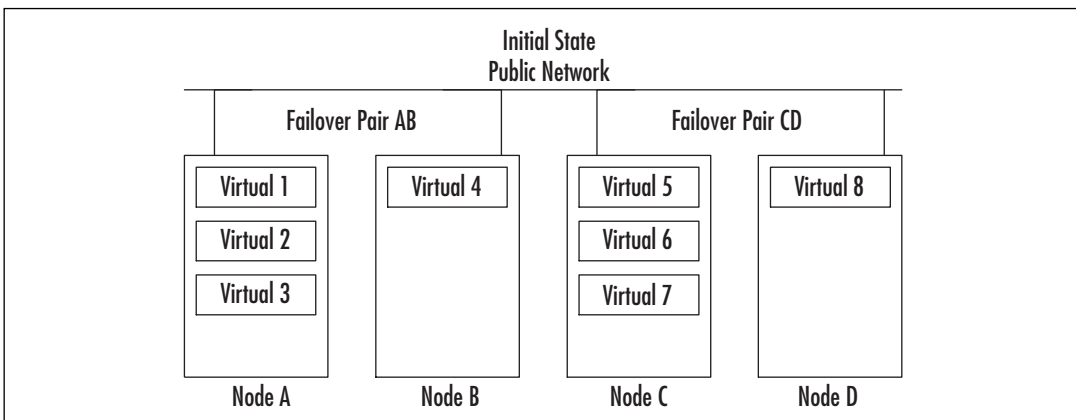
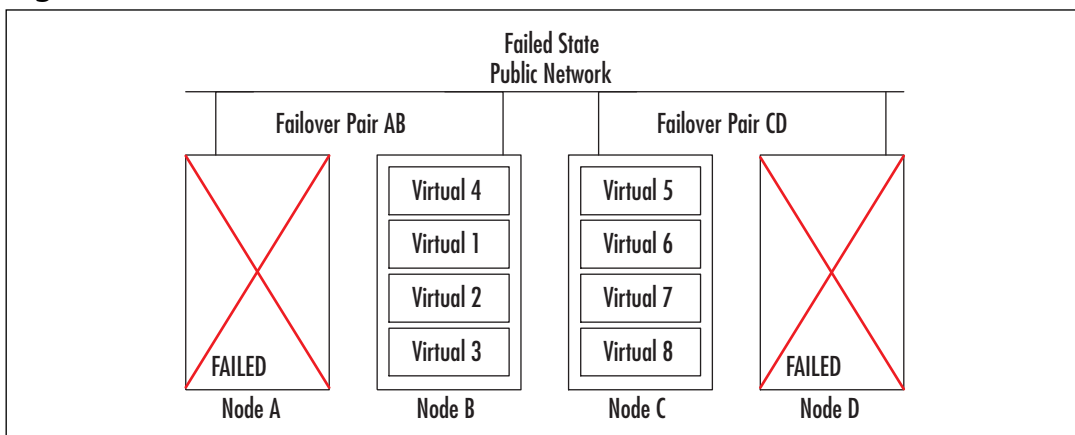


Figure 9.6 shows the same server cluster as Figure 9.5, but after two of the nodes failed. As you can see, node B has taken ownership of the virtual servers that were operating on its failover partner (node A). Node C has also taken ownership of node D's virtual servers.

Figure 9.6 N-Node Failover, Failed State



Note that Figures 9.5 and 9.6 depict a single quorum device server cluster. An MNS server cluster with four nodes could not operate with failed two nodes. The storage devices and Interconnects have been removed from the images for clarity.

Hot-Standby Server/N+I

The hot-standby server/N+1 deployment option is possible on server clusters with two or more nodes and is sometimes referred to as an *active/passive* design. In this design, you specify one node in the server cluster as a *hot spare*. This hot-spare node is normally idle or lightly loaded. It acts as the failover destination for other nodes in the cluster.

The main advantage of this option is cost savings. If a two-node server cluster is configured with one node running the application(s) (the *N* or *active* node) and one node standing idle, waiting for a failover (the *I* or *passive* node), the overhead cost in hardware is 50 percent. In an eight-node server cluster with seven *N* (active) nodes and one *I* (passive) node, the overhead cost is about 15 percent.

This option is not limited to a single hot-spare node. An eight-node server cluster could be configured with one *N* node and seven *I* nodes or any other possible combination. In these configurations, the overhead cost savings would be quite a bit less or nonexistent.

Configure this option by setting the Preferred Owners property of the group to the *N* node(s), as shown in Figure 9.7, and the Possible Owners of the resources to the *N* and *I* nodes. As mentioned earlier, the Possible Owners property is a property of the individual resource. The Preferred Owner property, however, applies only to cluster groups. Both the Possible Owners and Preferred Owners properties are configured via the server cluster

administrative tools, which are covered in the “Server Cluster Administration” section later in this chapter.

Figure 9.7 Setting the Preferred Owners Property

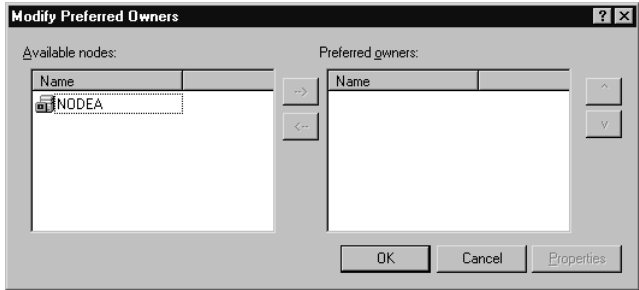


Figure 9.8 illustrates a four-node server cluster configured with three active (N) nodes and one passive (I) node in its normal operational state. Each active node supports various virtual servers.

Figure 9.8 Hot-Standby/N+I Configuration, Initial State

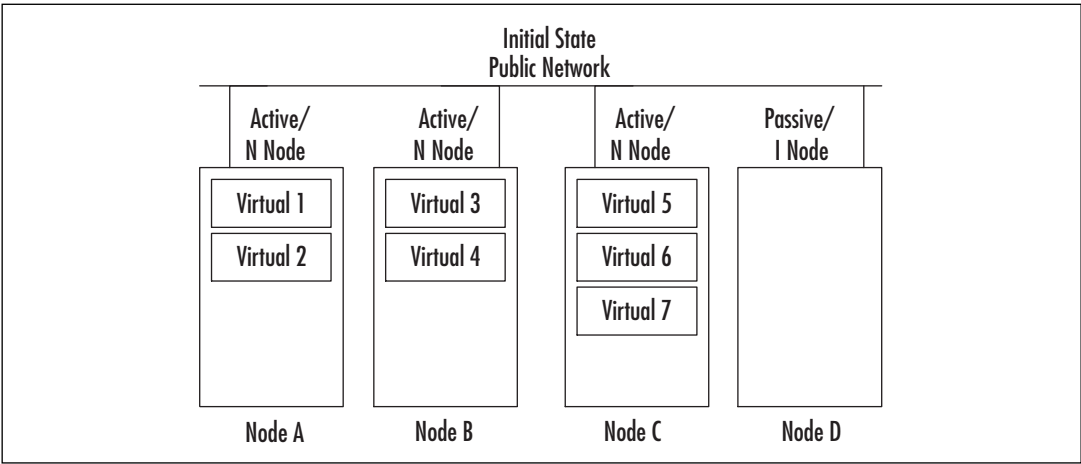
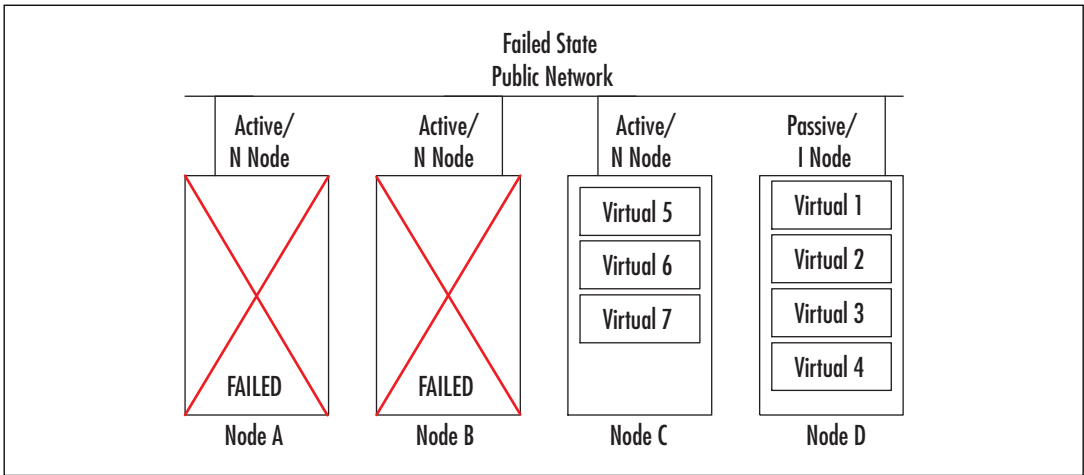


Figure 9.9 shows the same server cluster as Figure 9.8, but after the failure of two of the nodes. The virtual servers that were operating on the failed nodes have failed over to the I node. Again, if this were an MNS server cluster, there would not be enough nodes operating to support the server cluster. The MNS cluster would have failed when the second node failed, but the virtual servers from the first node would have been successfully failed over to the I node. Again, note that the storage devices and Interconnects have been removed from both images.

Figure 9.9 Hot Standby/N+I Configuration, Failed State



Failover Ring

A *failover ring* is mainly used when all nodes in a server cluster are active. When a failover occurs, applications are moved to the next node in line. This mode is possible if all nodes in the server cluster have enough excess capacity to support additional applications beyond what they normally run. If a node is operating at peak utilization, a failover to that node may reduce performance for all applications running on that node after the failover.

The order of failover is defined by the order the nodes appear in the Preferred Owner list (see Figure 9.7). The default node for the application is listed first. A failover will attempt to move the cluster group to each node on the list, in order, until the group successfully starts.

It is possible to limit the size of the failover ring by not specifying all the cluster nodes on the Preferred Owner list. In effect, this combines the N+I and failover ring options to produce a hybrid option. This hybrid option reduces the N+I overhead cost to zero, but you need to make sure that enough capacity is present to support your applications.

Figure 9.10 illustrates an eight-node server cluster in a failover ring configuration in its initial state. This server cluster is operating with eight nodes. To simplify the diagram, each node is running one virtual server. (The configuration of the failover ring in this scenario is very simple: each node fails over to the next node, with the last node set to fail over to the first, and so on.) Storage devices and Interconnects have been removed for clarity.

Figure 9.10 Failover Ring Configuration, Initial State


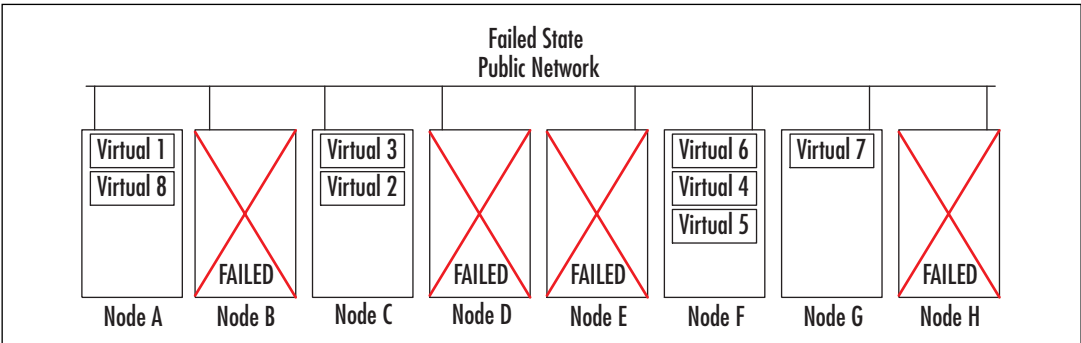


Figure 9.11 illustrates the failover ring configuration after the server cluster has experienced a failure of half of its nodes. Notice how node F has picked up the virtual servers from nodes D and E, and how node A picked up the virtual server from node H. Again, if this were an MNS server cluster, there would not be enough nodes left operational for the server cluster to function. As usual, storage devices and Interconnects have been removed from the image for clarity.

Figure 9.11 Failover Ring Configuration, Failed State



Random

The random deployment option makes the cluster service determine the destination for a failover. This option is used in large server clusters where each node is active and it is difficult to specify an order of failover because of the needs and complexity of the environment. When adopting this option, it is important to make sure that each node has sufficient excess capacity to handle additional load. Otherwise, a failover may reduce performance for applications running on a node that is at or near peak capacity.

This mode is configured by not defining a Preferred Owner for the resource group. The cluster service will attempt to determine a suitable node for the application in the event of a failover. Figure 9.12 illustrates a random failover configuration in the initial state.

It shows a server cluster of eight nodes supporting two virtual servers, each in its normal operating mode.

Figure 9.12 Random Configuration, Initial State

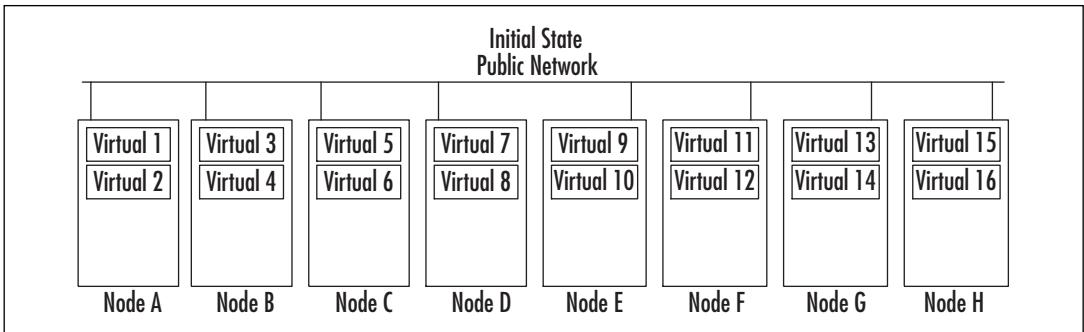
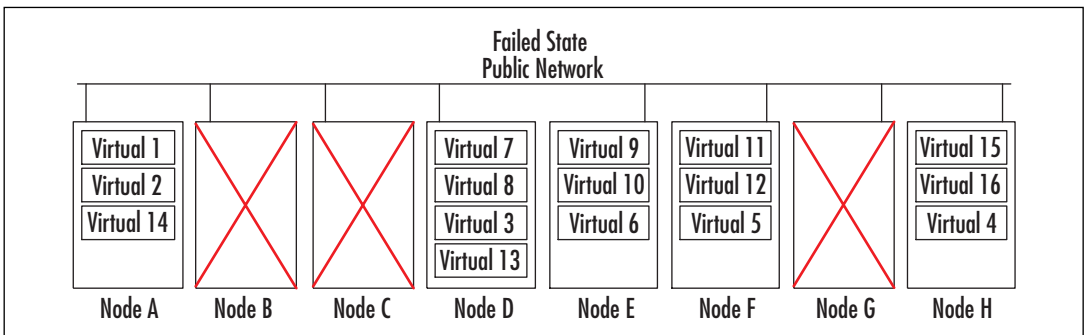


Figure 9.13 shows the same configuration after this server cluster has experienced a failure of three of its nodes. Notice how the virtual servers have been distributed seemingly at random to the surviving nodes. If this were an MNS server cluster, it would still be functioning.

Figure 9.13 Random Configuration, Failed State




Server Cluster Administration


After a server cluster is operational, it must be administered. There are two tools provided to you to accomplish this: Cluster Administrator, an interactive graphical utility, and Cluster.exe, provided for use at the command line and in scripts or batch files.

Using the Cluster Administrator Tool

To access the Cluster Administrator utility, select **Start | Administrative Tools | Cluster Administrator**. The Cluster Administrator utility, shown in Figure 9.14, allows you to create a new server cluster, add nodes to an existing server cluster, and perform administrative tasks on a server cluster.

Figure 9.14 The Cluster Administrator Window

At the **Open Connection to Cluster** dialog box, shown in Figure 9.15, you can enter the name of a server cluster or browse for it.

Figure 9.15 The Open Connection Dialog Box

If you wish to create a new server cluster, select **Create new cluster** in the **Action** drop-down list box and click **OK**. This will start the New Server Cluster Wizard, which will step you through the process of creating a new server cluster. Selecting **Add nodes to cluster** in the **Action** drop-down list will start the Add Nodes Wizard. This Wizard lets you add nodes to an existing server cluster.

Using Command-Line Tools

Cluster.exe is the command-line utility you can use to create or administer a server cluster. It has all of the capabilities of the Cluster Administrator graphical utility and more.

Cluster.exe has numerous options. Figure 9.16 shows the syntax of the **cluster.exe** command and the options you can use with it.

Figure 9.16 Cluster.exe Command Options

```

CLUSTER /LIST[:domain-name]

CLUSTER /CHANGEPASS[WORD] /?
CLUSTER /CHANGEPASS[WORD] /HELP
CLUSTER /CLUSTER:clusternam1[,clusternam2[,...]]
    /CHANGEPASS[WORD][:newpassword[,oldpassword]] <options>

<options> =
    [/FORCE] [/QUIET] [/SKIPDC] [/TEST] [/VERB[OSE]] [/UNATTEND[ED]] [/?] [/HELP]

CLUSTER [/CLUSTER:]cluster-name <options>

<options> =
    /CREATE [/NODE:node-name] [/VERB[OSE]] [/UNATTEND[ED]] [/MIN[IMUM]]
        /USER:domain\username | username@domain [/PASS[WORD]:password]
        /IPADDR[ESS]:xxx.xxx.xxx.xxx[,xxx.xxx.xxx.xxx,network-connection-
            name]
    /ADD[NODES][:node-name[,node-name ...]] [/VERB[OSE]] [/UNATTEND[ED]]
        [/MIN[IMUM]] [/PASSWORD:service-account-password]

CLUSTER [[/CLUSTER:]cluster-name] <options>

<options> =
    /CREATE [/NODE:node-name] /WIZ[ARD] [/MIN[IMUM]]
        [/USER:domain\username | username@domain] [/PASS[WORD]:password]
        [/IPADDR[ESS]:xxx.xxx.xxx.xxx]
    /ADD[NODES][:node-name[,node-name ...]] /WIZ[ARD] [/MIN[IMUM]]
        [/PASSWORD:service-account-password]
    /PROP[ERTIES] [<prop-list>]
    /PRIV[PROPERTIES] [<prop-list>]
    /PROP[ERTIES][:propname[,propname ...]] /USEDEFAULT]
    /PRIV[PROPERTIES][:propname[,propname ...]] /USEDEFAULT]
    /REN[AME]:cluster-name
    /QUORUM[RESOURCE][:resource-name] [/PATH:path] [/MAXLOGSIZE:max-size-
        kbytes]

```

Continued

Figure 9.16 Cluster.exe Command Options

```

/SETFAIL[UREACTIONS][:node-name[,node-name ...]]
/LISTNETPRI[ORITY]
/SETNETPRI[ORITY]:net[,net ...]
/REG[ADMIN]EXT:admin-extension-dll[,admin-extension-dll ...]
/UNREG[ADMIN]EXT:admin-extension-dll[,admin-extension-dll ...]
/VER[SION]
NODE [node-name] node-command
GROUP [group-name] group-command
RES[OURCE] [resource-name] resource-command
{RESOURCETYPE|RESTYPE} [resourcetype-name] resourcetype-command
NET[WORK] [network-name] network-command
NETINT[ERFACE] [interface-name] interface-command

<prop-list> =
  name=value[,value ...][:<format>] [name=value[,value ...][:<format>
    ] ...]

<format> =
  BINARY|DWORD|STR[ING]|EXPANDSTR[ING]|MULTISTR[ING]|SECURITY|ULARGE

CLUSTER /?
CLUSTER /HELP

```

Note: With the /CREATE, /ADDNODES, and /CHANGEPASSWORD options, you will be prompted for passwords not provided on the command line unless you also specify the /UNATTENDED option.

The following are some of the tasks that are impossible to do with Cluster Administrator or are easier to perform with Cluster.exe:

- Changing the password on the cluster service account
- Creating a server cluster or adding a node to a server cluster from a script
- Creating a server cluster as part of an unattended setup of Windows Server 2003
- Performing operations on multiple server clusters at the same time

EXAM
70-293
OBJECTIVE
4.3.2

Recovering from Cluster Node Failure

It is reasonable to assume that on any server cluster, you will have a component failure or need to take part of the server cluster offline for service. A properly designed and maintained server cluster should have no problems. But what if something causes the node to fail? For example, if a local hard disk in the node crashes, how do you recover?

Many of the same basic administrative tasks performed on nonclustered servers apply to clustered ones. Following the same practices will help prevent unplanned downtime and assist in restoring service when service is lost:

- **Have good documentation** Proper and complete documentation is the greatest asset you can have when trying to restore service. Configuration and contact information should also be included in your documentation.
- **Perform regular backups and periodically test restores** Clusters need to be backed up just like any other computer system. Periodically testing a restore will help keep the process fresh and help protect against hardware, media, and some software failures.
- **Perform Automated System Recovery (ASR) backups** When performing an ASR backup on your server cluster, make sure that one node owns the quorum resource during the ASR backup. If you need an ASR restore, this will be a critical component.
- **Develop performance baselines** A performance baseline should be developed for each node and the server cluster as a whole. This will help you determine if your server cluster is not performing properly or is being outgrown.

If a node experiences a failure, any groups that were on the failed node should be moved to another node (unless you are using the single-node model). You should then repair the failed components in the node in the same way as you would repair any computer system.

If repairing the node involves the replacement of the boot and/or system drives, you may need to do an ASR restore. As a precaution, you should physically disconnect the node from the cluster's shared storage devices first. Once the ASR restore is complete, shut down the node, reconnect it to the shared storage devices, and boot the node.

Server Clustering Best Practices

There are many ways to accomplish the setup and operation of a server cluster, but some methods are more reliable than others. Microsoft has published a number of “Best Practices” documents relating to its products and technologies, and server clusters are no exception.

Preparation Is Key

One of the great secrets of successfully building server clusters is extensive preparation. This can (and probably will) be tedious and time-consuming, but it will make your installation much more likely to succeed.

In addition to a good design and thorough documentation, appropriate hardware preparation is critical. Ensure that all the hardware components work correctly and have their firmware versions up-to-date. If you are using identically configured nodes, make sure that they are installed identically, even down to the slot the expansion cards sit in.

As an example, I was once required to create four clustered configurations in four days. Assembly and configuration of the hardware and updating firmware took three and half days. This was time well spent, because the actual installation of the cluster services and software took three hours to complete on all four server clusters.

Hardware Issues

The foundation of your server cluster is the hardware. It is critical to build reliable nodes at the hardware level. You cannot build high availability from unreliable or unknown components.

Compatibility List

Microsoft's position since it first began publishing cluster technology is that the hardware components used in a server cluster *and* the entire server cluster configuration itself must be listed on the Hardware Compatibility List (HCL) in order to receive support. With the introduction of Windows XP, Microsoft changed from the HCL to the Windows Catalog. Windows Server 2003-compatible hardware is listed in the Windows Server Catalog, but the concept and support requirements remain the same as they were with the HCL.

In order to receive technical support from Microsoft, ensure that your entire hardware configuration is listed as compatible in the Windows Server Catalog. Using unlisted hardware does not mean you cannot make the hardware work; it simply means that you cannot call Microsoft for help if the need arises.

Network Interface Controllers

A server cluster requires at least two network interfaces to function: one for the public network (where client requests come from) and one for the private interconnect network (for the heartbeat). Since a single private interconnect would present a single point of failure, it is good practice to have at least two interconnects. Do not use a *teamed configuration* with interconnects. A teamed configuration binds two or more physical interfaces together into one logical interface. Using teamed controllers preserves the single point of failure.

Network controllers should be identical. This includes not only the manufacturer and model, but also the firmware and drivers. Using identical controllers will also simplify the design of your server cluster and make troubleshooting easier.

Change the default name of each network interface to a descriptive name. Use Heartbeat, Interconnect, or Private for the interconnect interface. Similarly, use Public, Primary, or some similar name for the public interfaces. You should configure these names identically on each node. Following this procedure will make identifying and troubleshooting network issues much easier.

Storage Devices

No single resource in a server cluster requires more planning and preparation than shared storage. Poor planning can make management tasks quite difficult. Planning cluster disk resources requires attention to numerous details.

First, thorough planning must be done for the acquisition of the shared disk hardware. Develop capacity requirements and design disk layouts. Dynamic disks, volume sets, remote storage, removable storage, and software-based RAID cannot be used for shared cluster disks. Plan on using hardware RAID, and purchase extra hard disks for use as RAID hot spares.

If a single RAID controller is part of the design (likely in a single-node cluster), make sure that you keep an identical spare RAID controller on hand. The spare should be the exact brand and model and have the same firmware version as your production RAID controller.

If you are using Fibre Channel-based controllers, consider using multiple Fibre Channel host bus adapters (HBAs) configured in either a load-balanced or failover configuration. This will increase the cost of the cluster, but fault-tolerance will also increase. Before purchasing redundant HBAs, make sure that they are of the same brand, model, and firmware version. Also, ensure that the hardware vendor includes any necessary drivers or software to support the redundant HBA configuration.

If you are using SCSI-based controllers, ensure that each SCSI adapter on the shared storage bus is configured with a different SCSI ID. Also ensure that the shared SCSI bus is properly terminated. If either of these tasks is not done properly, data could be lost, hardware could be damaged, or the second cluster node may not properly join the cluster.

Use caution with write caching of shared disks. If power fails or a failover occurs before data is completely written to disk, data can be corrupted or lost. Disable write caching in **Device Manager** by clearing the **Enable write caching on the disk** check box on the **Policies** tab in the **Properties** of the drive, shown in Figures 9.17 and 9.18. If the RAID controller supports write caching, either disable the feature or ensure that battery backup for the cache or an alternate power supply for the controller is available.

Figure 9.17 Accessing Disk Drive Properties in Device Manager


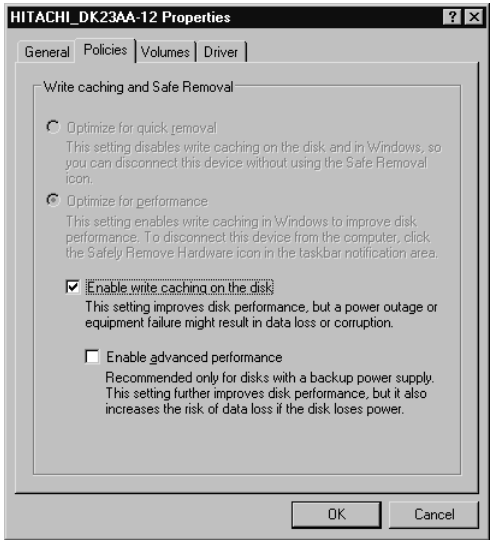


Figure 9.18 Disabling Write Caching on a Drive through Device Manager

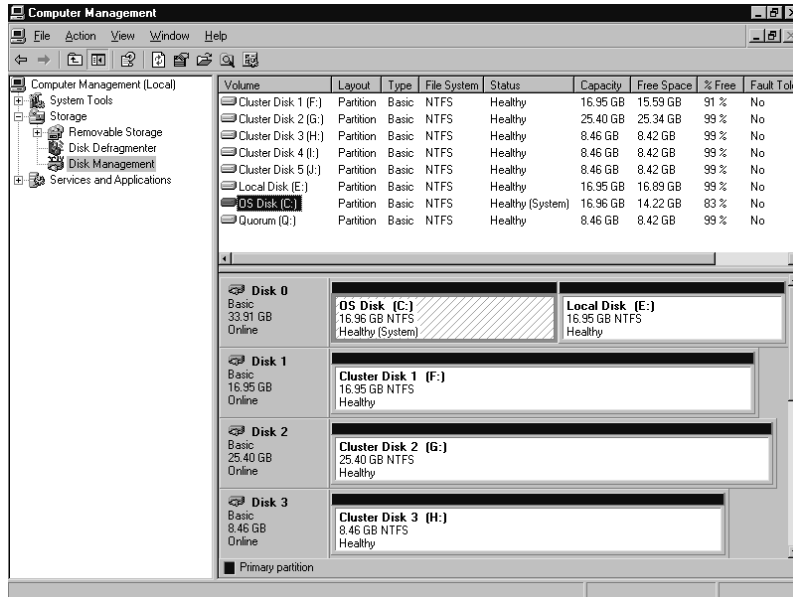


When starting the installation of the first node, ensure that the first node is the *only* node on the shared storage bus. This must be done to properly partition and format the drives in the shared storage. Until the cluster service is installed, other nodes can access the shared disks and cause data corruption.

If you are using a sophisticated disk system for shared cluster storage, use the features of the system to create logical drives that your nodes will access. This step is necessary because the *disk* is the smallest unit of storage that is recognized as a cluster resource. All of the partitions on a disk move with the disk between cluster nodes.

Once the first node is booted, format your shared drives. Only the NTFS file system is supported on clustered disks. The quorum drive should be created first. A minimum of 500MB should be assigned to the quorum drive, and no applications should reside on it. Partition and format the rest of your clustered drives as planned. Assign drive letters as you normally would, as shown in Figure 9.19, and document them. You can assign any drive letters that are not already in use, but it is a good idea to adopt the convention of assigning the quorum drive the same drive letter each time you create a cluster—Q (for quorum) is a good choice. Once you have assigned drive letters, you will need to match these drive-letter assignments on each node in the cluster.

Figure 9.19 Configuring Clustered Disks in Disk Management



In addition to drive-letter assignments, you also have the option of using NTFS mounted drives. A mounted drive does not use a drive letter, but appears as a folder on an existing drive. Mounted drives on clustered storage must be mounted to a drive residing on shared storage in the same cluster group and are dependent on this “root” disk.

Planning sufficient allocation of disk space for your applications is critical. Since you cannot use dynamic disks on shared storage without using third-party tools, it is difficult to increase the size of clustered disks. Be sure to allow for data growth when initially sizing your partitions. This is a situation where it is better to allocate a few megabytes too many than a few kilobytes too few.

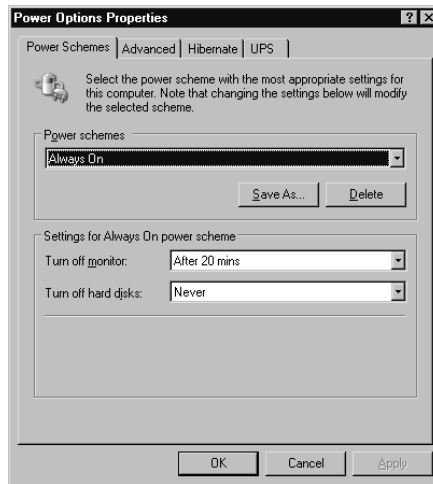
If you plan on using the *generic script* resource type, make sure the script file resides on a local disk, not a shared disk. It is possible for errant scripts to be the cause of a failover, and if a script resides on a clustered disk, the script “disappears” from under the node executing it. By keeping the scripts on a local disk, they remain available to the node at all times, and the appropriate error-checking logic can be used when errors are encountered.

Power-Saving Features

Windows Server 2003 includes power-management features that allow you to reduce the power consumed by your servers. This is very useful on laptop computers and some small servers, but can cause serious problems if used on clustered servers. If more than one node were to enter a standby or hibernation state, the server cluster could fail.

The power-saving options in Windows Server 2003 must be disabled for server clusters. Nodes should be configured to use the **Always On** power scheme, as shown in Figure 9.20. To access this option, select **Start | Control Panel | Power Options**. Using this power scheme will prevent the system from shutting down its hard drives and attempting to enter a standby or hibernation state.

Figure 9.20 Enabling the Always On Power Scheme



Cluster Network Configuration

Communications are a critical part of server cluster operations. Nodes must communicate with each other directly over the interconnects in order to determine each other’s health and, if necessary, initiate a failover. Nodes must also communicate with client systems over the public network to provide services. Both of these networks require proper planning.

When referring to server clusters, there are four types of networks:

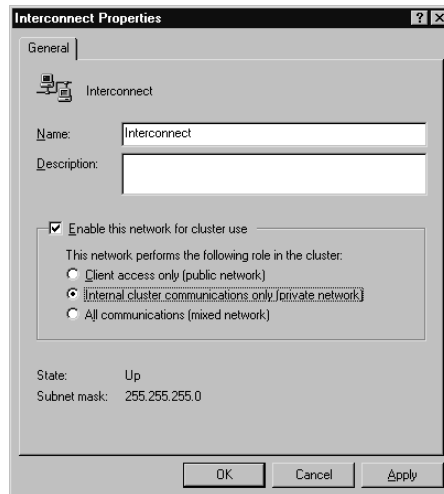
- **Internal cluster communications only (private network)** Used by nodes to handle their communication requirements only. No clients are present on this net-

work. This network should be physically separated from other networks and must have good response times (less than 500 ms) in order to avoid availability problems.

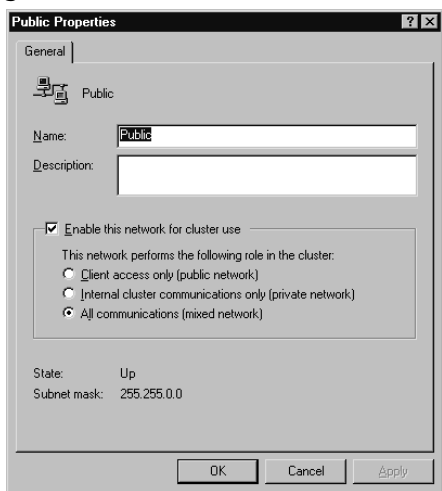
- **Client access only (public network)** Used to service client requests only. No internal cluster communication occurs over this network.
- **All communications (mixed network)** Can handle both categories of communications traffic. Normally, this network acts as a backup to a private network, but that is not required.
- **Nonclustered network (disabled)** Unavailable for use by the cluster for either servicing clients or for internal communications.

When you create the server cluster through the New Server Cluster Wizard, it will detect the different networks configured in the server. You will be asked to select the role each network will have in the server cluster. Select **Internal cluster communications only (private network)** for the interconnect(s), as shown in Figure 9.21, instead of accepting the default value (which will mix the server cluster heartbeat traffic with client communication traffic).

Figure 9.21 Configuring Interconnect Networks



If you are using only a single interconnect, you should configure at least one public network interface with the **All communications (mixed network)** setting, as shown in Figure 9.22. This allows the server cluster to have a backup path for internal server cluster communications, if one is needed. If you have multiple interconnects configured, you should set the public interfaces to the **Client access only (private network)** setting.

Figure 9.22 Configuring Public Networks

Multiple Interconnections

At least one interconnect between nodes is required. Node status messages are passed over this communication path. If this path becomes unavailable, a failover may be initiated. Because of this, multiple interconnects are recommended.

If your server cluster is configured with multiple interconnects, the reliability of the interconnects goes up. If a heartbeat message on one interconnect path goes unanswered, the node will attempt to use the other interconnect paths before initiating a failover. As with most components in a high-availability system, redundancy is good.

When using multiple interconnects, follow the same rules previously stated for configuration, but try to avoid using multiple ports on the same multiport network interface card (NIC). If the card fails, you will lose the interconnect. If you are using two dual-port cards, try to configure the system to use one port on each card for interconnects and the other port for your public network.

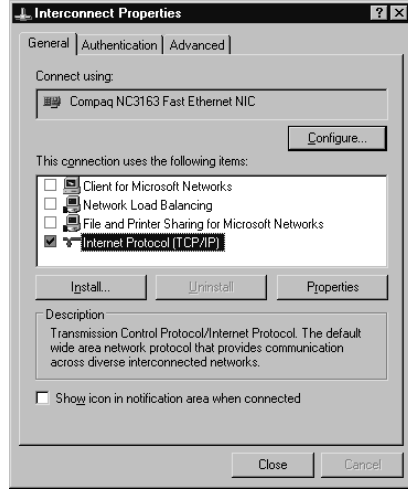
Node-to-Node Communication

The interconnects are used by the nodes to determine each other's status. This communication is unencrypted and frequent. Normal client activity does not occur on this network, so you should not have client-type services assigned to the network interface used for interconnects. Windows Server 2003 normally attaches the following services to each network interface:

- Client for Microsoft Networks
- Network Load Balancing
- File and Printer Sharing for Microsoft Networks
- Internet Protocol (TCP/IP)

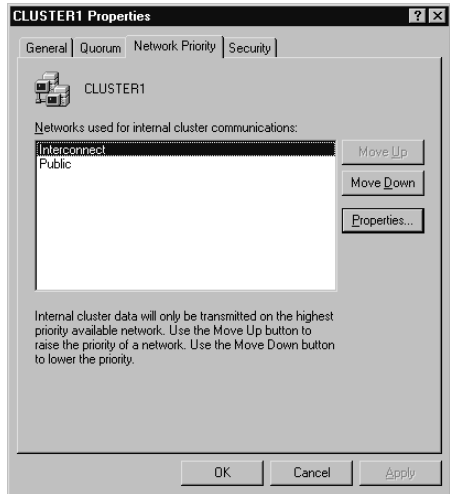
You should uncheck the first three services from each interconnect interface (the properties of a network interface are accessible via **Start | Control Panel | Network Connections**). Only TCP/IP should be assigned. Figure 9.23 shows a properly configured interconnect interface.

Figure 9.23 Configuring an Interconnect Interface



You should also make sure that the Network Priority property of the server cluster is configured with the interconnect(s) given highest priority, as shown in Figure 9.24. This ensures that internal cluster communication attempts are made on the interconnects first. To access this property, in Cluster Administrator, right-click the server cluster name and select **Properties**.

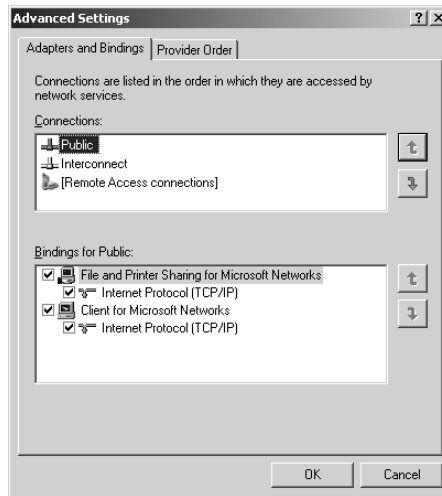
Figure 9.24 Setting the Network Priority Property of the Cluster



Binding Order

Binding is the process of linking the various communications components together, in the proper order to establish the communications path. To configure the binding order of communication protocols and services to the network interface, select **Start | Control Panel | Network Connections**. Click the **Advanced** menu and select **Advanced Settings....** When establishing the order of network connections, you should ensure that the public interfaces appear highest on the list, followed by interconnects, and then any other interfaces. Figure 9.25 shows this binding order.

Figure 9.25 Setting the Proper Binding Order of Interfaces



Adapter Settings

All network interfaces in a server cluster should be manually set for speed and duplex mode. Do not allow the network adapters to attempt to auto-negotiate these settings. If the controllers negotiate differently, your communications can be disrupted. Also, in many cases, a crossover cable is used on the interconnects. In these cases, an auto-negotiation may fail entirely, and the interconnect may never be established, affecting cluster operation.

As mentioned earlier, teamed network adapters must not be used for interconnects. However, they are perfectly acceptable for the public network interfaces. A failover or load-balanced configuration increases redundancy and reliability.

TCP/IP Settings

Static IP addresses (along with the relevant DNS and WINS information) should be used on public network interfaces. For the interconnects, you *must* use static IP addresses.

It is also a good practice to assign private IP addresses on interconnects from a different address class than your public class. For example, if you are using class A addresses (10.x.x.x) on your public interface, you could use class C addresses (192.168.x.x) on your intercon-

nects. Following this practice helps easily identify the type of network you may be troubleshooting just by looking at the address class. Using addresses this way is not required, but it does prove useful.

Finally, you should not configure IP gateway, DNS, or WINS addresses on your interconnect interfaces. Name resolution is usually not required on interconnects and, if configured, could cause conflicts with name resolution on your public interfaces. All public interfaces must reside on the same IP subnet. Likewise, all interconnect interfaces must reside on the same IP subnet.

The Default Cluster Group

Every server cluster has at least one cluster group: the default. This group contains the following resources:

- Quorum disk (which contains the quorum resource and logs)
- Cluster IP address
- Cluster name (which creates the virtual server)

When designing your server cluster, you should not plan on using these resources for anything other than system administration. If this group is offline for any reason, cluster operation can be compromised. Do not install applications on the quorum drive or in the default cluster group.

Security

Security is a consideration for any computer system. Server clusters are no exception. In fact, because they often contain critical information, they should usually be more closely guarded than a standard server.

Physical Security

Nodes should be kept in controlled environments and behind locked doors. More downtime is caused by accident than by intent. It is obvious that you would not want an unhappy or ex-employee to have access to your computer systems, but what about the curious user? Both can lead to the same end.

When setting up physical security, do not forget to include the power systems, network switches and routers, keyboards, mice, and monitors. Unauthorized access to any of these can lead to an unexpected outage.

Public/Mixed Networks

It is a good idea to isolate critical server clusters behind firewalls if possible. A properly configured firewall will also allow you to control the network traffic your server cluster encounters.

If there are infrastructure servers (DNS, WINS, and so on) that are relied on to access the server cluster, make sure that those servers are secured as well. If, for example, name resolution fails, it is possible that clients will not be able to access the server cluster even though it is fully operational.

Private Networks

The traffic on the private interconnect networks is meant to be used and accessed by nodes only. If high traffic levels disrupt or delay heartbeat messages, the server cluster may interpret this as a node failure and initiate a failover. For this reason, it is a good idea to place the interconnects on their own switch or virtual LAN (VLAN) and to not mix heartbeats with other traffic.

Do not place infrastructure servers (DNS, WINS, DHCP, and so on) on the same subnet as the interconnects. These services are not used by the interconnects and may cause the conflicts you wish to avoid.

Remote Administration of Cluster Nodes

Administration of your server cluster should be limited to a few controlled and trusted nodes. The administrative tools are quite powerful and could be used intentionally or accidentally to cause failovers, service stoppages, resource stoppages, or node evictions.

Use of Terminal Services on nodes is debatable. Terminal Services works just fine on nodes and actually includes some benefits. Evaluate your administrative, security, and operational needs to determine if installing Terminal Services on your nodes is appropriate for your situation.

The Cluster Service Account

The account that the cluster service uses must be a domain-level account and configured to be a member of the local Administrators group on each node. This account should not be a member of the Domain Admins group. Using an account that has elevated domain-level privileges would present a strong security risk if the cluster service account were to become compromised.

Do not use the cluster service account for administration, and be sure to configure it so that it can log on to only cluster nodes. Use different cluster service accounts for each cluster in your environment. This limits the scope of a security breach in the event that one occurs. If any of the applications running on your server cluster require accounts for operation, create and assign accounts specifically for the applications. Do not use the cluster service account for running applications. Doing so would make your cluster vulnerable to a malfunctioning application.

If you are required to permanently *evict* (forcibly remove) a node from a server cluster, you should manually remove the cluster service account from the appropriate local security groups on the evicted node. The cluster administrative tools will not automatically remove this account. Leaving this account with elevated permissions on an evicted node can expose you to security risks for both the evicted node and your domain.

Another possible method of securing a server cluster is to create a *domainlet*. A domainlet is a domain created just to host a server cluster. Each node in the server cluster is a domain controller of the domain. A domainlet allows you to better define and control the security boundary for the cluster. There are advantages and disadvantages to this approach. (For more information about domainlets, visit Microsoft's Web site.)

Client Access

Use the security features built into Windows Server 2003 and Active Directory (AD) to secure the applications and data on your server cluster. Turn on and use the auditing features of the operating system to see what activity is occurring on your server cluster.

Administrative Access

In larger organizations, it may be possible to have a different group of personnel responsible for administering clusters than those that perform other administrative tasks. Evaluate this possibility in your organization. If this strategy is adopted, assign these cluster administrators to a domain group and make that group a member of the appropriate local groups on the nodes. Also, assign NTFS permissions in a similar manner.

Cluster Data Security

As with any server, data should be accessed in a controlled manner. You do not want users accessing, deleting, or corrupting data. Assign appropriate NTFS file system permissions on a server cluster, just as you would assign them on a stand-alone server.

Disk Resource Security

Use NTFS permissions to ensure that only members of the Administrators group and the cluster service account can access the quorum disk. If you use scripts and the generic script resource type, you should assign appropriate NTFS Execute permissions to the scripts. A buggy script, or one run in an unplanned or uncontrolled manner, may cause data loss or a service outage.

Cluster Configuration Log File Security

When a cluster is created or a node is added to a cluster using the wizard, a file containing critical information about the cluster is placed in the `%systemroot%\System32\LogFiles\Cluster\` directory, unless you do not have administrative permissions on the node; in that case, the file is placed in the `%temp%` directory. The log file, `ClCfgSrv.log`, should have NTFS permissions that allow access to only the Administrators group and the cluster service account.

EXERCISE 9.01

CREATING A NEW CLUSTER

This exercise will walk you through the steps of creating a server cluster. Only the creation of the first node is covered. Each server cluster and network configuration is unique. You will need to substitute your TCP/IP addresses and account names, and adjust this process to fit your hardware.

1. Properly assemble your hardware. Ensure that only this first node is connected to and can access the shared storage unit(s).
2. Assign friendly names to your network interfaces and configure them with static IP addresses.
3. Log on to your domain with an account capable of creating user accounts. Open **Active Directory Users and Computers**. In the **Users** container, create an account called **ClusterAdmin** matching the settings shown in Figures 9.26 and 9.27. Close **Active Directory Users and Computers**.

Figure 9.26 Create a New Cluster Service User Account


The screenshot shows the 'New Object - User' dialog box. At the top, it says 'Create in: syngress.com/Users'. Below this are several input fields: 'First name' (empty), 'Initials' (empty), 'Last name' (ClusterAdmin), 'Full name' (ClusterAdmin), 'User logon name' (ClusterAdmin), and 'User logon name (pre-Windows 2000)' (SYNGRESS\ClusterAdmin). There are also buttons for '< Back', 'Next >', and 'Cancel'.

Figure 9.27 Assign a Password and Properties to New Cluster Service User Account

The screenshot shows the 'New Object - User' dialog box, specifically the password and properties section. It has 'Password' and 'Confirm password' fields, both masked with dots. Below these are four checkboxes: 'User must change password at next logon' (unchecked), 'User cannot change password' (checked), 'Password never expires' (checked), and 'Account is disabled' (unchecked). There are also buttons for '< Back', 'Next >', and 'Cancel'.

4. Log on to your first cluster node and start Cluster Administrator by selecting **Start | Administrative Tools | Cluster Administrator**.
5. When the **Open Connection to Cluster** dialog box is presented (Figure 9.28), select **Create new cluster** from the **Action** drop-down box and click **OK**.

Figure 9.28 Open Connection to Cluster




6. The **New Server Cluster Wizard** will start, as shown in Figure 9.29. Click **Next**.

Figure 9.29 The New Server Cluster Wizard’s Welcome Window



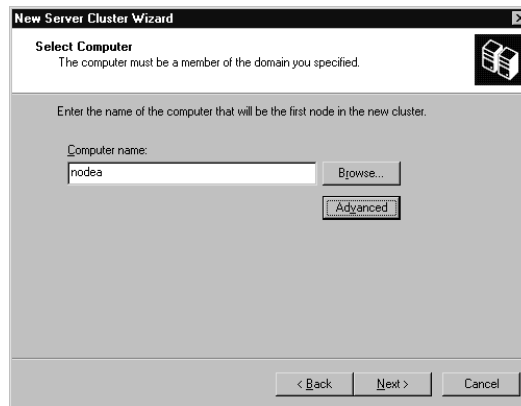
7. Select your domain in the **Domain** drop-down list and enter **cluster1** in the **Cluster name** text box, as shown in Figure 9.30. Click **Next**.

Figure 9.30 Specify the Cluster Name and Domain




8. Enter the name of the computer that will become your first node in the **Computer name** text box, as shown in Figure 9.31, and click **Next**.

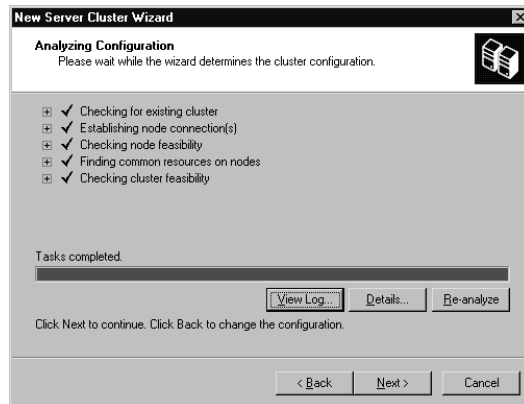
Figure 9.31 Select the Computer Name



9. The **Analyzing Configuration** window will appear, as shown in Figure 9.32, while the configuration of the node is verified. You can click the **View Log...** button to see the history of actions the Wizard has performed, or click the **Details...** button to see the most recent task.

Figure 9.32 Analyzing the Configuration of the Cluster Node

10. When the analysis is completed, the **Analyzing Configuration** window will show the tasks completed, as shown in Figure 9.33. Click the plus signs (+) to see the details behind each step. When you're finished examining the details, click **Next**.

Figure 9.33 Finished Analyzing the Configuration of the Cluster Node

11. You are asked what IP address you want assigned to the server cluster, as shown in Figure 9.34. Enter the appropriate **IP Address** and click **Next**.

Figure 9.34 Enter the Cluster IP Address

New Server Cluster Wizard

IP Address
Enter an IP address that cluster management tools will use to connect to the cluster.

IP Address:
10 . 20 . 200 . 9

< Back Next > Cancel

- In the **Cluster Service Account** window, shown in Figure 9.35, enter the **User name**, **Password**, and **Domain** for the cluster service account you created in step 3. Then click **Next**.

Figure 9.35 Enter the Cluster Service Account Information

New Server Cluster Wizard

Cluster Service Account
Enter login information for the domain account under which the cluster service will be run.

User name: ClustesAdmin
Password: *****
Domain: syngress.com

ⓘ This account will be given local administrative rights on all nodes of this cluster to allow for proper operation.

< Back Next > Cancel

- The Wizard will display the proposed server cluster configuration, as shown in Figure 9.36. Review the information.

Figure 9.36 Review the Proposed Cluster Configuration

New Server Cluster Wizard

Proposed Cluster Configuration
Verify that you want to create a cluster with the following configuration.

Cluster name:
cluster1.syngress.com

Cluster IP address:
10.20.200.5\255.255.0.0

Cluster network:
Public - Private and Public
Compaq NC3163 Fast Ethernet NIC
Primary Address: 10.20.200.4\255.255.0.0

Cluster service account credentials:
Name: ClustesAdmin
Password: *****

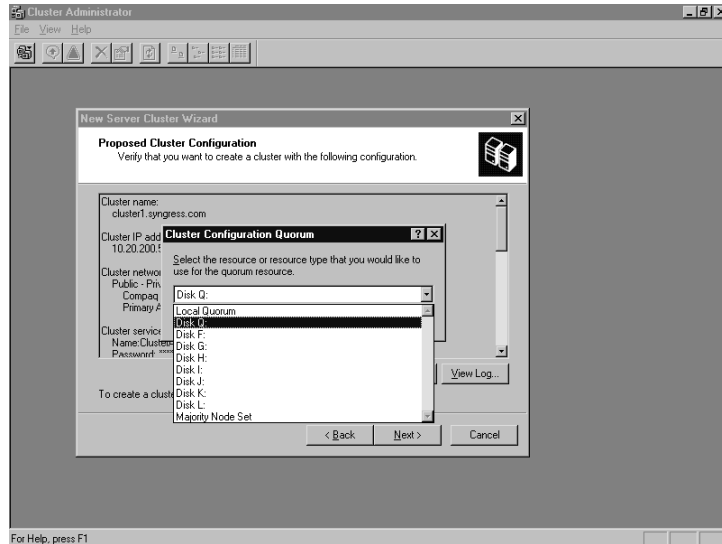
Quorum... View Log...

To create a cluster with this configuration, click Next.

< Back Next > Cancel

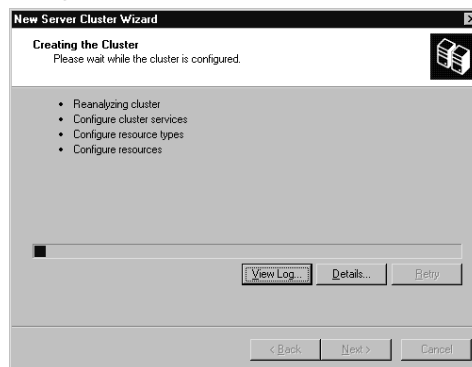
- Click the **Quorum...** button. Select the correct quorum disk for your configuration from the drop-down list, as shown in Figure 9.37, and select **OK**.

Figure 9.37 Select the Quorum Disk

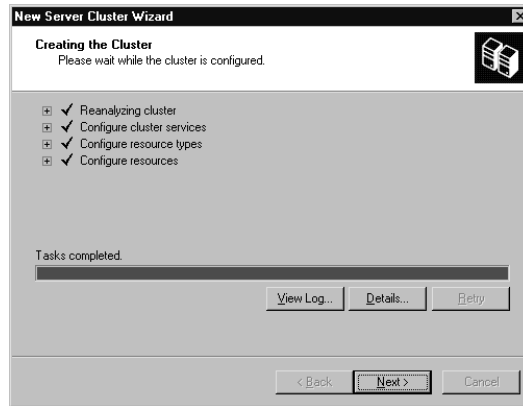


- The wizard will now create the server cluster, as shown in Figure 9.38. As the configuration progresses, you can click **View Log...** or **Details...** to see what the wizard is doing.

Figure 9.38 Creating the Cluster



- When the wizard finishes creating the server cluster, the **Creating the Cluster** window will show the tasks completed, as shown in Figure 9.39. Click the plus signs (+) to see details about each step performed. Click **Next**.

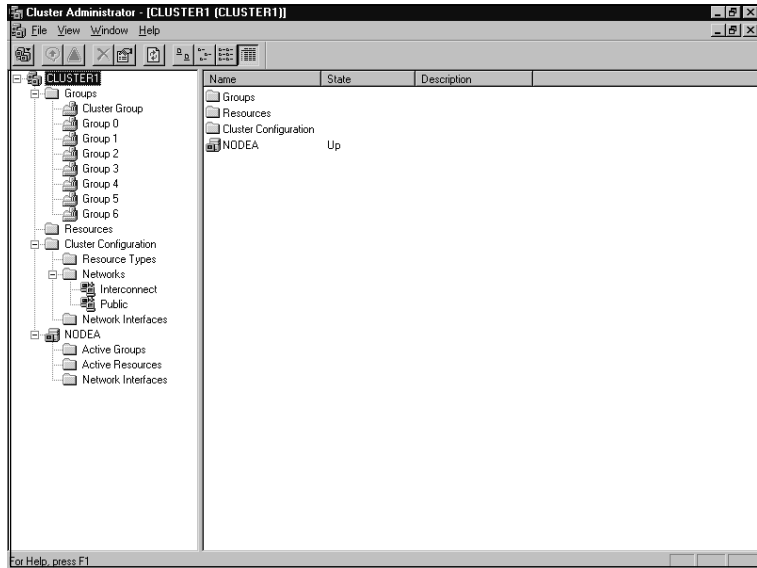
Figure 9.39 Completed Cluster Creation

17. The wizard informs you that the server cluster is created, as shown in Figure 9.40. You can click **View Log...** to examine all of the activity involved in the creation. Click **Finish** to exit the wizard.

Figure 9.40 The Wizard's Final Window


18. The **Cluster Administrator** utility appears. As shown in Figure 9.41, it displays the server cluster you just created.

Figure 9.41 The Newly Created Cluster



- Right-click the server cluster name (CLUSTER1) and select **Properties**. Click the **Network Priority** tab and move **Interconnect** to the top of the list, as shown in Figure 9.42. Click **Apply**.

Figure 9.42 Change Network Priorities



- Examine the **Quorum** and **Security** tabs to become familiar with the default settings on these tabs. When you have finished reviewing the configuration of these tabs, click **OK**. Then close Cluster Administrator.

EXAM
70-293
OBJECTIVE
4.1.2

Making Network Load Balancing Part of Your High-Availability Plan

The other high-availability tool included in Windows Server 2003 is Network Load Balancing (NLB). A primary use for NLB is increasing the scalability and availability of Internet applications (Web, FTP, VPN, firewall, proxy servers, and so on) by having multiple machines simultaneously answering and serving client requests. NLB is included in all versions of Windows Server 2003 and is installed automatically, although it must be configured and activated before it is usable.

Microsoft also considers NLB a clustering technology. The two clustering technologies are very different and serve different purposes. A server cluster requires specialized hardware, and there is typically one installed copy of each application, that moves between server cluster nodes. Only the node actively hosting the application responds to client requests. An NLB cluster does not require any specialized or additional hardware. *Every* host runs a *separate and independent copy* of the application and actively responds to client requests. Server clusters are used mainly for database-type applications. NLB clusters are used for traffic or communication oriented applications.



EXAM WARNING

Make sure you thoroughly understand the difference between the two clustering technologies and where each is primarily used. Exam questions may attempt to mix or confuse the two.

NLB has been available since Windows NT 4.0 when it was an add-in component called Windows Load Balancing Service (WLBS). You will still see NLB called this in some utilities and documentation. Unless specifically referred to in a historical context, the terms *WLBS* and *NLB* should be considered interchangeable.

Terminology and Concepts

NLB introduces some new terms for dealing with this form of clustering. Some terms are similar to those used with server clusters, but they have different meanings.

Hosts/Default Host

When referring to NLB, a *host* is a server running any edition of Windows Server 2003 that has been configured to respond to client requests via the NLB driver. Since NLB is automatically installed, any Windows Server 2003 server has the potential to be an NLB host.

The default host in an NLB cluster is the host with the highest currently active *priority*. The priority is a unique identifying number assigned to each host in an NLB cluster. An

NLB cluster can have up to 32 hosts, so the priorities range from 1 to 32. Hosts cannot be configured to have the same priority.

Load Weight

As previously mentioned, an NLB cluster can consist of up to 32 hosts. The hosts do not need to be identical in hardware or configuration. The *load weight* is a mechanism for distributing the traffic load within an NLB cluster to the hosts that are most suited to handle the load. Lighter loads can be configured for hosts with less capacity and heavier loads for more robust hosts.

The load weight is applicable only if specifically configured; otherwise, all hosts are configured with equal load weights. When used, each host is assigned a load weight from 0 (lowest weight) to 100 (highest weight). The weights from all active hosts in the cluster are averaged, and traffic is distributed accordingly. In this way, the load weight is a relative value within the NLB cluster.

Traffic Distribution

The way requests from clients are spread out among the hosts in an NLB cluster is referred to as *traffic distribution*. Each host in an NLB cluster is configured with at least two IP addresses. One address is reserved for the nonclustered traffic directed to the host, and the second IP address is shared among all nodes in the cluster and is called the *cluster IP address*. It is to this second IP address that clients direct their requests.

When a request is sent to the cluster IP address, all hosts in the cluster receive the request. The NLB driver passes the incoming traffic through the defined *port rules*. The host that the port rules specify to receive the request services the request, while all other hosts discard the request. Port rules are the mechanism used to direct incoming traffic on specific TCP/IP ports to specific hosts or groups of hosts. All hosts in an NLB cluster must have the same number and specific port rules. Port rules can apply to a specific cluster IP address, all port numbers, or a specific range of port numbers, and to the TCP, UDP, or both protocols.

In addition, each port rule contains a *filtering mode* for that rule. The filtering mode defines how the hosts in a cluster handle inbound traffic. The options for the filtering mode are as follows:

- **Disabled** All traffic matching the associated cluster IP address, port range, and protocol will be blocked. Applications on the NLB cluster will never see this traffic.
- **Single Host** All traffic matching the associated cluster IP address, port range, and protocol will be handled by one specific host in an NLB cluster. For example, this filtering mode could be used to direct all FTP traffic inbound to an NLB cluster to host 2 of that cluster, while Web traffic is served from all nodes.

- **Multiple Host** All traffic matching the associated cluster IP address, port range, and protocol will be distributed to multiple hosts in the NLB cluster. When using the multiple host filtering mode, you must also select an *affinity*. Affinity describes how multiple requests from the same client are directed among the multiple hosts. There are three affinity options:
 - **None** Any NLB host matching the port rule can service requests from clients. This is the most efficient affinity setting in terms of evenly distributing the workload, but it should not be used with the UDP or Both protocol settings to properly handle fragmented packets.
 - **Single** This is the default setting. Single affinity ensures that only one NLB host will handle traffic requests for the same client session. This setting is necessary if session state preservation is needed. (for example, for Web servers using server-side cookies). This setting reliably supports the UDP or Both protocol setting.
 - **Class C** This affinity setting specifies that all client requests originating from the same class C IP subnet will be directed to the same NLB host. This setting is useful in large NLB clusters handling traffic inbound from the Internet. This setting also reliably supports the UDP or Both protocol setting.

Convergence and Heartbeats

An NLB cluster can be a fluid environment. By design, a host can be added or removed from the operational cluster without affecting the services provided by the NLB cluster. However, each time a host is added to or removed from the NLB cluster, the cluster must reconfigure itself to allow for the new increased or decreased capacity, and calculate for traffic distribution accordingly. This process is called *convergence*. During convergence, the new stable state of the cluster and default host (the host with the highest priority) is determined.

Convergence normally occurs within 10 seconds, and client requests to operational hosts are unaffected. Requests to hosts that have failed or exited the cluster are redistributed to working hosts after convergence is completed.

NLB cluster hosts determine the status of each other by exchanging *heartbeat* messages. Heartbeats in an NLB cluster differ from those used in a server cluster but serve a similar purpose. In essence, the heartbeat messages generated by an NLB host are a way for the host to tell the other members of the cluster “I’m alive.” By default, if a host does not send a heartbeat message to the other NLB cluster hosts within five seconds, it will be considered failed and a convergence will be initiated.

How NLB Works

NLB requires between 2 and 32 host systems to be effective. Each host has its own copy of the applications being supported by the cluster. The hosts share one or more IP addresses. When the cluster is started, the hosts perform a convergence. Once convergence is complete, the hosts will begin responding to client requests. Client systems then issue requests directed to one of the cluster IP addresses. All of the cluster hosts receive the request. The host that is next in line to service the request does so, while the other hosts ignore it.

Once per second, a host issues heartbeat messages to the other hosts in the NLB cluster. If another host is added or if a host leaves, the cluster will perform another convergence.

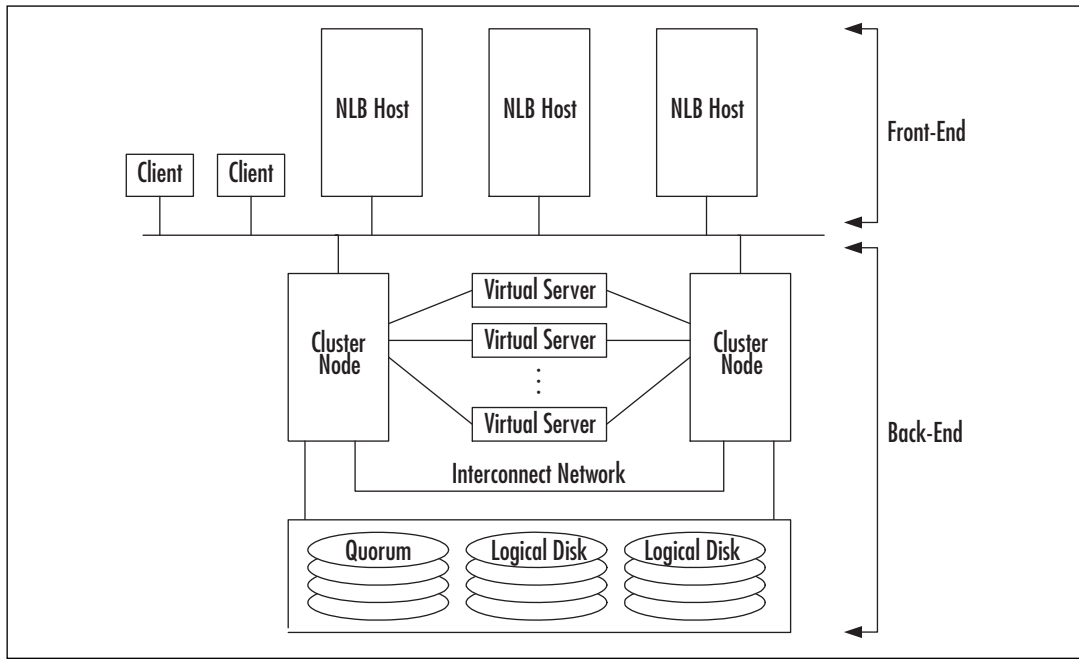
Relationship of NLB to Clustering

Server clustering and NLB clustering differ greatly. You cannot combine NLB and server clustering on the same hosts, but the two technologies can sometimes be used together to increase overall reliability and performance.

Server clustering is used primarily for database-type applications (such as SQL Server, Exchange Server, and Oracle) that run as a single instance of the application, and parallel or concurrent execution is impossible or impractical. Server-clustered databases often operate behind an NLB cluster. For this reason, a server cluster is sometimes referred to as the *back-end*.

NLB is used for applications whose primary resource is TCP/IP communication-related—such as Internet Information Server (IIS), ISA Server, virtual private Network (VPN) servers, and terminal servers—that can run in multiple instances or in a parallel fashion. By adding hosts to an NLB cluster, more requests can be serviced simultaneously, increasing responsiveness and performance. The applications on the NLB hosts would then issue requests to the back-end on the client's behalf, process the returned request, and then fulfill the original client request. Since the NLB cluster logically resides between the client and the server cluster, or in “front” of the server cluster, the NLB cluster is usually referred to as the *front-end*. The combination of these two high-availability technologies can be very powerful and reliable. Figure 9.43 illustrates this front-end/back-end structure.

Figure 9.43 Combining Network Load Balancing and Server Clustering into a Front-end/Back-end Architecture



EXAM
70-293
OBJECTIVE
4.4

Managing NLB Clusters

Windows Server 2003 includes some useful tools for creating and managing NLB clusters. The NLB Manager (new to Windows Server 2003) is provided to centrally create and manage NLB clusters from a graphical interface. For performing administrative tasks from the command-line interface, the `NLB.exe` utility is provided.

Using the NLB Manager Tool

Microsoft made many improvements and added many tools to Windows Server 2003, but NLB Manager should earn Microsoft a special thanks. This tool is extremely powerful. It takes what used to be a difficult manual process and simplifies it with a point-and-click interface. With NLB Manager, you can perform the following tasks:

- Create a new NLB cluster.
- Add and automatically configure a new host.
- Remove a host from an NLB cluster, automatically disabling NLB on the removed host.
- Configure all NLB-related properties on the cluster.

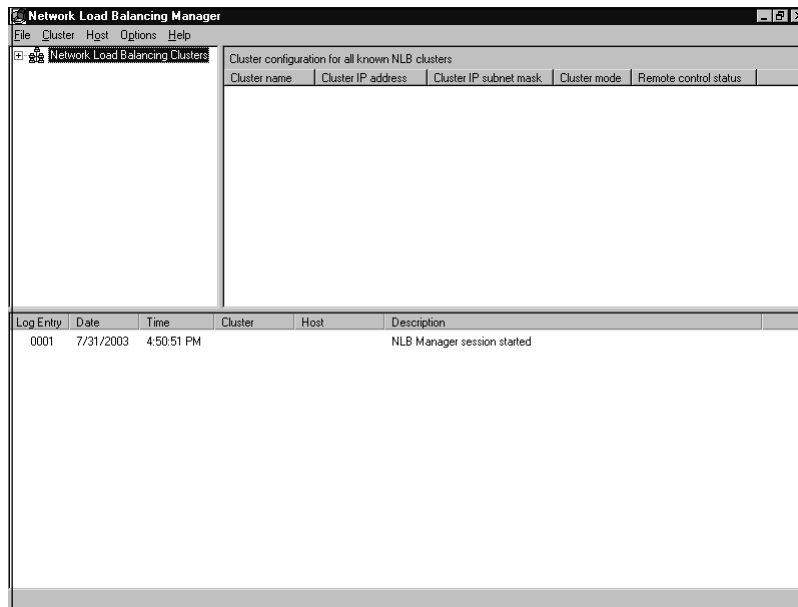
- Configure all hosts in the cluster.
- Replicate the NLB cluster configuration (but not applications) to other NLB hosts.
- Troubleshoot NLB clusters.

To run NLB Manager, you must be a member of the local Administrators group on the host you are adding, configuring, or removing from the cluster. You do not need to have elevated privileges for the system on which you are running NLB Manager.

The NLB Manager utility is a part of the Windows Server 2003 Administration Tools Pack, which can be found in `%systemroot%\System32\Adminpak.msi`. The Administration Tools Pack can be installed on a Windows XP workstation to allow remote administration.

To access NLB Manager, select **Start | Administrative Tools | Network Load Balancing Manager**. When the utility starts for the first time, you are presented with an empty session, as shown in Figure 9.44. From here, you can begin the process of creating or managing an NLB cluster.

Figure 9.44 Starting NLB Manager for the First Time



Remote Management

You must take a series of steps in order to remotely manage an NLB cluster or host with NLB Manager. NLB Manager uses Windows Management Instrumentation (WMI) interfaces. WMI requires Remote Procedure Call (RPC) and Distributed Component Object Model (DCOM) availability. You can verify that these services are available for NLB

Manager's use by selecting **Start | Administrative Tools | Services** and viewing the list of services.

If you are attempting to manage an NLB cluster that is on the other side of a firewall from your location, you will need to make sure that your firewall is configured to allow DCOM to pass. Microsoft has a white paper available that describes how to do this at www.microsoft.com/com/wpaper/dcomfw.asp.

Command-Line Tools

Before Windows Server 2003, the only way to manage a load-balanced cluster was with command-line tools. In some situations, this approach still makes sense, because command-line tools can be scripted and scheduled.

Microsoft includes the NLB.exe utility for this purpose. NLB.exe can perform many of the same functions as NLB Manager, but it uses a different mechanism that is disabled by default. NLB.exe uses the remote-control feature of NLB instead of RPC and DCOM. This may be advantageous in certain circumstances, but enabling the remote-control feature exposes the cluster to possible security risks. Microsoft recommends that remote control be disabled and suggests that you perform all NLB administration through NLB Manager.

If you need to use NLB.exe, make sure that you enforce strong passwords on the NLB cluster and keep your NLB cluster behind a firewall. The default UDP ports used by NLM.exe are 1717 and 2504.

Figure 9.45 shows the command-line parameters that can be used with NLB.exe.

Figure 9.45 Output of the NLB.exe /? Command

```
Usage: NLB <command> [/PASSW [<password>]] [/PORT <port>]
<command>
  help                - displays this help
  ip2mac <cluster>    - displays the MAC address for the
                       specified cluster
  reload [<cluster> | ALL] - reloads the driver's parameters
                       from the registry for the
                       specified cluster (local only).
                       Same as ALL if parameter is not
                       specified.
  display [<cluster> | ALL] - displays configuration parameters,
                              current status, and last several
                              event log messages for the
                              specified cluster (local only).
                              Same as ALL if parameter is not
                              specified.
```

Continued

Figure 9.45 Output of the NLB.exe /? Command

```

query      [<cluster_spec>]      - displays the current cluster state
                                for the current members of the
                                specified cluster. If not
                                specified a local query is
                                performed for all instances.

suspend    [<cluster_spec>]      - suspends cluster operations
                                (start, stop, etc.) for the
                                specified cluster until the resume
                                command is issued. If cluster is
                                not specified, applies to all
                                instances on local host.

resume     [<cluster_spec>]      - resumes cluster operations after a
                                previous suspend command for the
                                specified cluster. If cluster is
                                not specified, applies to all
                                instances on local host.

start      [<cluster_spec>]      - starts cluster operations on the
                                specified hosts. Applies to local
                                host if cluster is not specified.

stop       [<cluster_spec>]      - stops cluster operations on the
                                specified hosts. Applies to local
                                host if cluster is not specified.

drainstop  [<cluster_spec>]      - disables all new traffic handling
                                on the specified hosts and stops
                                cluster operations. Applies to
                                local host if cluster is not
                                specified.

enable <port_spec> <cluster_spec> - enables traffic handling on the
                                specified cluster for the rule
                                whose port range contains the
                                specified port

disable <port_spec> <cluster_spec> - disables ALL traffic handling on
                                the specified cluster for the rule
                                whose port range contains the
                                specified port

drain <port_spec> <cluster_spec> - disables NEW traffic handling on
                                the specified cluster for the rule

```

Continued

Figure 9.45 Output of the NLB.exe /? Command

whose port range contains the

```

                                specified port
queryport [<vip>:]<port>         - retrieve the current state of the
    [<cluster_spec>]             port rule. If the rule is handling
                                traffic, packet handling
                                statistics are also returned.
params [<cluster> | ALL]         - retrieve the current parameters
                                from the NLB driver for the
                                specified cluster on the local
                                host.
<port_spec>
    [<vip>: | ALL:](<port> | ALL) - every virtual ip address (neither
                                <vip> nor ALL) or specific <vip>
                                or the "All" vip, on a specific
                                <port> rule or ALL ports
<cluster_spec>
    <cluster>:<host> | ((<cluster> | ALL) - specific <cluster> on a
        specific
        (LOCAL | GLOBAL))       <host>, OR specific <cluster> or
                                ALL clusters, on the LOCAL machine
                                or all (GLOBAL) machines that are
                                a part of the cluster
<cluster>                       - cluster name | cluster primary IP
                                address
<host>                           - host within the cluster (default -
                                ALL hosts): dedicated name |
                                IP address | host priority ID
                                (1..32) | 0 for current DEFAULT
                                host
<vip>                             - virtual ip address in the port
                                rule
<port>                             - TCP/UDP port number
Remote options:
    /PASSW <password>           - remote control password (default -
                                NONE)
                                blank <password> for console prompt
    /PORT <port>                - cluster's remote control UDP port

```



EXAM WARNING

One particularly useful function of both NLB.exe and NLB Manager is the *drainstop* option. This feature allows you to plan the shutdown of an NLB host *without* affecting sessions already in progress (think *transparent to the user*). This function works by setting the host to not allow any new connections to it. As existing connections complete their conversations, their sessions are closed. If that same client starts a new connection, the NLB cluster directs the connection to another available host. It is likely that the exam will present questions regarding this scenario or function.

NLB Error Detection and Handling

The objective of NLB is increased availability. Consequently, Microsoft has included mechanisms in NLB to handle and manage error situations without affecting the reliability of the NLB cluster. If an error is encountered, details about the error are recorded in the Windows event log, and NLB isolates the host having the problem by preventing it from joining the cluster and servicing requests.

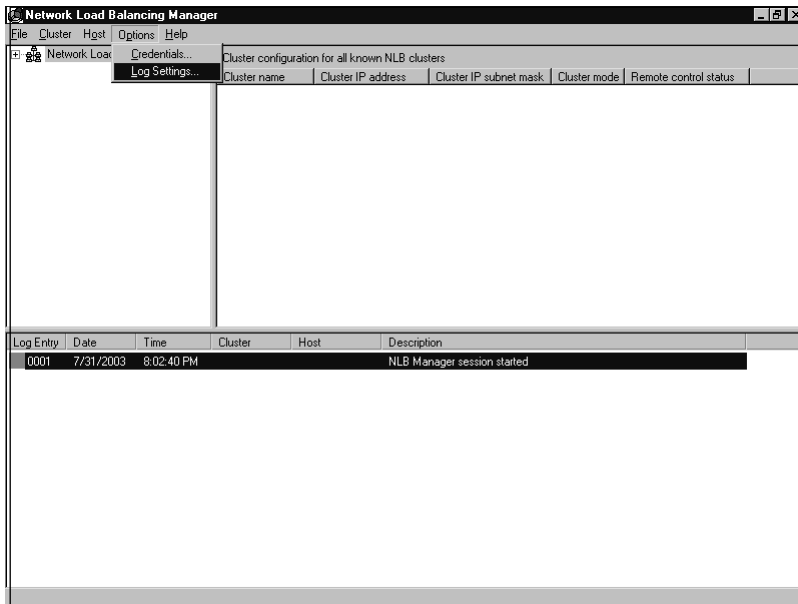
As previously stated, an NLB cluster performs a convergence when a host joins or leaves the cluster. When a host attempts to join, it notifies the other cluster hosts of its configuration. Likewise, the other hosts notify the joining host of their configurations. A check for consistency in operating parameters (host priority, port rules, and so on) is performed. If the host that is attempting to join does not have a configuration consistent with the hosts already in the cluster, the new host will not be allowed to join, and convergence will not occur. This process ensures that a misconfigured host does not compromise cluster operations.

Monitoring NLB

Events encountered by NLB (convergence, communication errors, and so on) are recorded by NLB in the System event log. You can use Event Viewer to examine these events.


NLB Manager does not use the Windows event logs. Instead, it includes its own logging function that records actions performed by the utility. This log file allows you to see what administrative activity has occurred on your NLB cluster. The log function must be activated before it can be used. To activate the log, start NLB Manager and select **Options** | **Log Settings...**, as shown in Figure 9.46.

Figure 9.46 Starting an NLB Manager Log



When the **Log Settings** dialog box appears, as shown in Figure 9.47, check **Enable logging** and enter a path and filename for the log. If no path is given, the log file is stored in the profile of the logged-on user account.

Figure 9.47 Enabling the NLB Manager Log



This log file contains sensitive information about your NLB cluster. You should secure it by restricting access to it with NTFS permissions. Be aware, however, that the account under which NLB Manager runs will require Full Control permissions to the log file.

Using the WLBS Cluster Control Utility

If you have enabled the remote-control feature of NLB, you can use the `NLB.exe` command-line utility to get status information from an NLB cluster. You can use the **NLB query** command to display the current configuration, status, and any recent event log mes-

sages for the NLB cluster. The **NLB display** command displays the current state of the NLB cluster and hosts.

NLB Best Practices

As with all technologies, there are certain ways to implement and operate NLB that are better than others. Microsoft publishes a number of items that fall into the best practices category for NLB.

Multiple Network Adapters

NLB can be implemented with a single network interface adapter in each host, but multiple adapters are recommended. A single network interface generates additional communications overhead for the NLB cluster, because all hosts see the network traffic destined for a specific host.

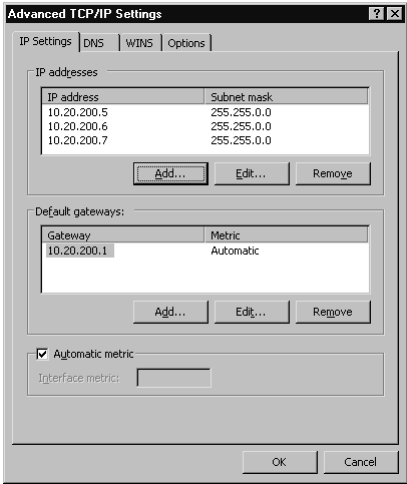
You are also limited in how you can perform administrative tasks. A host with a single NIC cannot perform regular (non-NLB) communications. This means that you cannot run the NLB administrative tools on an NLB host in this configuration. To avoid this situation, you must enable multicast or use multiple network adapters. When multiple network adapters are installed in each host, one adapter can be configured for NLB and the other for regular traffic. When using multiple adapters, you should configure only one adapter for use by NLB.

Protocols and IP Addressing

NLB supports only TCP and UDP communications. Do not attempt to attach any other protocols (IPX/SPX, AppleTalk, ATM, and so on) to the adapter. Only static IP addresses are allowed on an NLB cluster node. DHCP is not supported. This is true for the cluster IP address and the dedicated host IP address.

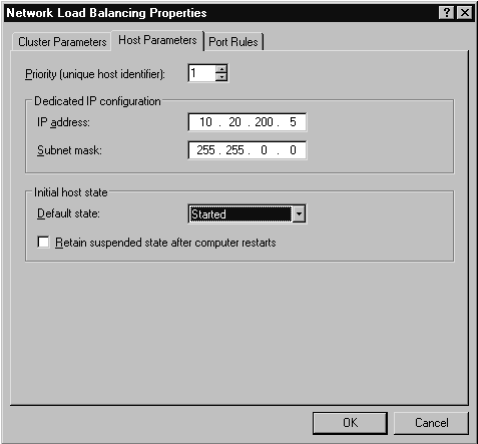
Each node in an NLB cluster must be on the same TCP/IP subnet. NLB does not support hosts residing on multiple subnets. When configuring the IP addresses for your hosts, keep in mind that multiple IP addresses can be assigned to an adapter, and *all* of those IP addresses will be load-balanced, *except* for the address configured as the host's dedicated address (the one that handles non-NLB traffic). The host's dedicated IP address must be first on the list of IP addresses assigned to a network interface, so that any outbound traffic from the host is sent from this IP address. Figure 9.48 shows a network adapter with multiple IP addresses configured in the **Advanced TCP/IP Settings** dialog box (to open this dialog box, click **Advanced** in the **Internet Protocol (TCP/IP) Properties** dialog box for the network interface properties).

Figure 9.48 Configuring a Network Adapter with Multiple IP Addresses



You will notice from the example in Figure 9.48 that the IP address **10.20.200.5** is listed first and is therefore the node’s dedicated IP address. This configuration is not complete, however, until the properties of the NLB driver are also configured with this IP address, as shown in Figure 9.49 (check **Network Load Balancing** in the property pages of the network interface, and then click the **Properties** button to open this dialog box).

Figure 9.49 NLB Dedicated IP Address Configuration



Security

Security is of greater concern in an NLB cluster than it is with a stand-alone server. NLB has no inherent security features, and it cannot be used as a firewall or in any other intru-

sion-prevention role. When improperly configured, NLB can open security holes into your environment. It is critical that you take proper security precautions when using NLB.

Host Security

Consider tightening the security of the operating system. Limit the number of users permitted to access the hosts. Place a secured PC in front of the NLB cluster and behind a firewall. Use this PC to run NLB Manager and administer the cluster.

Application Security

Because NLB provides no additional security functions, it is imperative to use any security features available in your load-balanced applications. If you are using IIS on an NLB cluster, follow the documented procedures and guidelines for securing IIS.

Physical Security

Like any server, an NLB host should be locked behind closed doors for protection, and so should the network equipment that the NLB cluster depends on. It is theoretically possible to cause a service disruption by forging cluster heartbeats.

Host List

If you are using the host list feature of NLB Manager, you should secure the host list file on your administrative system. Restrict access to appropriate users.

Remote Control Option

The remote-control feature of NLB is a known security risk. You should avoid using this feature. If you must enable remote control, ensure that strong passwords are used. It is also advisable to place the cluster behind a firewall and filter the port traffic going to the remote-control ports.

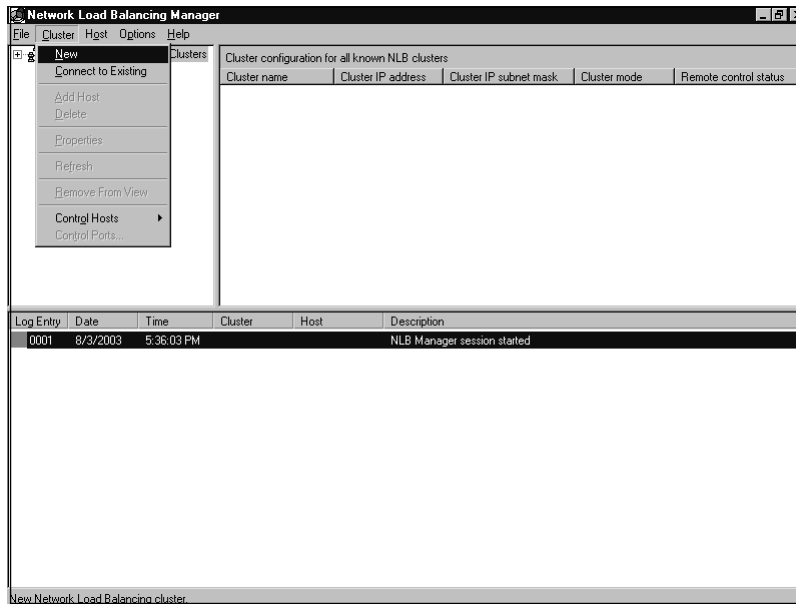
EXERCISE 9.02

CREATING A NETWORK LOAD BALANCING CLUSTER

This exercise will walk you through the process of creating a new NLB cluster using the NLB Manager administrative tool. Where appropriate, use your own TCP/IP addresses in this exercise.


1. Start NLB Manager by selecting **Start | Administrative Tools | Network Load Balancing Manager**.
2. Select **Cluster | New**, as shown in Figure 9.50.

Figure 9.50 Create a New NLB Cluster



- You will be presented with the **Cluster Parameters** window. Enter the **IP address**, **Subnet mask**, and **Full Internet name** (this is the fully qualified domain name) of the cluster in the **Cluster IP configuration** section, as shown in Figure 9.51.

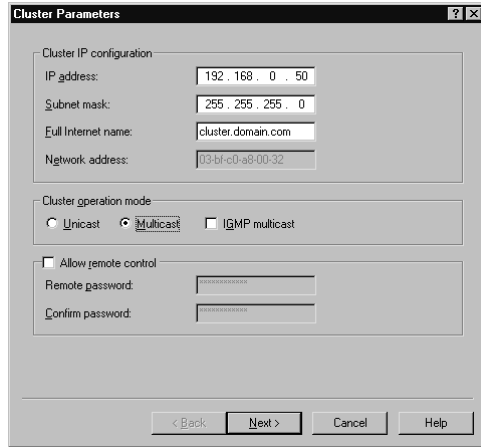
Figure 9.51 Configure Cluster Parameters



- Click the **Multicast** option in the **Cluster operation mode** section, and notice how the **Network address** entry changes, as shown in Figure


9.52. The network (media access control, or MAC) changes to fit the correct mode based on the communication mechanism you select. (We will leave **Multicast** selected for the exercise.)

Figure 9.52 Select Multicast Cluster Operation Mode




5. Select the check box next to **IGMP multicast**, as shown in Figure 9.53.

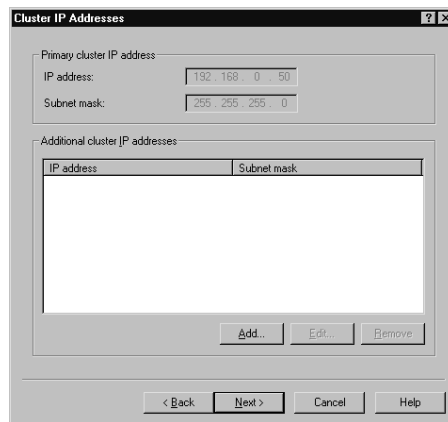
Figure 9.53 Select IGMP Multicast with the Cluster Operation Mode



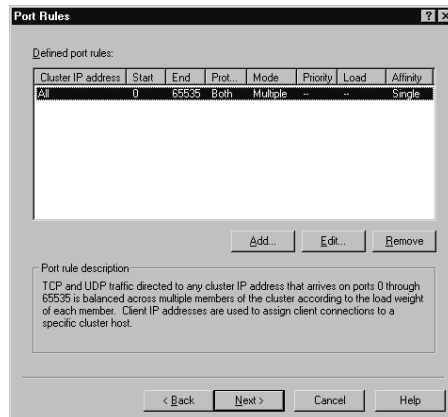
6. You will be presented with the warning message shown in Figure 9.54. This message is intended to remind you that additional configuration of your switches and NIC may be required if you select IGMP support. Click **OK** to close the **Warning** dialog box.

Figure 9.54 IGMP Warning Message

- You will be presented with the **Cluster IP Addresses** window, as shown in Figure 9.55. If you want to load-balance multiple IP addresses, you can click the **Add...** button and add them to the cluster at this point. For this exercise, we will work with only one address. Click **Next** to continue.


Figure 9.55 Cluster IP Addresses Window

- In the **Port Rules** window, you see the default port rule, as shown in Figure 9.56. This rule evenly distributes arriving traffic among all cluster hosts. Select the default port rule and click **Edit...**

Figure 9.56 The Port Rules Window

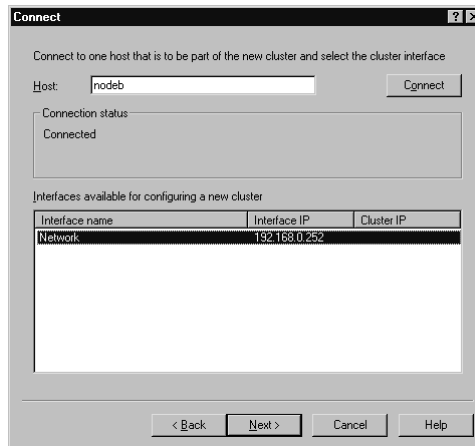
- The **Add/Edit Port Rule** dialog box appears, as shown in Figure 9.57. As you can see, the default port rule applies to all cluster IP addresses on all ports and protocols. It also directs all client requests to the same cluster host (Multiple host/Single Affinity). Click **Cancel** to avoid modifying the default port rule.

Figure 9.57 The Add/Edit Port Rule Dialog Box




- Click **Next** in the **Port Rules** window to advance to the **Connect** window.
- Enter the name of a host in the **Host** field and click the **Connect** button. When the host is identified, select the network interface to load-balance, as shown in Figure 9.58. Then click **Next**.

Figure 9.58 Connect to an NLB Node



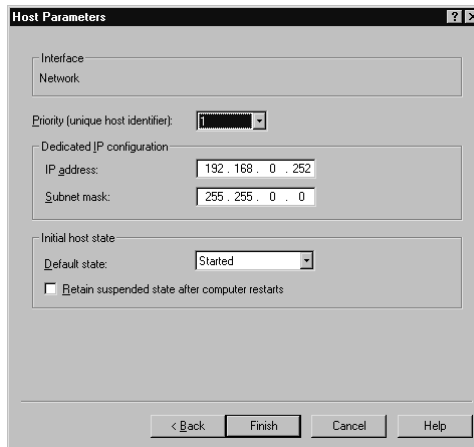
At this point, you may receive the warning message, as shown in Figure 9.59. If you receive this message, you are using DHCP to assign an IP address to your network interface. You must use static IP addresses on your network interfaces when using NLB. You must cancel the configuration, change from DHCP to static IP addresses, and begin this process again.

Figure 9.59 DHCP Warning Message



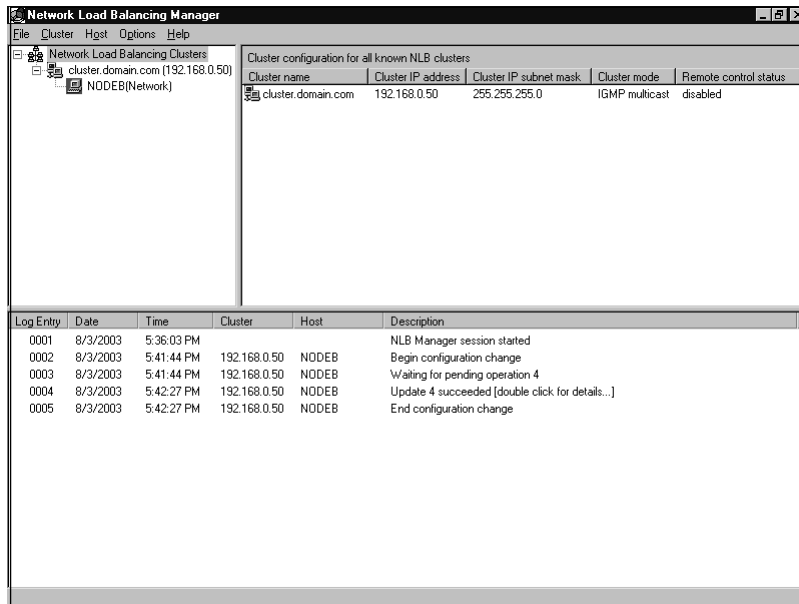
12. You are now presented with the **Host Parameters** window, as shown in Figure 9.60. Enter the **Priority**, **Dedicated IP address**, and **Subnet mask** for the cluster host. Set the **Default state** of the host to **Started**. (This setting will make the host automatically attempt to join the NLB cluster on startup). Click **Finish**.

Figure 9.60 Configure Host Parameters



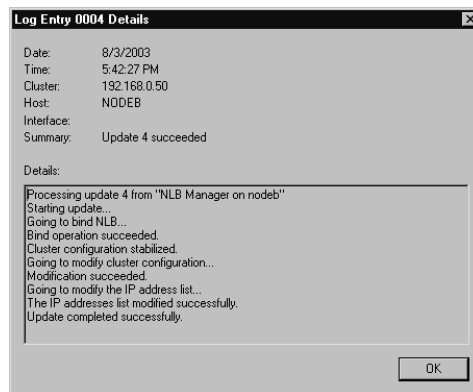
13. You are now taken back to the main window of the **NLB Manager** utility, which will look similar to Figure 9.61.

Figure 9.61 The Configured NLB Cluster



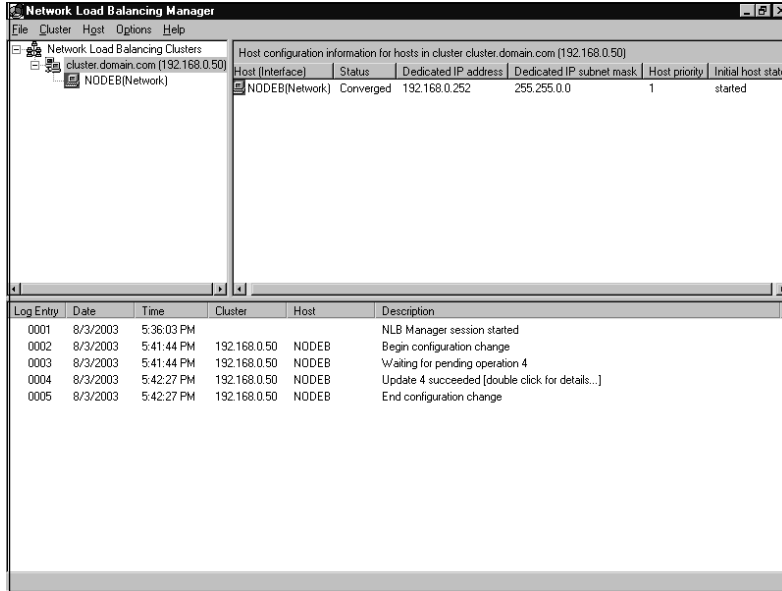
- The bottom pane of the window is the log of activities performed by the NLB Manager. Double-click an entry. Figure 9.62 shows an example of the details that appear when Log Entry 0004 was double-clicked. When you are finished viewing the log entry's details, click **OK**.

Figure 9.62 View NLB Manager Log Entry Details



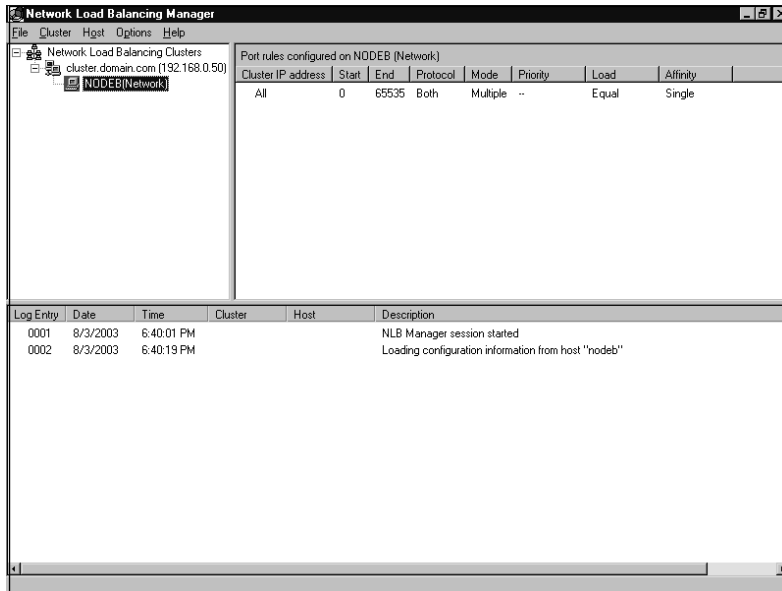
- Click the NLB cluster you just created. You will see current details about your cluster, similar to those shown in Figure 9.63.

Figure 9.63 Configured NLB Cluster Details



- Click the host you just configured. You will see the port rules, as shown in Figure 9.64.

Figure 9.64 Configured Port Rules on Cluster Node



Summary of Exam Objectives

The two high-availability technologies included in Windows Server 2003 are server clusters and NLB clusters. Server clusters provide failover support for applications, primarily databases. NLB provides parallel servicing for TCP/IP-based network applications, such as Web servers, firewalls, and DNS servers.

Server clusters use specialized high-end, high-capacity hardware, and shared storage components to increase availability. An application running on a server cluster runs as a single instance of the application, although multiple separate and unrelated instances may be running on a server cluster. NLB can be used on generalized or low-end hardware and requires nothing extra to implement, except for an optional second network interface adapter. Each host in an NLB cluster has separate copies of the applications, which are unaware of each other's existence.

A server cluster can be configured in one of three ways depending on its intended purpose and geographic operating environment. It can also be configured to follow a certain pattern of behavior when a failure is encountered. NLB clusters are fundamentally configured identically each time an NLB cluster is created, except for the way the cluster handles client traffic and internal communications. NLB cluster error-control behavior is limited to excluding a failed or improperly configured host from the cluster.

Server clusters and NLB cannot coexist on the same hardware. However, the two technologies can be used together to form an extremely reliable foundation for fast, reliable service. NLB is often used to create a front-end structure to handle incoming client requests. The hosts in an NLB cluster then make requests of applications residing on a back-end server cluster. This structure allows efficient processing of large numbers of client requests and reliable availability of data.

Exam Objectives Fast Track

Making Server Clustering Part of Your High-Availability Plan

- ☑ Individual servers in a server cluster are called *nodes*.
- ☑ A server cluster can consist of up to eight nodes.
- ☑ The individual physical or logical components managed by a cluster are called *resources*. Resources are combined to form *groups*, which are the basic administrative unit of a server cluster.
- ☑ Applications operating on a server cluster move between nodes in a process called *failover*. What happens during a failover is configurable and depends on the server cluster design.

- ☑ The resource that is used to control the server cluster is the *quorum*. Every server cluster has a quorum resource whose form is determined by the cluster design model adopted.
- ☑ Server clusters require very specific up-front planning for successful implementation.

Making Network Load Balancing Part of Your High-Availability Plan

- ☑ NLB is used to increase availability for applications that service TCP/IP traffic only. Other protocols are not supported.
- ☑ Individual servers in an NLB cluster are called *hosts*.
- ☑ NLB clusters determine their operating state through a process called *convergence*.
- ☑ NLB clusters can be administered from with a graphical tool (NLB Manager) or a command-line tool (NLB.exe). The graphical tool is more secure.
- ☑ An NLB cluster does not require multiple network adapters in each host, although this is recommended.
- ☑ NLB provides no additional security features and requires more stringent security practices than do stand-alone servers.

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 nodes can be in a server cluster?

A: One to eight nodes can be in a server cluster, depending on the cluster's design and applications.

Q: What server cluster model is most frequently used?

A: The single quorum device model is the most used cluster model.

Q: How many applications can a server cluster run?

A: There is no inherent limit to the number of applications a server cluster can run. Only the resources of the nodes limit the number of applications that can be run on a node.

Q: Can a server cluster be configured with a single network interface adapter?

A: No, at least one public interface and one private interface must be present in each cluster node.

Q: Can any application be run on a server cluster?

A: No, only applications that can operate as a system service, use the TCP/IP protocol, can be configured to store their data where specified, and whose client applications attempt to reconnect if a failure is encountered can be configured to operate on a server cluster.

Q: Can server clusters be created from any equipment?

A: No, server clusters require very specific and robust hardware configurations.

Q: Can any application be used with Network Load Balancing?

A: No, only network-type, TCP-based applications can be used with **Network Load Balancing** (NLB).

Q: Can server clustering and NLB be used together?

A: Yes, but not on the same hardware. They can be used in conjunction to form high-availability solutions.

Q: Can NLB hosts be created from any equipment?

A: Yes, an NLB host is not required to be expensive or high performance.

Q: How many hosts can be in an NLB cluster?

A: Up to 32 hosts can be in an NLB cluster.

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.

Making Server Clustering Part of Your High-Availability Plan

1. You have purchased a prepackaged solution that uses an eight-node majority node set (MNS) server cluster. Because you have so many nodes, you have decided to install three nodes in your Atlanta data center, three in your Denver data center, and the last two in your Seattle sales office. You notice fairly soon that the server cluster is experiencing some uptime issues. The nodes in your Atlanta data center seem to fail frequently during times of high WAN utilization. What is likely the problem?
 - A. All nodes in an MNS server cluster must be in the same data center.
 - B. The high WAN traffic is making the heartbeats take longer than 500 ms to get to all nodes and back.
 - C. The nodes in Atlanta are failing, and an MNS server cluster can have two nodes fail before losing quorum and shutting down.
 - D. The cluster cannot be in three geographic areas. An MNS server cluster can exist in a maximum of two geographic regions, and high-speed networks must connect the nodes in each region.
2. Your data center experiences a power failure, bringing all of your systems down. When power is returned, a single quorum device server cluster you have in use will not start. You examine the event logs and find error messages stating that the quorum drive cannot be found, yet you are able to view the contents of the quorum drive in Windows Explorer. Research reveals that either the disk signature on the quorum drive or the Registry key containing the disk signature for the quorum drive has been corrupted. What steps should you take to recover from this problem?

- A. Evict all other nodes from the server cluster, repartition and reformat the quorum drive, and rejoin the other nodes to the server cluster.
 - B. Do a restore of the quorum drive from tape.
 - C. Change the location of the quorum resource to another drive, repartition and reformat the quorum drive, and move the quorum resource back to the original quorum drive.
 - D. Shut down all nodes except one, perform an ASR restore on that node, and restart all the nodes.
3. As a consultant, you have been called in to attempt to fix a high-availability configuration that is not performing as designed. Your client wanted to provide high availability for a high-traffic Web site. The client purchased a preconfigured, mid-range, two-node server cluster and implemented IIS on the nodes. Response time for serving Web pages is unacceptable, although there have been no incidents of the application failing over. What is the correct fix for this situation?
- A. More nodes need to be added to the server cluster. Increase the number of nodes until performance reaches an acceptable level.
 - B. Add NLB to the server cluster to handle more requests from clients simultaneously.
 - C. Convert the server cluster to an NLB cluster.
 - D. Move the server cluster to high-end hardware to provide quicker response times.
4. You have been asked to design a server cluster. The server cluster will start small, but it may expand as more applications are added and predicted growth is experienced. Your proposal is for two nodes, a shared storage device, Fibre Channel host bus adapters, and switches for connectivity. When you present your proposal to management, you are asked to justify the high cost of the Fibre Channel solution. What justification can you provide for implementing Fibre Channel?
- A. Fibre Channel supports more than two nodes, allowing for the predicted growth.
 - B. Fibre Channel is the fastest connectivity solution and will therefore yield the highest performance.
 - C. Fibre Channel easily expands to allow more storage to be added to support the future applications.
 - D. All of the above.

5. You are configuring a large, single quorum device server cluster consisting of eight nodes and a dozen shared storage cabinets with 30 logical drives among them. Because of the large number of logical drives, you are using mount points instead of drive letters on most of the drives. After running the Wizard to create your first node, you can see only the drives that have been assigned drive letters. How is this resolved?
 - A. Install the second node, which will automatically create mount point resources.
 - B. Manually create the disk resources after the first node is created.
 - C. Reconfigure the shared storage to reduce the number of logical drives to less than 16.
 - D. Temporarily assign drive letters to the mount point drives, and then remove the drive letters after the Wizard finishes installing the first node.

6. You are configuring a two-node, single quorum device server cluster with a single public network interface and a single interconnect interface. The network interfaces and storage devices have been configured, and the interconnects on both nodes have been connected with a direct crossover Ethernet cable. The installation of the first node proceeds without incident, but when attempting to create the second node, the installation fails. The Wizard reports problems communicating with the first node over the interconnect. You have verified that the cables are functional and have been properly inserted into the connectors. What is the most likely problem?
 - A. The interconnect adapters are configured for auto-negotiation or for different speed and duplex settings.
 - B. The direct crossover cable method cannot be used with this cluster configuration.
 - C. A second interconnect is required with this cluster configuration.
 - D. A switch must be used to handle heartbeat traffic.

7. You have installed a third-party backup agent on your nodes. The agent is supposed to listen for requests from its control server and send data to it during a backup. Despite this, your backups are failing. The application on the control server reports that it cannot communicate with the agent. You check the node and see that the agent is running properly. What is the most likely problem?
 - A. The agent is not server cluster-compatible and cannot be used on a node.
 - B. The control server is attempting to communicate with the agent over the interconnect network.
 - C. There is a firewall between the control server and the node running the agent.
 - D. The agent has configured itself to listen on the interconnect instead of the public network.

8. You have created a small, two-node, single quorum device server cluster to act as a print server for several hundred printers. The shared storage is a 4GB drive. Because of the small size of the shared storage and the transient nature of the data, the print spool resource is on the quorum drive. The server cluster operates acceptably for a period of time, and then both nodes are taken down by a sudden power failure. When power is restored, the nodes boot, but the cluster service will not start. How do you fix this problem and prevent it from happening again?
 - A. Delete the files under the spool directory, remove the spooler resource, add external storage, and re-create the spool on a different drive.
 - B. You cannot resolve this issue. Once the quorum drive is filled, all nodes must be evicted and the server cluster re-created.
 - C. Reformat the quorum drive and apply disk quotas to prevent the spooler from filling the drive again.
 - D. Perform an ASR restore on the nodes.

9. You are configuring a large, single quorum device server cluster consisting of eight nodes and a dozen shared storage cabinets with 30 logical drives among them. The storage cabinet that contains the quorum drive also contains eight other logical drives and is connected to the last port on your 32-port Fibre Channel switch. While running the Wizard to create your first node, you cannot see any of the drives in the quorum drive's cabinet, including the quorum drive. Which of the following is a possible cause of the problem?
 - A. The maximum number of logical drives recognizable in by a server cluster configuration has been exceeded.
 - B. The cabinet containing the quorum drive is not properly connected or powered on.
 - C. The cabinet containing the quorum drive must be relocated to a lower numbered Fibre Channel switch port.
 - D. The maximum number of storage devices recognizable by a cluster has been exceeded.

Making Network Load Balancing Part of Your High-Availability Plan

10. You have installed an NLB cluster onto a 10/100 Mbps switch. Other devices, including some older 10 Mbps-only devices, are also attached to the switch. Your NLB hosts are configured for 100 Mbps and full duplex. Soon, you notice that communications with the 10 Mbps devices have failed. After troubleshooting, you discover that apparently the increased traffic on the switch is preventing the 10 Mbps devices from having sufficient bandwidth for reliable communications. What is the best fix for this problem?
 - A. Change the operating mode of the NLB cluster to multicast and enable IGMP support.
 - B. Relocate all of the NLB hosts to a different virtual LAN (VLAN).
 - C. Relocate all of the 10 Mbps-only hosts to the same VLAN.
 - D. Install a firewall between the NLB hosts and the 10 Mbps-only devices and filter all NLB-oriented traffic.

11. You have configured an NLB cluster with 10 hosts. The default port rule has been changed from all possible ports to just port 80. No other port rules have been defined. You have configured each node with IIS and followed the appropriate procedures for installing and securing IIS. After clients begin using the cluster, you notice that clients requesting normal Web pages are being served equally across the cluster, but clients requesting secured Web pages (SSL) and FTP sessions are all going to the host with priority 1. What is the best way to resolve this issue and to balance the SSL and FTP requests?
 - A. Do nothing. SSL and FTP traffic cannot be load-balanced.
 - B. Split the NLB cluster into three clusters and serve the SSL and FTP sessions from different clusters.
 - C. Add new port rules for the SSL and FTP traffic.
 - D. Change the default port rule back to encompass all possible ports

12. You are a consultant. You have been called in to troubleshoot a malfunctioning NLB cluster that serves IIS Web pages. The cluster in question consists of six hosts, but only four successfully join the cluster. Two of the hosts never successfully join. When the rest of the hosts are shut down and those two hosts are started up together, they successfully perform convergence and form a cluster. This two-host cluster, however, seems to favor certain types of incoming traffic on each host, rather than equally among the two hosts. What is the most likely reason for this behavior?

- A. The two malfunctioning hosts are configured with different cluster IP addresses and a different host name than the four correctly operating hosts.
 - B. The two malfunctioning hosts are underpowered and cannot join the cluster due to poor performance.
 - C. The two malfunctioning hosts are configured with different port rules than the four correctly operating hosts.
 - D. The two malfunctioning hosts are configured with the same priority.
13. You are a consultant. You have been called in to troubleshoot a malfunctioning NLB cluster that is supposed to serve Web pages with IIS. The cluster contains four hosts, but only one host at a time will successfully form the cluster. Clients appear to have no problems connecting to any of the single-host cluster configurations. What is the most likely cause of the problem?
- A. The hosts are configured with duplicate priorities.
 - B. The hosts are configured with different port rules.
 - C. The hosts are configured with different cluster IP addresses.
 - D. The hosts are configured with duplicate cluster IP addresses.
14. One of your hosts in multiple-host NLB cluster requires maintenance. The cluster is heavily used and central to the profitability of your company. You want to bring the node down for service in the least disruptive way. How should you accomplish this goal?
- A. Use the drainstop option on the host needing maintenance.
 - B. Use the drainstop option on all the hosts in the cluster not needing maintenance.
 - C. Use the suspend option on the host needing maintenance.
 - D. Use the suspend option on all the hosts in the cluster not needing maintenance.
15. You have been asked to develop a design for an NLB cluster for an IIS-based Web site. The specifications given to you state that the Web application will be using server-side cookies to keep track of a visitor's session state. Which port-rule filtering mode should you configure to support the application?
- A. Single host
 - B. Multiple host/Affinity: None
 - C. Multiple host/Affinity: Single
 - D. Multiple host/Affinity: Class C

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. D | 10. A |
| 3. C | 11. C |
| 4. D | 12. C |
| 5. B | 13. A |
| 6. A | 14. A |
| 7. D | 15. C |
| 8. A | |

MCSE 70-293

Planning, Implementing, and Maintaining Internet Protocol Security

Exam Objectives in this Chapter:

- 3.3.1 Create and implement an IPSec policy.
- 5 Planning and Maintaining Network Security
- 5.3 Plan for network protocol security.
- 5.6 Plan security for data transmission.
- 5.6.1 Secure data transmission between client computers to meet security requirements.
- 5.6.2 Secure data transmission by using IPSec.
- 5.1 Configure network protocol security.
- 5.1.1 Configure protocol security in a heterogeneous client computer environment.
- 5.1.2 Configure protocol security by using IPSec policies.
- 5.2.1 Configure IPSec policy settings.
- 5.3.2 Plan an IPSec policy for secure network communications.
- 5.7 Troubleshoot security for data transmission. Tools might include the IP Security Monitor MMC snap-in and the Resultant Set of Policy (RSOP) MMC snap-in.
- 5.2 Configure security for data transmission.

Introduction

Securing sensitive or mission-critical data is an important part of the network administrator's job. Data is especially vulnerable to interception as it travels across the network. Windows Server 2003 includes Microsoft's implementation of the Internet standard IP Security (IPSec) protocol, for the purpose of protecting data in transit. This chapter deals with how to work with Windows Server 2003's IPSec. We start by introducing IPSec terminology and concepts and explaining how IPSec works "under the hood" to secure data in transit over the network. We discuss the purposes of IPSec encryption: authentication, integrity, and confidentiality. You'll learn about how IPSec operates in either of two modes: tunnel or transport.

Although we refer to IPSec as a protocol, it is actually a *framework*, or a collection of protocols and standards designed to protect IP data in transit. In this chapter, you'll learn about the protocols used by IPSec. These include the two primary protocols: the Authentication Header (AH) protocol and the Encapsulating Security Payload (ESP) protocol. We'll also discuss the roles of additional protocols used by IPSec, including the Internet Security and Key Management Protocol (ISAKMP), Internet Key Exchange (IKE), the Oakley key-determination protocol, and the Diffie-Hellman key-agreement protocol. You'll learn about Windows Server 2003's IPSec components—the IPSec driver and the IPSec Policy Agent service. We'll also discuss the relationship of IPSec to Internet Protocol version 6 (IPv6).

Next, we'll show you how to deploy IPSec on your network, taking into consideration organizational needs and security levels, and help you determine the appropriate authentication methods. You'll learn about managing IPSec, and we'll walk you through the process of using the IP Security Policy Management Microsoft Management Console (MMC) snap-in as well as the command-line tools. We'll discuss the role of IPSec policies, including default and custom policies, and we'll show you how to assign and apply policies. We'll also talk about IPSec security considerations and issues, including the use of a strong encryption algorithm (Triple Data Encryption Standard, or 3DES), authentication methods, firewall packet filtering, unprotected traffic, Diffie-Hellman groups, and the use of pre-shared keys. We'll show you how to use the Resultant Set of Policy (RSoP) and the RSoP MMC snap-in to view policy assignments and to simulate policy assignments for deployment planning.

Understanding IP Security (IPSec)

Microsoft's modern operating systems, which include Windows 2000 Server and Professional, Windows Server 2003, and the Windows XP Client operating system, give you the ability to enforce smart security policies without excessive overhead and expense, as well as to encrypt data traveling across the network, using IPSec.

The Internet Engineering Task Force (IETF) designed the IPSec specifications. The IP Security Working Group of the IETF developed IPSec as an industry standard for encrypting TCP/IP traffic within networking environments. The two main goals of IPSec

EXAM
70-293

OBJECTIVE
3.3.1
5
5.3
5.6
5.6.1
5.6.2

are to protect IP packets and to give network administrators the ability to use packet filtering as a defense against network attacks. Microsoft's Windows Server 2003 IPSec deployment includes the following features:

- Enhanced IPSec security monitoring with the MMC
- IPSec integration with Active Directory that allows for security policies to be centrally administered
- Use of Kerberos 5 authentication as the default method by IPSec policies to verify the authenticity of connecting computers
- Backward compatibility with the Windows 2000 Security Framework
- Client and application transparency, because IPSec works at the Network layer of the OSI model
- Automatic security negotiation

IPSec in transport mode is based on an end-to-end security model, meaning that security and trust are established from the source IP address and end with the destination IP address. Each computer handles security at its respective end with the assumption that the medium over which the communication takes place is not secure. IPSec is not required to be supported by any intermediary computer that routes data from the source to destination IP address, unless *network address translation* (NAT) or packet filtering has been implemented on the firewall. IPSec can be deployed with IPSec policy in Windows Server 2003 under any of these circumstances:

- Client-to-client and peer-to-peer support
- Gateway-to-gateway and router-to-router support
- Remote-access client dial-up and Internet access from private networks

The IP Security Policy Management MMC allows network administrators to set security policy settings and options that will allow the systems to negotiate with other systems regarding the traffic that is sent and received from that system.

When Not to Use IPsec

IPsec is useful in a number of different scenarios, as we discuss in the text. However, Microsoft recommends that IPsec *not* be used in certain situations, because deployment can be difficult and cause access problems on your network. Specifically, you should not use IPsec to secure the messages that pass between a domain controller and members of the domain. Getting this to work requires setting up very complex policies, and according to Microsoft, it is best avoided.

In addition, you generally should not try to configure IPsec to secure all traffic between all clients and all servers on a network. Broadcast traffic and multicast traffic cannot be secured via IPsec, and there are many types of application traffic that won't work with IPsec, including traffic that is generated by real-time communications programs, peer-to-peer applications, and applications that rely on Internet Control Message Protocol (ICMP).

There are other cases where using IPsec is generally not a good idea, including using IPsec to secure 802.11 wireless communications (the 802.1X security protocol is recommended instead).

Terminology and Concepts

Before you create a security plan to implement IPsec within your organization, it is important to understand the terminology and concepts relating to IPsec. This chapter deals entirely with IPsec and is not a generic security chapter; however, we will briefly touch on some basic security terms to clarify how these concepts are integrated with IPsec in Windows Server 2003.

IPsec uses two primary protocols: Authentication Header (AH) and Encapsulating Security Payload (ESP). These protocols can be used individually or together, and provide data security on wide area networks (WANs), local area networks (LANs), remote offices, corporate workgroups, domain servers, and client computers. Because IPsec works at the Network layer of the OSI model rather than at the Application layer as many other security applications (for example, the Secure Sockets Layer (SSL) protocol) do, it is easy to implement without making changes to client computers (other than setting client policy).

The main components of IPsec are shown in Table 10.1. In the following sections, we will discuss how each of these components is involved in securing data with IPsec.

Table 10.1 IPsec Terminology

Term	Definition
Authentication Header (AH)	One of the two primary IPsec protocols. AH is used to provide data authentication and integrity. It does not provide data confidentiality.

Continued

Table 10.1 IPSec Terminology

Term	Definition
Encapsulating Security Payload (ESP)	One of the two primary IPSec protocols. ESP provides authentication and integrity services via a keyed hash that is computed for just the ESP header, trailer and payload. It provides data confidentiality.
Security Association (SA)	Consists of an agreement of security settings associated with keying material.
Internet Key Exchange (IKE)	Used to manage and exchange cryptographic keys between client machines and negotiate a common set of security settings between client machines, eliminating the need for the two client machines to have exact policies configured.
Internet Security Association and Key Management Protocol (ISAKMP)	An add-on protocol for IPSec.
Triple Data Encryption	A strong encryption algorithm that is standard on all Windows Standard (3DES)Server 2003 computers and client machines running Windows. 3DES uses 56-bit keys.
Oakley key-determination protocol	A secondary protocol for IPSec by which two authenticated parties agree on secret key material. It uses the Diffie-Hellman algorithm.
Diffie-Hellman groups	A method used for key agreement, to establish a shared key over insecure media. Diffie-Hellman groups are based on the number of bits in the base prime numbers used in key exchange. Group 1 provides 768-bit key strength. Group 2 uses 1024 bits, and group 2048 uses 2048 bits. Windows Server 2003 supports group 2048, but it is not supported by Windows 2000, Windows XP, or other Microsoft operating systems.
Resultant Set of Policy (RSOP)	A Windows Server 2003 tool that is used to view advanced IPSec policy assignments for clients who belong to a particular Group Policy container in Windows Server 2003.

How IPSec Works

Before secure data can be exchanged, a security agreement between the two communicating computers must be established. This security agreement is called a security association (SA). Both IPSec-enabled computers agree on how to send and receive data, as well as how to protect the information contained in the data packets. Because IPSec SAs are unidirectional, at least two separate SAs are established to protect the data for every communication: one for inbound traffic and one for outbound traffic. There is a unique SA for each direction and for each protocol. Thus, if you are using both AH and ESP, there will be two SAs for AH and two for ESP.

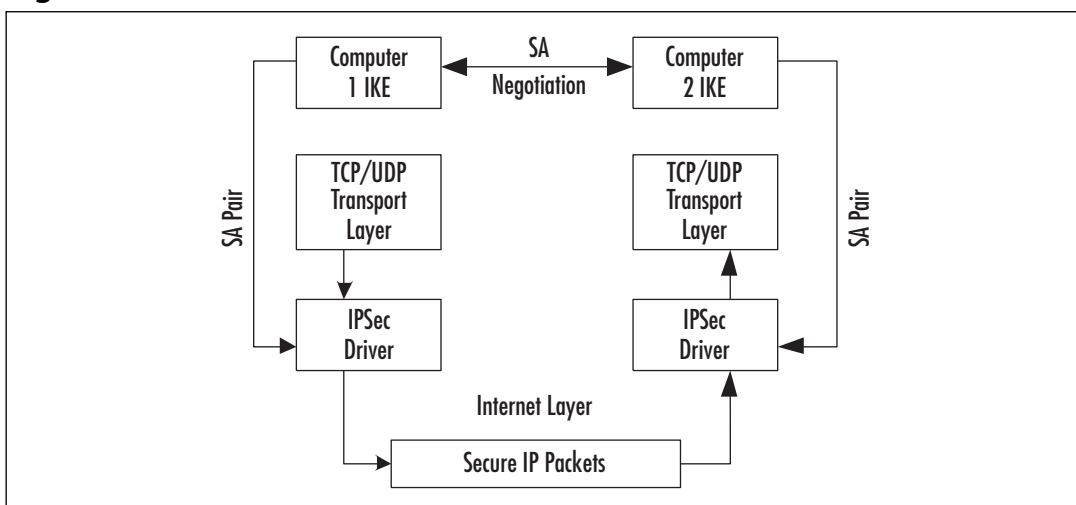


NOTE

In addition to the SAs used by IPsec itself, there is another type of SA (called a *main mode SA*) that protects the IKE negotiation. The SAs used by IKE are bidirectional, so a single SA can handle both outgoing and incoming traffic.

Using the IP Security Policy Management console, you can configure the security policy to block, permit, or negotiate security within your networked environment. Because this security is transparent to users, it is easy to implement and administer. Figure 10.1 shows how this process takes place.

Figure 10.1 How the SA Process Functions



Securing Data in Transit

An SA is a combination of three things:

- Security protocols
- A negotiated key
- A security parameters index (SPI)

These items together define the security settings that are used to protect the communication from the source IP to the destination IP. The SPI is a unique entry in the IPsec header of each packet and is used to identify which SA is being used to secure data. As mentioned earlier, there will always be separate SAs for inbound and outbound traffic. If a computer is communicating with multiple machines (for example, a database server with multiple clients sending queries), many SAs will exist. The receiving computer uses the SPI to determine which SA should be used to process incoming IP packets.

Purposes of Encryption

IPSec functions by using cryptographic techniques. The term *cryptography* refers to methods of making data unreadable or undecipherable by anyone except the authorized recipient in the event that the message is intercepted by someone else. IPSec uses cryptography to provide three basic services:

- Authentication
- Data integrity
- Data confidentiality

There are times when only one or two of these services is needed, and other times when all of these services are needed. We will take a look at each of these services individually.

Head of the Class...

IPSec Encryption Algorithms

IPSec provides computer-level authentication, as well as data encryption, for virtual private network (VPN) connections that use the Layer Two Tunneling Protocol (L2TP). One important purpose of IPSec encryption is to provide for data confidentiality so that the messages that travel through the VPN tunnel cannot be read by unauthorized persons. This is the “private” part of virtual private networking.

Before an L2TP connection is established, IPSec is negotiated between the client computer and the VPN server that uses L2TP. When the negotiation is completed, the data and the password are secure. One point of negotiation is the encryption algorithm that will be used. Windows Server 2003 supports the following encryption algorithms:

- **Data Encryption Standard (DES)** This method uses a single 56-bit key encryption level.
- **Triple Data Encryption Standard (3DES)** This method uses three 56-bit keys for encryption.

In today’s security-conscious environments, most servers are set to allow encryption and allow the client machines to select their encryption methods (algorithms). You can also set the server settings to deny encryption, select the specific encryption strength, or allow the client computer to select the encryption strength. Data encryption is very important if you want to ensure that your data is not readable in the event that it is captured by a “sniffer” or otherwise intercepted as it travels across the network.

Authentication

Authentication is the process of verifying the identity of a data sender or recipient. This allows the message recipient to know that the message was actually sent from the sender

and not from someone posing as the sender. IPSec can use different methods to authenticate identities, including pre-shared keys, digital certificates, and Kerberos authentication. Authentication is needed when it is important to verify that a message came from the person who claims to have sent it.

A concept closely related to authentication is *nonrepudiation*, which refers to a way of ensuring that the sender cannot later deny sending the message.

IPSec can also provide *anti-replay*. This refers to ensuring that an unauthorized person cannot capture the authentication credentials as they're sent across the network and "replay" them to establish a communications session with the server.



NOTE

The use of pre-shared keys is not recommended, because it is the least secure of the authentication methods supported by Windows Server 2003 IPSec. The biggest problem with any shared secret such as a pre-shared key is the difficulty of sharing the key with both parties without compromising it.

Data Integrity

Data integrity refers to the ability to ensure that the data that is received at the endpoint of the communication is exactly the same data that was sent from the originating computer, and it has not been modified in any way in transit. IPSec uses the *hash functions* to ensure that the contents of the data packet have not changed between the time it was sent and the time it was received.

Head of the Class... Hashing and Hash Algorithms

A hash algorithm used for encryption is a mathematical calculation that has been proven to be *one-way* so that it cannot be reverse-engineered (discovery of the original message using the hash result). (Two-way hashes are sometimes used for purposes other than encryption.) The result of the application of the algorithm is called the *hash result*.

Hashing uses a secret key to create a *message digest*, which is a combination of the message itself and the hash result. The message digest is sent to the recipient, and the same key is applied to it. The recipient applies the same key to the message, and the result will be identical if there has been no alteration.

The Message Digest 5 (MD-5) and Secure Hash Algorithm (SHA) algorithms are two popular hashing algorithms.

Data Confidentiality

Data *confidentiality* refers to the ability to “scramble” the data using encryption algorithms so that it cannot be understood by an unauthorized person who intercepts it. IPsec provides data confidentiality only through the ESP protocol. AH does *not* provide for encryption of the data. ESP uses the 3DES and DES algorithms to ensure data confidentiality.

IPsec Modes

IPsec in Windows Server 2003 has two different modes: tunnel mode and transport mode. Your choice of which IPsec mode to use depends on your organizational needs. We will take a look at how each of these works and when each is appropriately used.

Tunnel Mode


Tunneling refers to a method of encapsulating a data packet inside another packet and routing the new packet across a network. Tunnels are used to create VPNs that allow data to go across the Internet (or another public or nonsecure network) without compromising security, because the inner packet and its header information are not visible on the public network.

In tunnel mode, IPsec encrypts the IP header and the payload, thereby securing the entire IP packet. It is used primarily when end systems or gateways do not support the L2TP/IPsec or the Point-to-Point Tunneling protocol (PPTP). In other words, tunnel mode allows you to use IPsec to create a tunnel, in addition to encrypting the data within the tunnel, with servers that cannot use the traditional VPN tunneling protocols (L2TP and PPTP). However, Windows Server 2003 does not support using IPsec as the tunneling protocol for remote access VPNs; it is only supported between gateways, routers, and servers. Remote access clients must use PPTP or L2TP for VPN connections.

The entire packet is encrypted by either AH or ESP. These two protocols will be discussed in more detail in the “IPsec Protocols” section. The outer IP header contains the addresses of the tunnel endpoints, and the encapsulated IP header contains the ultimate source and destination addresses, as illustrated in Figure 10.2.

Tunnel mode is used to protect data traveling between different networks that must pass through an untrusted network (such as the Internet). Tunnel mode works in the following configurations:

- Gateway to gateway
- Server to gateway
- Server to server

Figure 10.2 The IPSec Tunnel Mode

Transport Mode

Transport mode, the default mode for IPSec, provides for end-to-end security. It can secure communications between a client and a server. When using the transport mode, only the IP payload is encrypted. AH or ESP provides protection for the IP payload. Typical IP payloads are TCP segments containing a TCP header and TCP segment data, User Datagram Protocol (UDP) messages containing a UDP header and UDP message data, and ICMP messages containing an ICMP header and ICMP message data.



EXAM DAY WARNING

Know and understand the differences between tunnel and transport modes in IPSec. Be aware of how each is used to make secure communications possible.

IPSec Protocols

As we mentioned earlier, IPSec itself is merely a framework within which a number of components work together. Those components include services, drivers, and protocols. IPSec uses many different protocols to provide various types of security for traffic that is passed through the network. The protocols that are used in a given IPSec communication session depend on several factors, such as whether you need data confidentiality or only authentication and integrity.

The primary IPSec protocols are ESP and AH. You can configure IPSec to use both of these protocols together to secure the data if you need both data encryption and integrity/authentication for the entire packet. Other IPSec protocols include ISAKMP, IKE, and Oakley, which uses the Diffie-Hellman algorithm.

Primary IPSec Protocols

ESP and AH can be used with both tunnel and transport mode. Which you choose depends on whether you wish to have data confidentiality. In the following subsections, we discuss each of these protocols in more depth.

EXERCISE 10.01

USING NETWORK MONITOR TO DETERMINE IPSEC PROTOCOL

In this exercise, you will learn how to determine which IPSec protocol is in use by using the Network Monitor. This exercise assumes that the Network Monitor has been installed via **Control Panel | Add/Remove Programs**.

1. Select **Start | Programs | Administrative Tools | Network Monitor**.
2. When the Network Monitor opens, begin the capture by either clicking the **Capture** button and selecting **Start** or by pressing the **F10** key.
3. Allow the capture to run for a few minutes. To stop it, either click the **Capture** button and then the **Stop and View** button, or press the **F11** key.
4. To view the IPSec protocol traffic on the captured packets, choose the **Display** and then select the **Captured Data** option.
5. Choose **Display | Filter Data**. Then choose **Edit Expression** option and select the **Protocol** tab.
6. All protocols are enabled by default. You can chose to **Disable All** and then reenale the AH and ESP traffic. Enabled traffic will appear in the left pane, and disabled traffic will appear in the right pane.
7. Click **OK** after the IPSec protocols have been enabled.
8. Select the **OK** option again, and the frames should be displayed in the Network Monitor window. Notice that when you open a packet that is IPSec-secured, you are unable to read the data inside.

ESP

ESP provides confidentiality (in addition to authentication, integrity, and anti-replay protection) for the IP payload. ESP in transport mode does not sign the entire packet. Only the IP payload (not the IP header) is protected. ESP can be used alone or in combination with AH (in order to provide for signing of the entire packet).



NOTE

IPSec is based on machine certificates, thus authentication pertains to only the *computer* from which the message was sent. IPSec cannot verify that data was sent from a particular *user* (although there are other mechanisms for doing so).

The ESP header is placed before the IP payload, and an ESP trailer and ESP authentication trailer are placed after the IP payload. The ESP header contains the following fields:

- **Security Parameters Index (SPI)** Used to identify which SA is used in conjunction with the security protocol and destination address. This value is used by the receiver to determine the packet identification.
- **Sequence Number** Provides anti-replay protection for the packet. The sequence number starts at 1 and increases in 32-bit increments. It is used to indicate the packet number sent over the quick mode SA for the communication. This number cannot be repeated. If a recipient gets a number that has been repeated, it will not accept the packet.

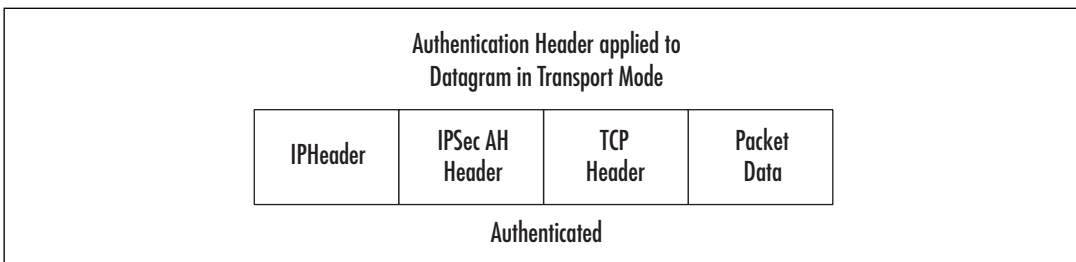
The ESP trailer contains the following fields:

- **Padding** Validates that byte boundaries are present on encrypted payloads. This process is required by the encryption algorithm.
- **Padding Length** Used to show the length, in bytes, of the Padding field.
- **Next Header** Used to identify whether the payload data is TCP or UDP.

The ESP authentication trailer contains the **Authentication Data** field, which holds the message authentication code, also known as the integrity check value (ICV). The ICV is used for message verification and authenticity. The ICV is calculated by the packet receiver and checked against the sender's value for integrity verification.

Figure 10.3 illustrates how ESP affects the data. You can see that the IPSec AH header has been placed after the IP header and before the TCP header.

Figure 10.3 The Effects of the ESP Header in Tunnel Mode



AH

AH does not provide confidentiality, which means that the data is not encrypted. Without data encryption, unauthorized people could use a sniffer-type program on your network to capture and read the packets, but they could not modify the data. AH works by using keyed hash algorithms, which are used to sign the packet for integrity verification.


Here is the process by which AH works:

1. Computer A sends data to Computer B.
2. The IP header, the AH header, and the data are signed to provide integrity.
3. The recipient at Computer B can be assured that the data was sent from Computer A and that the data arrived at the destination unmodified.

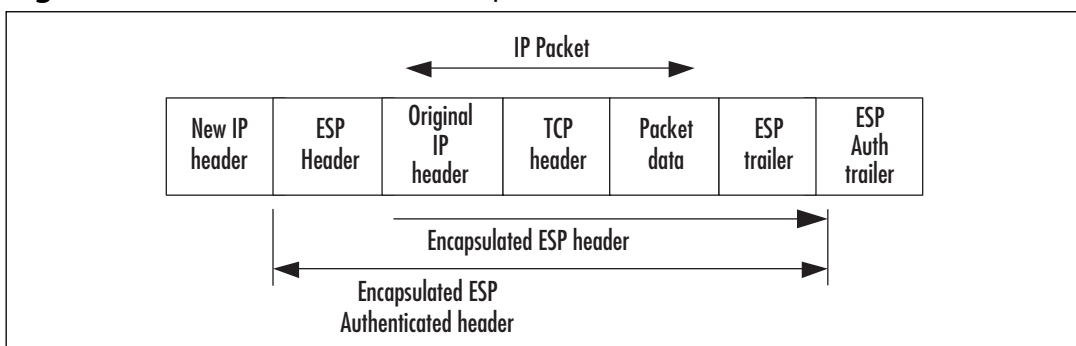
The AH header is placed between the IP header and IP payload to ensure integrity and authentication. AH can be used alone or combined with ESP. The AH header contains the following fields:

- **Next Header** Used to identify the IP payload via the IP protocol ID. The value here indicates the protocol (for example, TCP is represented by a value of 6).
- **Length** Used to indicate the length of the AH header.
- **SPI** A combination field that contains the destination address and the security protocol. This field is used to identify the correct SA for communication.
- **Sequence Number** Used to provide the packet with anti-replay protection. The sequence number starts at 1 and then increases in increments. The value in this field is a 32-bit number. For the life of the quick mode SA, the sequence number cannot repeat itself. If the receiver does a check on this field and finds that an SA with this number has been received in the past, the packet is denied.
- **Authentication Data** Used to verify message integrity and authentication using the ICV. The ICV value is checked and calculated by the receiver over the IP header, the AH header, and the IP payload.
- **Packet Signature with the AH Header** Used by AH to sign the entire packet. The packet is checked for integrity. The AH header will be inserted before any additional IPsec headers if other IPsec headers are present.
- **Packet Signature and Encryption** Used to protect IP payloads, as shown in Figure 10.4. The signed portion of the packet indicates the packet has been signed for integrity and authentication, and the encrypted portion of the packet indicates that the information itself is confidential.

The AH packet signature is shown in Figure 10.4.

Figure 10.4 AH Using Transport Mode

If you need both data integrity and authentication for the IP header, use ESP and AH in combination, as illustrated in Figure 10.5.

Figure 10.5 ESP Used with AH Transport Mode

EXAM DAY WARNING

Be able to differentiate between the AH and ESP IPsec protocols, and know how each of these protocols operate to make the data secure.

Additional Protocols

In addition to AH and ESP, the ISAKMP, IKE, and Oakley protocols and the Diffie-Hellman algorithm are used with IPsec. In the following subsections, we will discuss each of these in more detail.

ISAKMP and IKE

ISAKMP is used by IPsec as a key management system by combining the ISAKMP protocol and another protocol named IKE. IKE is used to centralize SA management and to generate and manage the secret shared keys that are used to secure data in transport.

There are two parts, or phases, involved in IKE SA establishment. The first phase is referred to as the *main mode SA*. In this first phase, the communicating IPsec-enabled sys-

tems create a secure channel. Based on the policies set on each system, they negotiate to determine which encryption algorithm, integrity algorithm, Diffie-Hellman group, and authentication method to use. The encryption algorithm can be DES or the more secure 3DES. The integrity algorithm will be one of two hashing algorithms: Message Digest 5 (MD5), which uses a 128-bit key, or the Secure Hash Algorithm 1 (SHA1), which uses a 160-bit key. Diffie-Hellman group 1, 2, or 2048 can be used, and the authentication method can be a pre-shared key, Kerberos, or digital certificates.

The systems will negotiate to use the most secure parameters that are supported by both. Thus, the final negotiation can range from a less secure channel that uses DES, Diffie-Hellman group 1 that provides only 768 bits of keying material, and a pre-shared key for authentication, to a more secure channel that uses 3DES, Diffie-Hellman group 2048, and certificate-based authentication in a Public Key Infrastructure (PKI) environment.

The main mode SA lifetime can be set to as short as 5 minutes up to a maximum of 48 hours. As more traffic is sent, a new *quick mode* is negotiated to create two new IPSec SAs for application traffic protection. When the main mode SA expires, by default, it is renegotiated as needed.

Often, firewalls, proxy servers, and security gateways must be configured to allow IPSec and IKE traffic to be forwarded. If the packets are not encrypted, the firewall, proxy server, or security gateway can inspect the packet contents or the TCP and UDP ports. If any type of modification has been made to the contents of these packets, the receiving IPSec computer will detect the modification and discard the packets.

In Windows 2000, a major drawback of IPSec was that it could not be used when one of the communicating computers was behind a NAT system. That is because NAT changes the IP headers when it translates multiple internal private IP addresses to a single public external address (which it does so that many computers can access the Internet via one public address). NAT has been an important mechanism for addressing the growing shortage of available public IP addresses, which is a limitation of the IPv4 protocol currently used for most Internet communications. Thus, many networks use NAT to reduce their need for additional public IP addresses.

However, Windows Server 2003's implementation of IPSec provides support for a new Internet specification that allows IPSec packets to be modified by a network address translator (NAT). This is called *NAT traversal*. IPSec's ESP packets can pass through NATs that allow UDP traffic. The IKE protocol automatically detects the presence of a NAT and uses UDP-ESP encapsulation to allow IPSec traffic to pass through the NAT.



TEST DAY TIP

The Windows Server 2003 family's implementation of IPSec finally provides support for NAT traversal, an Internet standard that allows IPSec packets to be modified by NAT. IPSec ESP packets can pass through NATs, which are configured to allow UDP traffic.

Oakley

Oakley is a key-determination protocol. It is used to define how to acquire keying material after it has been authenticated. The Diffie-Hellman algorithm is the basic mechanism for the Oakley protocol.

Diffie-Hellman

The Diffie-Hellman key-exchange algorithm is a secure algorithm that offers high performance, allowing two computers to publicly exchange a shared value without using data encryption. This exchanged information is protected with a hash function. The key itself is never exchanged by the two communicating machines, but each machine can generate the identical shared key.

The exchanged keying material that is shared by the two computers can be based on 768, 1024, or 2048 bits of keying material, known as Diffie-Hellman groups 1, 2, and 2048, respectively. The Diffie-Hellman key that is computed from the exchange is proportional to the strength of the Diffie-Hellman group. Longer key lengths that are created in conjunction with strong Diffie-Hellman groups make it more mathematically difficult to “crack” a secret key by brute force or other methods. Note that Diffie-Hellman does not provide authentication. For protection against man-in-the-middle attacks, identities are authenticated after the Diffie-Hellman exchange occurs. Diffie-Hellman algorithms can be embedded within a protocol that does provide for authentication.

IPSec Components

In addition to the protocols that operate within the IPSec framework, there are a number of operating system components involved in Microsoft’s implementation of IPSec. The major IPSec components that are installed with Windows XP and Windows Server 2003 family are the IPSec Policy Agent service and the IPSec driver.

IPSec Policy Agent

The IPSec Policy Agent is a service that resides on each computer running the Windows Server 2003 operating system. It is shown in the Service console as IPSec services. The IPSec Policy Agent begins when the system is started. This service has the following main functions:

- For Active Directory clients, the IPSec Policy Agent captures the appropriate IPSec policy. Domain member computers will have central IPSec policy information stored in Active Directory. It will then be cached in the local Registry of the computer to which the policy applies.
- For nondomain member computers, the IPSec Policy Agent retrieves the IPSec policy from the local Registry. The local Registry is used to store IPSec policy information for all nondomain member machines.

- The IPsec Policy Agent surveys the IPsec policy configuration for any changes. For computers not connected to the domain, the cached IPsec policies will be replaced with newer IPsec policies when the computer reconnects to the domain controller.
- The IPsec Policy Agent routes information to the IPsec driver.

For all domain member computers, the IPsec policy will be retrieved by the IPsec Policy Agent when the machine boots up or at the default Winlogon polling interval, unless an IPsec policy is in place that has the interval already set. Active Directory can be manually polled by typing the command **gpupdate /target:computer** at the command prompt.

If the IPsec Policy Agent is unable to find or connect to the Active Directory domain, it will wait for the policy to be activated or assigned. This is also true if there are no IPsec policies in Active Directory or the Registry.

IPsec Driver

The IPsec driver is used to match all packets against filters in the filter list. Once it finds a packet that matches the filter, it applies the appropriate filter action. If a packet does not match any filter, the packet is not changed and is sent back to the TCP/IP driver. The packet will then be either received or transmitted. After the transmission has been allowed by the filter action, the packet will be sent or received and not modified. If the packet is blocked by the filter action, it will be discarded. If the action requires security negotiation, main mode and quick mode SAs will be negotiated. The IPsec driver uses a database to store all current quick mode SAs. Any outbound packet that matches an IP filter list that is in need of security negotiation will be queued. After the packet has been queued, IKE is notified and will begin the security negotiation. After the negotiation has been successfully completed, the sending computer's IPsec driver will receive the session key from IKE. It will look in its database and locate the outbound SA, and then insert the SPI into the AH or ESP header. The packet will be signed, and if confidentiality is required, it will be encrypted and sent to the IP layer so it can be forwarded to the destination machine.

For inbound packets that match IP filters, the IPsec driver will receive the SA, session key, and SPI, and find the inbound SA in its database. The signature is then checked. If the packet was encrypted, it will be decrypted. The packet will search for the filter, and when it is found, it will send the packet to the TCP/IP driver so it can be forwarded to the receiver.

In summary, the IPsec driver plays a role in the following negotiation process:

- The sending computer's IPsec driver receives the SA containing the session key from IKE. Then it will locate the outbound SA in its database.
- The SPI is then inserted from the SA into the AH or ESP header.
- If confidentiality is required, the packets are encrypted; if not, the packets are signed.

- If a negotiation failed for some reason, the packets are not used and are discarded.

The IPSec driver will perform the following when the IPSec-secured inbound packet matches a filter in the IP filter list:

- The SPI and SA are received from IKE, and then the inbound SA is located in the database by destination address and SPI.
- The signature is checked and the packet is decrypted if needed.
- The IP packet then searches for a filter that matches the filter to make certain that no traffic that has not been agreed upon during the negotiation has been received.
- The packet is then sent to the TCP/IP driver for delivery to the receiving application.

When an unsecured IP packet is received, the IPSec driver looks for a matching filter in the filter list. If one is found, and the filter action for that filter either requires IPSec or blocks the packet, the packet will be discarded.

IPSec and IPv6

IPSec is an important part of the specifications for IPv6, which is supported by Windows Server 2003. As noted earlier, IPv6 is the “next generation” of IP, and its primary design goals were to create a larger address space to alleviate the shortage of IP addresses available under IPv4 and to provide for security of IP communications. IPSec is the means by which IPv6 provides the following:

- Authentication via the mechanism of digitally signing IPSec traffic with the shared encryption key so that the recipient of the data packet can verify that it was sent by the IPSec client transmitter
- Integrity via signing of the packet to ensure that any modifications made in transit will be detected by the recipient

IPSec and IPv6 work together to provide these services by using cryptographic security services. The Windows Server 2003 implementation of IPv6 does not support making data confidential by using data encryption. Keep this in mind when considering deploying IPSec and IPv6 within your network.

Deploying IPSec

With Windows Server 2003, Microsoft has made it relatively easy to deploy security for transmitted data throughout your organization by using the IP Security Policy Management MMC. However, before you begin to deploy IPSec on your network, you need to do your homework and determine the needs of your particular organization.

EXAM
70-293

OBJECTIVE

3.3.1

5

5.1

5.1.1

5.1.2

5.2.1

5.3.2

5.6

5.6.2

Determining Organizational Needs

It is very important to find a balance between protecting unauthorized access to data and choosing to make the information available to the largest group of users. The network administrator's dilemma is that security and accessibility are always at opposite ends of the continuum, and increasing one inevitably decreases the other.

To determine your organization's security policy needs, you should take the following steps:

- Assess the risk level.
- Determine the appropriate amount of security for your organization, based on risk level, acceptable risk, and accessibility needs.
- Define security policies that use your risk-management criteria and protect the identified information.
- Determine how you can best implement security policies within your organization.
- Identify the valuable and sensitive information on your network.
- Strive to provide all users with both secure and efficient access to the appropriate resources based on their computing needs.

After you've identified your organizational needs, you can begin to configure your policy. Only one policy configuration can be assigned at each of the following levels: domain, site, Organizational Unit (OU), and local level. Each IPSec policy consists of one or more IPSec rules. Each IPSec rule consists of the following:

- Selected filter list
- Selected filter action
- Selected authentication method or methods
- Selected connection type
- Selected tunnel setting

To configure IPSec policy, you can create a new policy, and then define the set of rules for the policy by adding filter lists and filter actions. Alternatively, you can create the set of filter lists and filter actions first, and then create the IPSec policies. Finally, you add rules that combine the appropriate filter list with the appropriate filter action. Additionally, you specify authentication methods, connection types, and tunnel settings.

Security Levels

When you begin to consider security levels within your organization, you must take into account the type of data each computer typically will be processing. For example, the configuration you would need for a Web server is different from the one you would need for a

domain controller. When planning to deploy IPSec on your network, take into account the following general guidelines for each type of computing environment:

- **Minimal security** No sensitive data is exchanged and IPSec is not active by default.
- **Standard security** This guideline is most appropriate for file servers and similar computers. You need these servers to be secure because of the data that is stored on them, but you need users to be able to access the data without complications. You can implement the Client (Respond Only) option or Server (Request Security) option for your IPSec policies. These policies enforce security when the client supports it, but they are also efficient because they do not require security if the client is not IPSec-enabled.
- **High security** The computers that need high security are the ones that contain sensitive or valuable data and/or are located in a public network setting. You can implement the Secure Server (Require Security) default policy on these machines. This requires IPSec protection for all traffic being sent to or received from the server (except initial inbound communication) with stronger security methods. Because unsecured communication with a non-IPSec-aware computer is not allowed, the server has a high level of security.

EXAM
70-293
OBJECTIVE
3.3.1
5
5.6.2

Managing IPSec

Windows Server 2003 comes with two handy tools for managing IPSec. These include the IP Security Policy Management MMC snap-in and the netsh utility (for those who love to use the command-line to execute commands).

IPSec policies are used to apply security at various levels within a network. IPSec policies can be applied to a computer, application, OU, domain, or site. You can create, modify, and add IPSec policies to these Active Directory objects using the IP Security Policy Management console. Each IPSec policy can be configured to store more than one rule, so different traffic types can be affected by each policy.

The following sections describe the IPSec management tools, and then how to manage IPSec policies.

Using the IP Security Policy Management MMC Snap-in

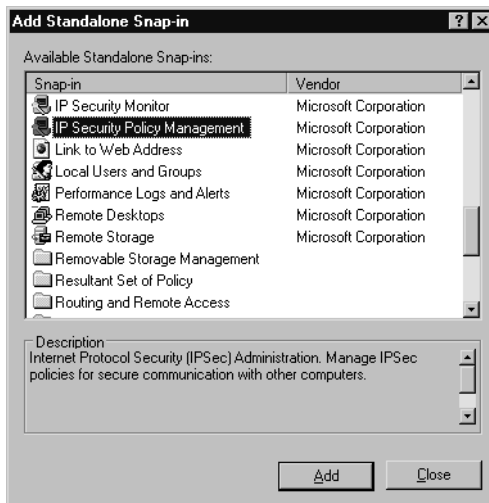
You can use the IPSec console to manage IPSec policies and to add and remove filters applied to the IPSec policies. IPSec filtering is used to permit or block certain types of IP traffic. With IPSec filtering, you can secure workstations from outside security hazards.

Follow these steps to install and access the IP Security Policy Management console:

1. Select **Start** | **Run**, type **mmc**, and click **OK**.

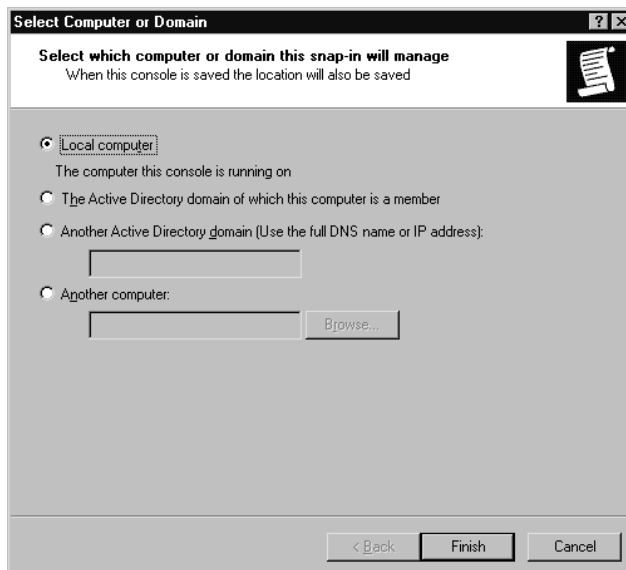
2. In the empty console, select **File | Add/Remove Snap-In**.
3. Click the **Add** button and scroll down to the **IP Security Policy Management** snap-in, as shown in Figure 10.6.

Figure 10.6 Add the IP Security Policy Management Console to the MMC




4. Click the **Add** button. The next window asks you to select the appropriate computer or domain that this snap-in will be used to configure. For this example, choose **Local computer**, as shown in Figure 10.7. Then click the **Finish** button.

Figure 10.7 Select the Computer or Domain to Manage



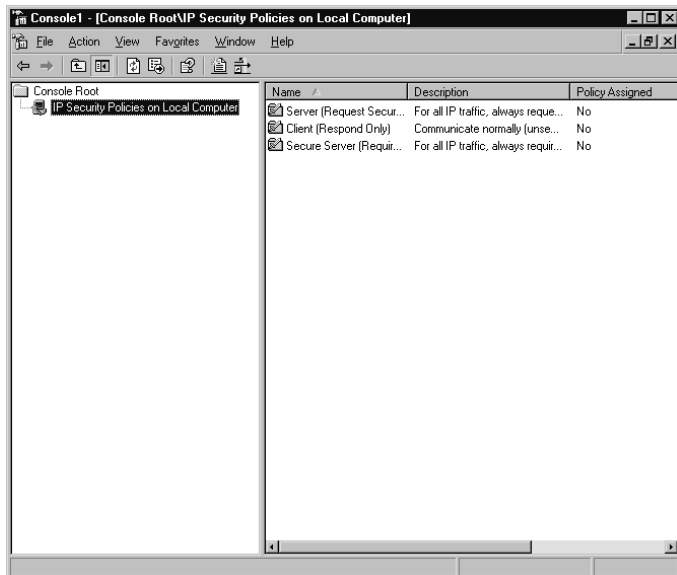
5. Select **Close**, and then click **OK**. The IP Security Policy Management console will open, as shown in Figure 10.8.

Figure 10.8 The Newly Created IP Security Policy Management Console



6. Double-click **IP Security Policies on Local Computer**. The three basic policy templates are now displayed in the right pane, as shown in Figure 10.9.

Figure 10.9 The Three Standard IPsec Policies in the IP Security Policy Management Console



Now you can use the IP Security Policy Management console to define, assign, and manage IPSec policies.

Using the netsh Command-line Utility

netsh is a command-line utility provided in Windows Server 2003 that you can use to control IPSec. This command can be used for managing advanced features of IPSec, including the following:

- Enabling IPSec driver event logging
- Configuring startup security on computers
- Viewing details of IPSec policies
- Troubleshooting IPSec configurations
- Setting default traffic exemptions

To use the netsh utility to manage IPSec, you need to change it to the *ipsec* context. Open a command prompt window (select **Start | Run**, type **cmd**, and click **OK**). In the command prompt window, type **netsh ipsec**. The IPSec command syntax you use at the prompt will depend on whether you are using IPSec static or dynamic mode commands. These two command modes have different functions in IPSec, as follows:

- **netsh ipsec static mode commands** Used to perform the same functions as the IP Security Policy Management and IP Security Monitor consoles. These commands allow you to create, modify, and assign IPSec policies, without affecting the current IPSec policy configuration.
- **netsh IPsec dynamic mode commands** Used to display the current state of IPSec; using this configuration will immediately affect the configuration of the IPSec policy.

Some netsh commands and switches are shown in Table 10.2. To view all of the available switches, type **netsh /?** at the prompt. All computers on which you wish to use the netsh utility for IPSec policy configurations must be members of the Windows Server 2003 family.



NOTE

You cannot use the netsh utility to configure IPSec on Windows XP machines. Instead, you must obtain the Windows XP installation CD and go to the Support/Tools folder, where you will find the Ipseccmd.exe utility. For IPSec policy configuration, you must use ipsecpol.exe, which is located in the Windows 2000 Server Resource Kit.

Table 10.2 netsh Command Switches

Command	Description
netsh ipsec static add policy <i>name</i>	Creates an IPsec policy with the specified name
netsh ipsec static delete all	Removes all IPsec policies, filter lists, and filter actions
netsh ipsec dynamic set policy <i>name</i>	Immediately sets a policy name
netsh ipsec dynamic delete policy <i>name</i>	Immediately removes a policy name
netsh ipsec dynamic export policy <i>name</i>	Immediately exports all IPsec policies to a specific file
netsh ipsec dynamic set policy <i>name</i>	Immediately sets a policy name



EXAM WARNING

Because the ability to use the *ipsec* context to manage IPsec with the netsh command-line utility is new to Windows Server 2003, it is likely that Exam 70-293 will contain questions on this topic. Be certain that you have a good understanding of all the *netsh ipsec* commands and know the difference between static mode and dynamic mode.

Default IPsec Policies

IPsec has a predefined set of default policies that can be implemented via the IP Security Policy Management console. The set includes Client (Respond Only), Server (Request Security), and Server (Require Security). The following sections explain the usage and settings for each default policy.

Client (Respond Only)

Client (Respond Only) is the least secure default policy. You might wish to implement this policy for intranet computers that need to respond to IPsec requests but do not require secure communications. If you implement this policy, the computer will use secured data communications when requested to do so by another computer.

This policy uses the default response rule, which creates dynamic IPsec filters for inbound/ outbound traffic based on the port/protocol requested. The policy settings are as follows:

- IP Filter List: All
- Filter Action: None
- Authentication: Kerberos

- Tunnel Setting: None
- Connection Type: All

Server (Request Security)

The Server (Request Security) policy consists of three rules and can be used when a computer needs to be configured to accept unsecured traffic from other computers that are not IPSec-enabled. However, it will always check for secure communication and use it if the other computer is able to use IPSec. The policy settings for the three rules are shown in Table 10.3.

Table 10.3 Policy Settings for Server (Request Security) Rules

Setting	First Rule	Second Rule	Third Rule (Default Response Rule)
IP Filter List	All IP Traffic	All ICMP Traffic	Dynamic
Filter Action	Request Security (Optional)	Permit	Default Response
Authentication	Kerberos	N/A	Kerberos
Tunnel Setting	None	None	None
Connection Type	All	All	All

Secure Server (Require Security)

The Secure Server (Require Security) policy consists of three rules and can be used for computers that require high security. Filters used in this policy require all outbound communication to be secured. This allows only initial inbound communication requests to be unsecured. The policy settings for the three rules are as shown in Table 10.4.

Table 10.4 Policy Settings for Secure Server (Require Security) Rules

Setting	First Rule	Second Rule	Third Rule (Default Response Rule)
IP Filter List	All IP Traffic	All ICMP Traffic	Dynamic
Filter Action	Require Security	Permit	Default Response
Authentication	N/A	Kerberos	Kerberos
Tunnel Setting	None	None	None
Connection Type	All	All	All



TEST DAY TIP

In order for Windows 2000 computers to use the 3DES algorithm, they must have the High Encryption Pack or Service Pack 2 or later installed. If a Windows 2000 computer receives a 3DES setting without having Service Pack 2 or the High Encryption Pack installed, the 3DES setting in the security method will be set to the weaker DES setting. Remember that DES is far less secure than 3DES, so this will not provide a level of security as high as when 3DES is supported.

Custom Policies

In addition to the default policies that can be implemented with the IPsec Security Policy MMC, you can also create your own custom policies for implementation by using the New IPsec Policy option in the IP Security Policy Management MMC.

To create your own custom policies with the IP Security Policy Management MMC, open the console and select the policy you wish to customize. See Exercise 10.02 for instructions on creating custom policies.

EXERCISE 10.02

CUSTOMIZING AN IPSEC SECURITY POLICY

1. Open the **IP Security Policy Management** console and click **IP Security Policies**.
2. Locate the policy you wish to customize in the right pane and double-click it, or right-click it and select **Properties**.
3. Click on the **Rules** tab, locate the rule you wish to modify and click **Edit**. Switch to the **Filter Action** tab, double-click the filter action that you want to modify.
4. Next, switch to the **Security Methods** tab, and do one of the following:
 - To add a new security method, select the **Add** option.
 - To modify an existing security method, select the security method that you want to modify and click the **Edit** option.
 - To remove a security method, click the security method that you wish to delete and select the **Remove** option.
5. To add or modify a security method, select the **Security Method** tab, choose the **Custom** option button, and then click **Settings**.

6. Set the security method as follows, depending on your policy's need for encryption:
 - Select the **Data and address integrity without encryption (AH)** check box if you need to provide data integrity for the packet's IP header and the data. Then for **Integrity algorithm**, select either **MD5** (which uses a 128-bit key) or **SHA1** (which uses a 160-bit key).
 - If you need to provide both integrity and encryption for data confidentiality, select the **Data integrity and encryption (ESP)** check box. Then under **Integrity algorithm**, click **None** (for no data integrity; if you have AH enabled and for increased performance, you can choose this), **MD5**, or **SHA1**. Under **Encryption algorithm**, choose **None**, **DES**, or **3DES**.
 7. You can also change the default session key lifetime settings, as follows:
 - You can set the number of kilobytes of data that is transferred before a new key is generated by choosing the **Generate a new key every** check box and typing in a value in kilobytes.
 - You can choose the **Generate a new key every** option to enter the number of seconds to elapse before a new session key is to be generated.
-



TEST DAY TIP

If you set the policy to use shorter key lifetime values, this will not increase the security level at which the data is protected. These short key lifetimes work by decreasing the amount of data that is revealed if an attacker discovers one encryption key.

Using the IP Security Policy Wizard

You can open the IP Security Policy Management console by clicking **Start** | **Run** and typing **mmc**, and then clicking **OK**. Select **File** | **Add/Remove Snap-in**, and then click **Add**. Click **IP Security Policy Management**, and then click **Add**. For each computer scenario, you need to select a specific option. Table 10.5 shows the scenario and specific snap-in you would need to use.

Table 10.5 IPSec Policy Management Scenarios

Scenario	Snap-In to Choose
Manage IPSec policy for local computer	Select the Local computer snap-in
Manage IPSec policies for any domain members	Select The Active Directory domain of which this computer is a member snap-in
Manage IPSec policies for a domain that this computer is not a member of	Select the Another Active Directory domain snap-in
Manage a remote computer	Select the Another computer snap-in

After you've chosen the snap-in, you can close the management console by selecting **Finish**, choosing **Close**, and clicking the **OK** button. To save your console settings select **File | Save**.

You can also access the IP Security Policy Management console from the Group Policy console. To do this, select **Start | Administrative Tools | Active Directory Users and Computers** and right-click the domain or OU for which you need to set Group Policy. (To open Active Directory Users and Computers utility, select **Start | Control Panel | Administrative Tools | Active Directory Users and Computers**.)



NOTE

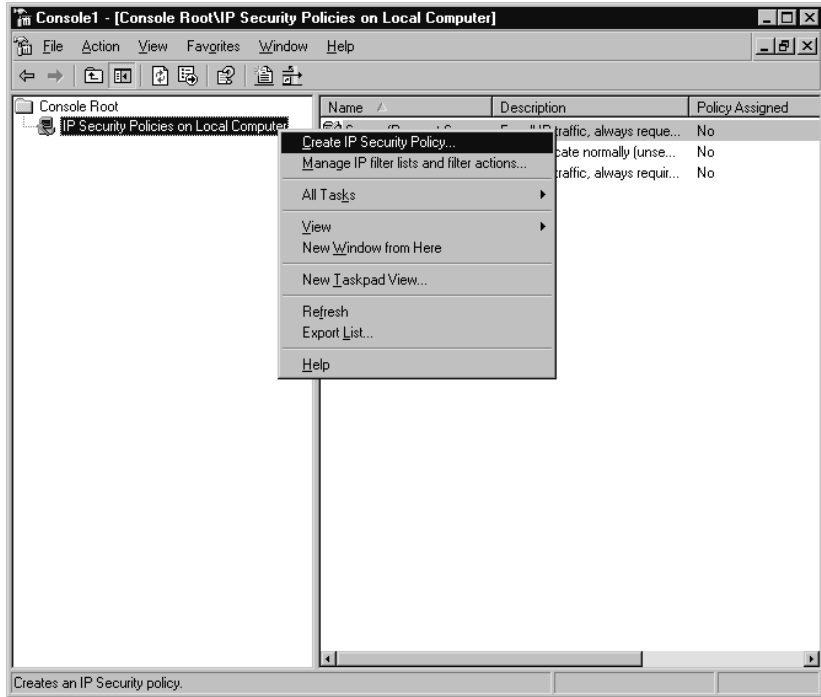
To save console settings, on the **File** menu, click **Save**, and then type in a name for the console.

Creating an IPSec Policy with the IP Security Policy Wizard

To create your own IPSec policy using the IP Security Wizard, follow these steps:

1. Open the IPSec Security Management Snap-in, right-click **IP Security Policies** in the left console pane, and then choose **Create IP Security Policy** from the context menu, as shown in Figure 10.10.

Figure 10.10 Creating a Custom IPsec Policy



2. The IP Security Policy Wizard Welcome window appears, as shown in Figure 10.11. Click the **Next** button.

Figure 10.11 The IP Security Policy Wizard.



- The IP Security Policy Name window appears, prompting you to give your IPSec policy a name and description, as shown in Figure 10.12. You can choose to accept the default name (not recommended, as it's not very descriptive), or you can enter a new name and description. Then click the **Next** button.

Figure 10.12 Enter a IP Security Policy Name

- The next window allows you to specify how the policy will respond to requests, as shown in Figure 10.13. Accept the default (**Activate the default response rule**) or clear the check box, and then click the **Next** button

Figure 10.13 Specify How the Policy Will Respond to Secure Communication Requests

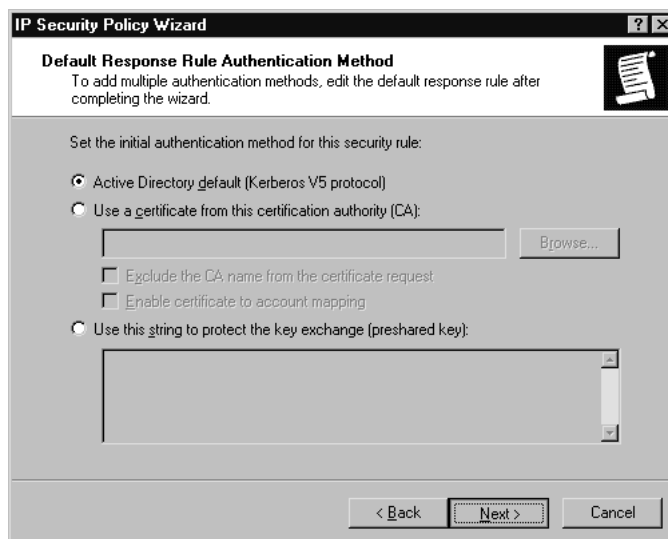
- The Default Rule Authentication Method window appears, as shown in Figure 10.14. Select a different authentication method or accept the default, **Active Directory default (Kerberos V5 protocol)**, and then click **Next**.



NOTE


Nothing special is required to use Kerberos authentication. If you select to use a certificate for authentication, you will need a PKI implementation and you must specify the certification authority to issue the certificate. If you select to use a pre-shared key, you must enter a string of characters that is also known to the party with which you are communicating.

Figure 10.14 Select the Default Rule Authentication Method




- The Completing the IP Security Policy Wizard window appears, as shown in Figure 10.15. You can choose to edit the properties of the policy (the default) or clear the check box if you do not wish to edit the properties at this time. Click **Finish** to complete the wizard. For this example, we will leave the **Edit properties** box selected.

Figure 10.15 Completing the IP Security Policy Wizard



- When you select the option to edit properties, the **New IP Security Policy Properties** dialog box opens, as shown in Figure 10.16. This dialog box allows you to edit the IP security rules and change the general properties of the rule, such as the name and description. Click the **Edit** button in this dialog box.

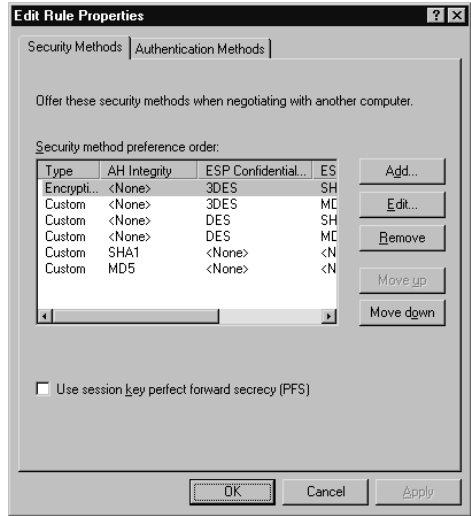
Figure 10.16 IP Security Policy Properties



- The Edit Rule Properties dialog box opens, as shown in Figure 10.17. Here, you can add, edit, or remove security methods; set the security methods that can be used when working with another machine; and select to use session key perfect forward secrecy (PFS). You can also arrange the order of precedence by using the

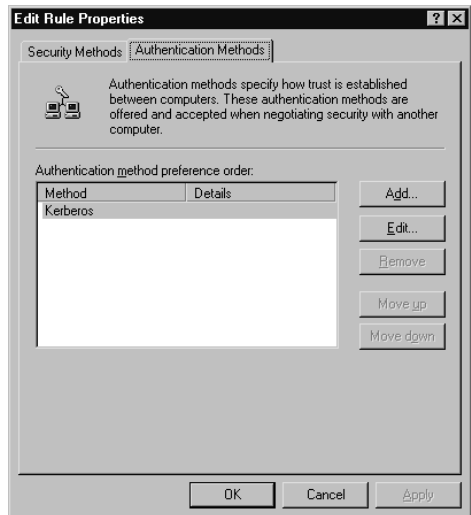
Move up and **Move down** buttons to change a method's position in the list. After making your selections, you can close the dialog box, or continue and select authentication methods. For this example, click the **Authentication Methods** tab.

Figure 10.17 Edit the IP Security Policy Security Methods



9. The **Authentication Methods** tab, shown in Figure 10.18, allows you to choose a trust method for communicating client computers. Click **Add** to add a method (again, your selections include using a certificate or a pre-shared key). You can change the order of precedence for these authentication methods in the same manner as described in Step 7. Click **OK** to close the dialog box.

Figure 10.18 Edit the IP Security Policy Authentication Methods



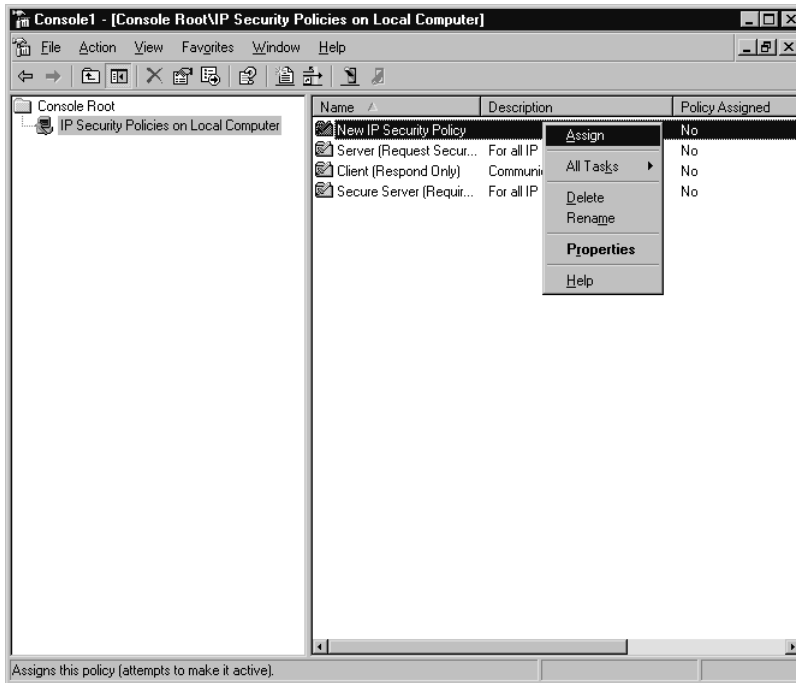
- After the policy has been edited, you need to assign the policy. Before you assign the policy, make sure that you have the IPsec service started. To assign the policy, right-click the policy name in the right pane and select **Assign**, as shown in Figure 10.19.



NOTE

The policy must be assigned before it can be used, and the IPsec service must be started before you assign the policy.

Figure 10.19 Assign the Newly Created IP Security Policy



EXAM WARNING

Ensure that you have the appropriate rights assigned to the account you will use to manage IPsec policies. To manage Active Directory-based IPsec policies, you must be a member of the Domain Admins group in Active Directory. To administer IPsec policies on a local or remote computer, you must be a member of the Administrators group on the local or remote computer.

Perfect Forward Secrecy

You can use *perfect forward secrecy* (PFS) to force reauthentication and negotiation of a new master key any time a new session key is required. There are two types of PFS used in Microsoft's IPsec implementation: master key PFS and session key PFS. Master key PFS should be used when it's needed for interoperability. By default, it is disabled. One reason is that it requires a lot of resources on the domain controller to perform the reauthentications (assuming Kerberos is the authentication protocol). Session key PFS is not as resource-intensive. Reauthentication is not required. You can configure PFS separately for master and session keys.

PFS doesn't determine when a new key is generated (as do key lifetimes). Instead, it is used to determine *how* new keys are generated, so that if one key is compromised, this won't compromise the entire communication. With PFS enabled, additional keys cannot be created from the keying material used to generate a particular key.

Defining Key Exchange Settings

You can define key exchange settings that apply to IP security policy. Open the MMC containing the security policy, and follow these instructions for modifying the policy:

1. Select the policy you wish to modify by double-clicking that policy.
2. Select the **General** tab and click the **Settings** button.
3. To force reauthentication and the negotiation of new master key keying material each time a new session key is required, click **Master key perfect forward secrecy (PFS)**.
4. To cause the reauthentication and new master key regeneration based on number of minutes, type in a value for **Authenticate and generate a new key after every number minutes**.

If you require a different setting, you can add a value in the **Authenticate and generate a new key after every number sessions**. This will set a maximum limit on the number of times a master key or its base keying material can be reused to generate the session key. When this limit is reached it will force a reauthentication with a new master key generation.

If you have enabled **Master key perfect forward secrecy (PFS)**, the number of sessions is set to **1** by default and cannot be reconfigured. For special requirements on the master key exchange, select the methods and use master key PFS where it is required for interoperability. By default, this setting is disabled, which should be appropriate in most environments. If you set the session limit to **0**, it will cause rekeys to be determined based

only on time. If you work in a performance-based environment, keep in mind that if you enable master key PFS, it could affect performance because each quick mode will require a new main mode negotiation.

Managing Filter Lists and Filter Actions

To manage IP filter lists and filter actions, open the IP Security Policy Management MMC and select the policy you wish to modify by double-clicking that policy. In the **Rules** tab, select the rule you wish to modify that contains the IP filter and double-click it. Select the **IP Filter List** tab and double-click the IP filter that contains the filter list you want to configure. Then do one of the following:

- Click **Add** to add a filter list.
- Select an additional filter that needs modifying and select **Edit**.
- To delete an existing filter, choose the filter and click the **Remove** button.

To edit or modify a filter in the IP Filter properties window, double-click the filter, choose the **Addresses** tab, and then select the **Source Address** drop-down box. Choose a source address as follows:

- **My IP Address** Secures packets from all IP addresses on the computer.
- **Any IP Address** Secures packets from any computer.
- **A specific DNS name** Secures packets from the Domain Name System (DNS) name that you specify in **Host name**. This is available only when creating new filters.
- **A specific IP address** Secures packets from only the IP address that you enter in **IP address**.
- **A specific IP subnet** Secures packets from the IP subnet indicated by the IP address that you specified in **IP address** and the subnet mask that you specify in **Subnet mask**.
- **DNS Servers dynamic** Secures packets from the DNS server that the computer is using. The filter is updated as needed, and it will automatically detect changes in the DNS server addresses.
- **WINS Servers dynamic** Secures packets from the WINS server that the computer is using. The filter is updated as needed, and it will automatically detect changes in the WINS server addresses.
- **DHCP Server dynamic** Secures packets from the DHCP server that the computer is using. The filter is updated as needed, and it will automatically detect changes in the DHCP server addresses.

- **Default Gateway dynamic** Secures packets from the default gateway that the computer is using. The filter is updated as needed, and it will automatically detect changes in the default gateway server addresses.

Select the **Destination Address** and repeat the same steps for the destination address. Next, select the desired **Mirrored** setting, as follows:

- To create two filters based on the filter settings, with one filter for traffic to the destination and one filter for traffic from the destination, select the **Mirrored** check box.
- To create a single filter based on filter settings, uncheck the **Mirrored** box.
- To create a filter for an IPSec tunnel, uncheck the **Mirrored** box and create two filter lists. The first filter list describes outbound traffic, and the other filter describes inbound traffic. Also, create two rules that use the inbound and outbound filter lists in the IP security policy.



NOTE

Mirrored IPSec filters are used to create two filters: one for traffic going to the destination and another filter for traffic coming from the destination computer.

Enter a description for the filter in the Description tab. To filter by a specific port or protocol, select **Configure advanced filter** settings on the Protocol tab.

When modifying IPSec rules, remember the following:

- Outbound packets that do not match any filter are sent unsecured.
- Inbound packets not matching any filters are allowed.
- Filters are applied in order, with the most specific followed by least specific.
- Filters are not applied in the order in which they appear in the filter list.
- Only address-based filters are supported.
- Protocol-specific filters are not supported.
- Port-specific filters are not supported.
- Tunnel filters should not be mirrored.
- IKE security requests result in the source IP address of the request being used to find a matching filter.
- IKE response is determined by the security action and tunnel settings that are associated with that particular filter.

- Filters used in tunnel rules are matched first.
- End-to-end transport filters are matched after tunnel rule filters have been matched.

Setting Up an IPSec Test Lab

You should set up an IPSec test lab with a server and a few client machines running the same operating system that your clients are using, so you can test IPSec policy configurations before deploying them on your production network. Use the lab to ensure that you can perform basic IPSec management tasks after you get the IPSec policies and filters set up.

Some of these tasks include the following:

- Secure Web traffic
- Secure ping
- Communication with a fallback server
- Communication with a secured server and communication with an IPSec/VPN connection

In a test lab, you can test and make changes to the environment without the possibility of causing a work stoppage on your live network. Be careful when rolling out IPSec, because misconfigured IPSec policies can shut down communications on your network.

Assigning and Applying Policies in Group Policy

Now we will take a look at how to assign or unassign IPSec policy in Group Policy for Active Directory. These settings will take effect the next time Group Policy is refreshed, and if a new policy is assigned over an existing policy, the current policy is automatically unassigned. Use the IP Security Policies on Active Directory within the Group Policy console to assign policies to apply to Active Directory objects. Follow these steps to assign or unassign IPSec policy in Group Policy for Active Directory-based Group Policy:

1. Click **Start | Administrative Tools | Active Directory Computers and Users** and right-click the domain or OU for which you want to set Group Policy.
2. Click **Properties**, and then click the **Group Policy** tab.
3. Select the Group Policy Object (GPO) you wish to modify and choose **Edit**. Alternatively, select **New** to create a new GPO (and type a descriptive name for it), and then click **Edit**.

4. From the Group Policy console tree in the left pane of the Group Policy Object Editor, under **Computer Configuration**, expand **Windows Settings**, and then expand **Security Settings**.
5. Select **IP Security Policies on Active Directory**.
6. In the right pane, click the IPSec policy that you want to assign or unassign. Click the **Action** menu (or right-click the policy), and then click **Assign** or **Unassign**.

To assign or unassign a local computer policy, select **Start | Run**, type **mmc**, and click **OK**. Then choose **File | Add/Remove Snap-in** and click **Add**. Click the **Group Policy Object Editor** and click **Add**. Choose **Finish**, click **Close**, and then click **OK**.



TEST DAY TIP

When dealing with IPSec policies, ensure that you unassign the IPSec policy before you delete the GPO or Group Policy. This is because an IPSec policy can remain active even after the GPO or IPSec policy that it has been assigned to has been deleted. To prevent these types of problems, unassign the IPSec policy and then make sure the change is effective by waiting at least 24 hours. Then delete the GPO or IPSec policy.

Active Directory Based IPSec Policies

Any IPSec policy that is applied for the domain will take precedence over local IPSec policy that is located on the member computer. After the IPSec policy has been applied to one of the Active Directory Group Policy Objects, it will be broadcast to all of the computer accounts that are affected by that GPO. When you wish to apply an IPSec policy within your Active Directory network, remember the following guidelines:

- OU IPSec policy assignments will take precedence over domain-level policies for members of that OU.
- Although the entire list of IPSec policies is available to assign at any level in the Active Directory structure, only a single IPSec policy can be assigned at a specific level (site, domain, or OU) in Active Directory.
- An IPSec policy that is assigned to the lowest level OU in the domain structure will override an IPSec policy that is assigned to a higher-level OU for computers that belong to that OU.
- Unless a policy is blocked or unassigned, OUs will inherit the policies of their parent OUs.
- IPSec policies from different OUs can never merge.

- The highest possible level of the Active Directory structure should be used to assign policies. Just as with Group Policy assignment, an IPSec policy might remain active even after the GPO to which it was assigned has been deleted. Ensure that you unassign the policy before deleting the GPO. You should unassign the IPSec policy in the GPO, wait 24 hours, ensure that the change has taken effect, and then remove the GPO.

Group Policy has backup and restore tools that you can use to save policy information on assigned GPOs. These tools *do not* back up the IPSec policies. To back up and restore IPSec policies, use the Export Policies and Import Policies command in the IP Security Policy Management console. The Group Policy console will back up and restore only information pertaining to the IPSec policy assignments in relation to GPOs.

The IPSec Policy Agent on client computers running Windows XP Professional or a Windows Server 2003 operating system will poll Active Directory for updates to the assigned IPSec policy. This does not detect domain or OU changes or whether new IPSec policies have been assigned. The Winlogon service polls for these changes every 90 minutes. If a change has been made, the Winlogon service will notify the IPSec Policy Agent, and the IPSec policy changes will be applied.



NOTE

You cannot administer Active Directory-based IPSec policies from Windows XP Home Edition computers. Only Windows XP Professional Edition computers can be members of the domain.

Cached IPSec Policy

A copy of the currently assigned IPSec policy for a site, a domain, or an OU is cached in the local Registry of each computer to which it applies. If the computer that has the IPSec policy assigned cannot log on to the domain for any reason, the cache copy will be applied. The cache copy of the IPSec policy cannot be changed or managed.

Local Computer IPSec Policy

All Windows Server 2003 servers and Windows XP Professional computers have one local GPO called the local computer policy. With this local policy, Group Policy settings can be stored on individual computers, even when they are not Active Directory domain members. You can manage the local IPSec policy by using the IP Security Policy Management console. Alternatively, you can use the following netsh command at the prompt:

```
netsh ipsec static set store location=local
```

If a computer on which you've applied local IPsec policies later joins an Active Directory domain that has IPsec policies applied, the domain policies will override the local IPsec policy.

IPsec Monitoring

It is important for network administrators to monitor IPsec settings and traffic on a regular basis after deploying IPsec. You can perform monitoring with the netsh command-line utility or with the IP Security Monitor MMC snap-in. In the following sections, we will look at each of these tools.

Using the netsh Utility for Monitoring

Earlier in the chapter, we discussed the use of the netsh command-line utility as equivalent to the IP Security Policy Management console. However, the netsh utility provides some features that are not available with the IP Security Policy Management console. These include the following:

- IPsec diagnostics
- Client computer startup security
- Client computer startup traffic exemptions
- Default traffic exemptions
- Strong certificate revocation list checking Certificate Revocation List
- IKE/Oakley logging

netsh Dynamic Mode Policy

If you want the IPsec rules you have configured to take effect without any wait time, you can use the *netsh ipsec dynamic* commands at the command prompt to add, modify, and assign IPsec policies immediately. Dynamic policies, as their name implies, are not saved; they will be lost if the IPsec service is stopped. However, not all dynamic policies take effect immediately. In some cases, you must restart the computer or the IPsec service first. If you need to make these changes permanent, you need to use the *netsh ipsec dynamic set config* command. This will ensure that the changes are not lost if the computer is restarted.



WARNING

Use of dynamic mode commands is recommended only for network administrators who understand IKE main and quick mode policies. You can cause problems by creating invalid IPsec policies with the dynamic mode commands if you do not have a good understanding of what you're doing.

IPSec Diagnostics

You can use the *netsh diag* command with additional diagnostics at the command prompt. The following are the additional diagnostics switches:

- **netsh diag connect** Used to connect to mail, news, and proxy servers.
- **netsh diag dump** Used to display a script that is used for configuration.
- **netsh diag show** Used to show computer, operating system, network, news, mail, and proxy server information.
- **netsh diag gui** Used to display diagnostics on a Web page. Once this command has been run, you can scan the computer for network diagnostics.



NOTE

Remember that you must type the **netsh ipsec** command at the command prompt, to enter the *ipsec* context, before typing any additional commands.

Here are two important things to remember when using the *netsh* utility:

- If you stop the IPSec service when configuring a dynamic policy, you will lose the settings.
- Use caution because some commands will require you to stop and restart the IPSec service.


Using the IP Security Monitor MMC Snap-in

Microsoft provides the IP Security Monitor MMC snap-in for monitoring IPSec activity. To use the IP Security Monitor, open the MMC and add the IP Security Monitor to the console. We will discuss the use of the IP Security Monitor in more detail in the next section, which covers troubleshooting IPSec.



NOTE

Unlike the **netsh ipsec** commands, which can be used only with Windows Server 2003 computers, you can use the IP Security Monitor to monitor IPSec activities on Windows XP computers as well as Windows Server 2003 systems. For computers running Windows 2000, however, you must use the **ipsecmon** command.



Troubleshooting IPSec

Troubleshooting is always a big part of any network administrator's job. The following sections will cover how to troubleshoot your IPSec configuration. We include tables that will list specific tools and scenarios you can use to perform the troubleshooting tasks. The IP Security Monitor and the Network Monitor are important tools for troubleshooting IPSec problems, as are the IP Security Policy Management MMC and the netsh utility. An additional tool that is introduced in this section is the Network Diagnostics Tool, netdiag.exe.

Using netdiag for Troubleshooting Windows Server 2003 IPSec

The netdiag tool is provided on the Windows Server 2003 family servers, Windows XP, and Windows 2000 machines. However, it is stored in different locations on each platform as described below:

- **Windows Server 2003 family** On the Windows Server 2003 installation CD, locate the Support/Tools folder and run the **Suptools.msi** installation package with the **Complete** option to install the tool.
- **Windows XP Professional** On the Windows XP Professional installation CD, locate the Support/Tools folder and run the **Setup.exe** file with the **Complete** setup option to install the tool.
- **Windows 2000** Download the updated version of the tool from the Microsoft Web site.

New & Noteworthy...

Stateful Filtering

In the Windows Server 2003 version of IPSec, more enhanced security is provided during computer startup by using the *stateful filtering* feature. This filtering occurs during startup and allows only the following three types of traffic:

- DHCP
- Outbound traffic that the machine has initiated during startup
- Inbound traffic that is sent in response to the allowed outbound traffic

Another option for enhanced security is to configure the computer to not allow any traffic before an IPSec policy has been applied. With any of these options, you can exempt specific types of traffic from filtering if you wish. The stateful filtering option can be configured only at the command prompt with the netsh utility. The command for performing this task is *netsh ipsec dynamic set bootexemptions*. After this command has been executed, you will need to restart the computer.

Viewing Policy Assignment Information

The Policy Assignment option allows you to view policy assignment and precedence. For troubleshooting, it is often important to be able to view IPsec policy assignments and determine the precedence in which policies are applied. Table 10.6 shows a list of the tools to be used with different Microsoft operating systems for viewing the IPsec policy name viewing the Group Policy object to which the IPsec policy is assigned.

Table 10.6 Viewing the IPsec Policy Precedence on Windows Server 2003 Family Machines

Operating System	IPsec Viewing Tools	IPsec Policy Assignment for Group Policies
Windows Server 2003	IP Security Monitor console or the netsh command: <i>netsh ipsec static show gpoassignedpolicy</i>	Resultant Set of Policy (RSOP) console or the netsh command <i>netsh ipsec static show gpoassignedpolicy</i>
Windows XP	IP Security Policy Management console for local IPsec policy viewing	<i>netdiag.exe netdiag /test:ipsec</i> command. <i>netdiag.exe</i> command <i>netdiag /test:ipsec:ipsec</i>
Windows 2000	<i>netdiag.exe</i> command: <i>netdiag /test:ipsec</i> Go to the properties option in the TCP/IP network connections and select Properties Advanced Options IPsec . The assigned IPsec policy that is shown is the global policy.	<i>netdiag.exe</i> command: <i>netdiag /test:ipsec</i> <i>gpresult.exe</i> Group Policy Results <i>gpoutil.exe</i> Group Policy Verification Tool (these can be downloaded from the Windows 2000 Server Resource Kit Web site)

Additionally, you can view all IPsec policies that are available by using the IP Security Policy Management console. Just because an IPsec policy is available, this does not mean that it has been assigned or applied to a computer. In the Windows Server 2003 family, you can determine the assigned (but not applied) policies on IPsec clients by using the RSOP console. RSOP is discussed in more detail later in this chapter, in the “Using RSOP for IPsec Planning” section.

NOTE

If you try to use the RSOP console in Windows XP Professional, it will not display the IPsec policies, and the *gpresult /scope computer* command will not display the GPO that contains the IPsec policy assignment. Use the *netdiag /test:ipsec* command to view the GPO to which the IPsec policy is assigned on Windows XP Professional or Windows 2000 Client machines. The Group Policy Tool, *gpoutil.exe*, is used to monitor the health of GPOs on Windows 2000 domain controllers only.

Viewing IPSec Statistics

To view IPSec statistics and items such as filters and security associations, use the tools listed in Table 10.7. These tools work on Windows Server 2003, Windows 2000, and Windows XP Professional machines.

Table 10.7 Viewing IPSec Policy and IP Statistic Details

Operating System	Group Membership Required	Tools
Windows Server 2003 family	Administrators group on that server	IP Security Monitor console or the netsh command <i>netsh ipsec dynamic show all</i>
Windows XP Professional	Administrators group on the local computer	IP Security Monitor console or the IPseccmd.exe command <i>ipseccmd show all</i> at the command prompt
Windows 2000	Administrators group for the debug command. If you need to view ActiveDirectory-based IPSec policies, you must be a member of the Domain Admins group in Active Directory. IPsecmon.exe displays outbound quick mode security associations.	Netdiag.exe command <i>netdiag /test:ipsec /v /debug ipsecmon.exe</i>

To monitor IPSec policies on a remote computer that is running Windows XP or Windows Server 2003, you can use the Remote Desktop Connection (RDC) to connect to that computer and view its policies as if you were sitting at its desktop. You can do this from any computer that has the RDC client or the Windows 2000 Terminal Services client installed. You can connect remotely to a Windows 2000 server that is running Terminal Services in the same way. However, you cannot connect remotely to the desktop of a computer running Windows 2000 Professional or Windows 9x.

Using IP Security Monitor to View IPSec Information

For Windows Server 2003 and Windows XP, the IP Security Monitor is implemented as an MMC snap-in. This MMC snap-in allows administrators to view details regarding active IPSec policies that have been applied by the domain or applied locally, the quick mode and main mode statistics, and the active IPSec SAs. You can use the IP Security Monitor to search for specific main mode or quick mode filters and to troubleshoot complex IPSec policy configurations, as well as for filter searches that match a certain traffic type. To view IPSec information on computers running Windows 2000, you need to use the **ipsecmon.exe** command at the **run** prompt.

To access the IPSec Security Monitor on Windows Server 2003 and Windows XP clients, follow these steps:

1. Select **Start** | **Run**, enter **mmc**, and click **OK**.
2. In the console, select **File** | **Add/Remove Snap-In**.
3. Click the **Add** button, scroll down and click the **IP Security Monitor** snap-in.
4. Select **Add**, select the **Close** button, and click **OK**.
5. You can now add the local computer or browse to a computer on the network by right-clicking the IP Security Monitor console and selecting the **Add Computer** option.
6. When the computer has been added, you can view active policy information by double-clicking **Active Policy**.
7. You can view main mode and quick mode statistics by double-clicking these options in the console.



EXAM WARNING

Only computers running Windows XP Professional or the Windows Server 2003 operating system can use the Security Monitor. When monitoring IPSec remotely, the computer that is being monitored by the IP Security console must run the same version of the Windows operating system as the computer that the IP Security Monitor console is running. For Windows 2000 clients, type **ipsecmon** at the command prompt to open the console.

Using Event Viewer to Troubleshoot IPSec

Event Viewer is a great troubleshooting tool to use to view IPSec information. However, most IPSec-related information will be contained in the Security log, which is not enabled by default. Verify that security auditing is enabled so security events will be entered in the Security log. For domains, use the Group Policy Editor. For local computers, use the Local Security Policy setting for this procedure. When enabling auditing for Windows Server 2003 machines, you can also turn on the auditing for the security policy database (SPD). Next, you need to edit the audit policy on your domain or local computer. Enable success or failure auditing for **Audit logon events** to allow Event Viewer to record this information.

After you have enabled security auditing and configured the audit policy, Event Viewer will record as separate events the following information:

- Success or failure of each main mode negotiation
- Success or failure of each quick mode negotiation

- Establishment of each negotiation
- Termination of each negotiation

Your Security log will fill up with IKE events, so you might wish to edit the Registry and disable auditing of IKE events by creating the **DisableIKEAudits** value.



NOTE

Remember to exercise extreme caution when editing the Registry. One misstep can render your system unbootable. It is always a good idea to back up the Registry before editing it.

To disable auditing of IKE events, perform the following steps:

1. Open the Registry Editor by selecting **Start | Run**, typing **regedit** or **regedt32**, and clicking **OK**.
2. Navigate to **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Audit**.
3. Right-click the **Audit** key, select **New**, and then choose **DWORD Value**.
4. In the right pane, change the default name of the new value to **DisableIKEAudits**.
5. Double-click the new value, or right-click and select **Modify**.
6. In the **Edit DWORD Value** dialog box, under **Value data**, type **1**. Then click the **OK** button and close the Registry Editor.

After this modification has been completed, you can stop and restart the IPSec service or restart the system to have the new Registry information read.

Using Packet Event Logging to Troubleshoot IPSec

You can enable packet event logging for the IPSec driver in Windows Server 2003, Windows XP Professional, and Windows 2000 Server by modifying the Registry. This will cause the System log to capture logging information on all dropped inbound and outbound packets. This information can be useful in troubleshooting IPSec problems.

To enable logging of inbound and outbound packets, perform the following steps:

1. Open the Registry Editor by selecting **Start | Run**, typing **regedit** or **regedt32**, and clicking **OK**.
2. Navigate to **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec**.
3. Right-click the **IPSec** key and select **New**, and then choose **DWORD Value**.

4. In the right pane, change the default name of the new value to **EnableDiagnostics**.
5. Double-click the new value, or right-click and select **Modify**.
6. In the **Edit DWORD Value** dialog box, under **Value data**, type 7 and click the **OK** button.
7. Close the Registry Editor.

After you've made this change, restart the computer.

You can also enable IPsec driver logging of dropped inbound and outbound packets by using netsh command-line tool utility. From a command prompt window, issue the following command:

```
netsh IPsec dynamic set config ipsecdiagnostics 7
```

Next, restart the computer so that the settings will take effect.

By default, the IPsec driver will write to the System log on an hourly basis, or after the event threshold value has been met. For troubleshooting purposes, you can change this setting to an interval of 60 seconds. To change this setting, you can modify the Registry by creating the following DWORD value:

1. Open the Registry Editor by selecting **Start | Run**, typing **regedit** or **regedt32**, and clicking **OK**.
2. Navigate to **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPsec**.
3. Right-click the **IPsec** key and select **New**, and then select **DWORD Value**.
4. In the right pane, change the default name of the new value to **LogInterval**.
5. Double-click the new value, or right-click and select **Modify**.
6. In the **Edit DWORD Value** dialog box, under **Value data**, type **60**.
7. Under **Base**, click the **Decimal option** button.
8. Click the **OK** button.
9. Close the Registry Editor.

After you've made this change, you can restart the system.

Again, you can use a netsh command to change this setting. Open the command prompt window and type the following command:

```
netsh ipsec dynamic set config ipsecloginterval 60
```

Then restart the computer so the changes can take effect.

Packet event logging is disabled by default. After you create the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPsec\EnableDiagnostics** value as described earlier, you can control the logging level by editing the value. Table 10.8 lists the possible values that you can set. To disable logging altogether after the

DWORD value has been created without deleting the value (if you will want to enable it again later), set the value to **0**.

Table 10.8 Value Settings and Level of Logging

Value	Logging Performed
1	Bad SPI, IKE negotiation failures, and invalid packet syntax are logged.
2	System log records the inbound per-packet drop events.
3	Unexpected cleartext events and level 1 and level 2 logging are performed.
4	Outbound per-packet drops are recorded.
5	Level 1 and level 4 logging are performed.
6	Level 2 and level 4 logging are performed.
7	All logging is performed.

The value of 7 enables all logging, creating a great deal of information in the logs. Before you enable logging of this magnitude, realize that your system logs will fill up quickly. To prevent problems, do one or more of the following:

- Set your system log size to at least 10MB.
- Clear all events so the log is empty before you start logging.
- Save the current log to a file.

Using IKE Detailed Tracing to Troubleshoot IPSec

Enabling audit logging for IKE events and viewing the events in Event Viewer provide the fastest and simplest way to troubleshoot failed main mode or quick mode negotiations. If you need a more detailed analysis of these negotiations, you can enable tracing for IKE negotiations. This is an extremely detailed log intended for troubleshooting IKE interoperability under controlled circumstances. Before you try to decipher the log, you will need to have expert-level knowledge of RFCs 2408 (defining ISAKMP) and 2409 (defining IKE).

The IKE tracing log is 50,000 lines long and will overwrite if necessary. This log is located in the *systemroot*\Debug\Oakley.log file. Each time the IPSec service is started, the previous version of the file is renamed *Oakley.log.sav*, and a new *Oakley.log* file is created. If the *Oakley.log* file becomes full before the IPSec service is started, the full log will be named *Oakley.log.bak*, and a new *Oakley.log* file will be created.

You might wish to minimize the number of negotiations because many of these can occur at the same time. This will make your log file easier to read. See Table 10.9 for scenarios and explanations regarding the IKE tracing log. The *Oakley* key does not exist in the specified Registry tree. To use these settings, you must first create a new key named **Oakley**, and then create the new **EnableLogging** DWORD value within that key.

Table 10.9 IKE Tracing Log Scenarios

Enable/Disable IKE Tracing Log	Operating System	Registry Setting to Enable IKE Tracing	netsh Command to Enable the IKE Tracing Log	IPSec Service Status
Enable	Windows Server 2003	N/A	<i>netsh ipsec dynamic set config ikelogging 1</i>	Remain started
Disable	Windows Server 2003	N/A	<i>netsh ipsec dynamic set config ikelogging 0</i>	Remain started
Enable	Windows XP Professional	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging DWORD Registry setting to a value of 1	N/A	Stop and restart the IPSec service by using <i>net stop policyagent</i> and <i>net start policyagent</i> at the command prompt
Disable	Windows XP Professional	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging DWORD Registry setting to a value of 0	N/A	Stop and restart the SIPSec service by using <i>net stop policyagent</i> and <i>net start policyagent</i> at the command prompt
Enable	Windows 2000	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging DWORD Registry setting to a value of 1	N/A	Stop and restart the IPSec service by using <i>net stop policyagent</i> and <i>net start policyagent</i> at the command prompt
Disable	Windows 2000	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\PolicyAgent\Oakley\EnableLogging DWORD Registry setting to a value of 0	N/A	Stop and restart the IPSec service by using <i>net stop policyagent</i> and <i>net start policyagent</i> at the command prompt

Using the Network Monitor to Troubleshoot IPsec

The Windows Server 2003 Network Monitor is a protocol analyzer (also called a *packet sniffer*) that Microsoft includes with its server operating systems.



NOTE

The version of Network Monitor that is built into Windows can be used to view IPsec traffic only on the computer on which you are running the Network Monitor utility. If you need to view network traffic on other computers, you can use the version of Network Monitor that is included in Microsoft's Systems Management Server (SMS), which allows you to place the computer's NIC in promiscuous mode so that it will capture traffic on the network that is not sent to or from the local computer.

The Network Monitor includes parsers for the AH, ESP, and ISAKMP (IKE) IPsec protocols. However, the Network Monitor cannot parse the encrypted portions of IPsec-secured ESP traffic when encryption is software-based. If you are using encryption on a hardware offload network adapter, ESP packets are decrypted when the Network Monitor captures them and therefore can be parsed and interpreted into the upper-layer protocols. The following types of traffic should be exempt from filtering:

- Broadcast
- Multicast
- IKE
- Kerberos
- RSVP

IPsec will exempt all multicast, broadcast, RSVP, Kerberos, and IKE traffic if you are using Windows XP and Windows 2000. The Windows Server 2003 family only exempts IKE traffic from traffic filtering by default. Actions such as block, configure, and permit filter actions can be configured just for broadcast and multicast traffic. SAs will not be negotiated for broadcast and multicast traffic. If you wish to change the filtering behavior on your Windows Server 2003 machines to match the default behavior on Windows 2000/XP machines (that is, to exempt multicast, broadcast, RSVP, and Kerberos traffic, along with IKE), you can use the following netsh command at the prompt on the Windows Server 2003 machine:

```
netsh ipsec dynamic set config ipsecexempt 0
```

After issuing this command, you will need to reboot the computer for the changes to take effect.



NOTE

To display monitoring information such as policy settings and statistics on Windows XP machines, use **ipseccmd.exe** with the **show all** command.

By design, Windows 2000 and Windows XP default exemption settings for IPSec are configured for low-risk environments, such as corporate LANs, because the risk of attack is minimal. The Windows 2000 and Windows XP default exemption settings should be used in only low-risk environments and be applied only when necessary for troubleshooting purposes.

To exempt all multicast, broadcast, RSVP, Kerberos, and IKE traffic from IPSec filtering, you need to edit the Registry to create a DWORD value called **NoDefaultExempt** in the **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSEC** Registry key and set its value to **0**. Follow the instructions given previously for creating new DWORD values.

Disabling TCP/IP and IPSec Hardware Acceleration to Solve IPSec Problems

IPSec offload is a process by which some network adapters can do the processing for the mathematical calculations involved in encrypting IPSec data and TCP checksums. This speeds up, or *accelerates*, the process because it is being handled by a chip on the network interface card (NIC) instead of by the operating system software. NICs that are capable of offloading IPSec cryptographic functions can also perform a *large-send offload*, which is the processing of very large TCP segments for accelerated transmissions. If a Plug and Play NIC has this capability, its driver can make an advertisement to IPSec and TCP/IP. This results in the protocols passing these tasks to the NIC driver.

Although hardware acceleration speeds up processing, it can sometimes cause problems with packet processing. Exercise 10.03 walks you through the steps of disabling hardware offload functions.

EXERCISE 10.03

DISABLING HARDWARE OFFLOAD FUNCTIONS

Before you begin to test your network adapter, verify that you have the latest software drivers for the adapter. To disable TCP/IP hardware acceleration, follow these steps:

1. Open the Registry Editor by selecting **Start | Run**, typing **regedit** or **regedt32**, and clicking **OK**.

2. Navigate to **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters**.
3. Right-click the **Parameters** key, Select **New**, and choose **DWORD Value**.
4. In the right pane, change the default name of the new value to **DisableTaskOffload**.
5. Double-click the new value, or right-click and select **Modify**.
6. In the **Edit DWORD Value** dialog box, under **Value data**, type **1** and click the **OK** button.
7. Close the Registry Editor.

To disable IPSec hardware acceleration, follow these steps:

1. Open the Registry Editor by selecting **Start | Run**, typing **regedit** or **regedt32**, and clicking **OK**.
2. Navigate to **HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IPSec**.
3. Right-click the **IPSec** key, select **New**, and then choose **DWORD Value**.
4. In the right pane, change the default name of the new value to **EnableOffload**.
5. Double-click the new value, or right-click and select **Modify**.
6. In the **Edit DWORD Value** dialog box, under **Value data**, type **0** and click the **OK** button.
7. Close the Registry Editor.

After making these modifications, you will need to restart the computer.

EXAM 70-293
OBJECTIVE
3.3.1
5
5.2
5.7

Addressing IPSec Security Considerations

As you begin to deploy IPSec throughout your organization, you will need to decide on the encryptions methods you wish to implement and whether to use firewall packet filtering. The following sections provide some guidelines to use when considering IPSec security.

Strong Encryption Algorithm (3DES)

Earlier in the chapter, we discussed the two encryption algorithms supported by IPSec for data encryption: DES and 3DES. The 3DES algorithm is the strongest of these, using three

unique 56-bit keys. In a high-security environment, the 3DES algorithm is the appropriate choice for encrypting your data.

DES and 3DES are *block ciphers*. This refers to an algorithm that takes a block of plaintext of a fixed length and changes it into a block of *ciphertext* (encrypted data) of the same length. The key length for DES is 64 bits total, but because 8 of the bits are used for parity information, the effective length is only 56 bits. With 3DES, the DES process is performed three times with different 56-bit keys, making the effective key length 168 bits. When using 3DES in encrypt-encrypt-encrypt (EEE) mode, 3DES works by processing each block as follows:

1. A block of plaintext is encrypted with key one.
2. The resulting block of ciphertext is encrypted with key two.
3. The result of step 2 is encrypted with key three.

When using 3DES in encrypt-decrypt-encrypt (EDE) mode, step 2 is run in decryption mode. When 3DES is decrypting a packet, the process is done in reverse order. 3DES offers you the best mode for data confidentiality.

Firewall Packet Filtering

To allow for secured packets to be passed through a firewall, you need to configure the firewall or other device, such as a security gateway or router, to allow these packets to pass through the external interface.

The following ports and protocols can be used for firewall filtering:

- IP protocol and port 50, ESP traffic
- IP protocol and port 51, AH traffic
- UDP port 500, IKE negotiation traffic

Diffie-Hellman Groups

As we discussed earlier in the chapter, Diffie-Hellman groups are used to define the length of the base prime numbers that are used during the key-exchange process. There are three types of Diffie-Hellman groups, as follows:

- **Diffie-Hellman group 1** This is the least secure group and it provides only 768 bits of keying strength.
- **Diffie-Hellman group 2** This group is set to a medium level, at 1024 bits of keying strength.
- **Diffie-Hellman group 3** This group is set to the highest level, at 2048 bits of keying strength.

Diffie-Hellman group 3 is available only on Windows Server 2003 family machines. If you wish to use this algorithm on Windows 2000 machines, you must have either Service

Pack 2 or the High Encryption Pack installed. If you configure one client machine for a Diffie-Hellman group 1 key exchange and another client machine for the Diffie-Hellman group 3 exchange, negotiation will fail.

For the best security, use the highest Diffie-Hellman group 3 key exchange. When using the quick mode, new keys are created from the Diffie-Hellman main mode master key material. If you have the master key or session key PFS enabled, a new master key will be created by performing a Diffie-Hellman exchange. The master key PFS will require a reauthentication of the main mode SA in addition to the Diffie-Hellman exchange. The session key PFS will not require this reauthentication.

Pre-shared Keys

To authenticate L2TP protocol and IPsec connections, you can select to use a pre-shared key. This is the simplest of three choices of authentication methods that you have with IPsec. The other two authentication methods are Kerberos and digital certificates. Before selecting to use a pre-shared key, you should be aware of all the implications of doing so.

A pre-shared key is a string of Unicode characters. You can use the Routing and Remote Access management console to configure connections to support authenticated VPN connections using the pre-shared key. A server that has the Windows Server 2003 operating system installed may also be configured to use a pre-shared key to authenticate connections from other routers via the Routing and Remote Access console.

As we discussed earlier in the chapter, when you create IPsec policies for a computer, you can define the authentication method to be used. In order for two computers to communicate via IPsec, they must have a common authentication method configured. To increase the chances that this will happen, you can configure a machine to use multiple authentication methods. You might want to set up a computer to be able to use a pre-shared key for this reason.

Advantages and Disadvantages of Pre-shared Keys

Pre-shared key authentication does not have the overhead costs that a PKI implementation does. This type of authentication is relatively easy to configure using the Routing and Remote Access console (for L2TP/IPsec connections) or the IP Security Policy Management console (for IPsec secured communications).

Pre-shared keys are stored as plaintext. This means the key can be compromised if a hacker is able to access the file on the computer. Thus, the pre-shared key is the weakest of the three IPsec authentication methods.

Another drawback of pre-shared keys in relation to L2TP/IPsec connections is that a remote access server can use one pre-shared key for all L2TP/IPsec connections that require a pre-shared key for authentication. In this case, you need to issue the same pre-shared key to all L2TP/IPsec VPN clients that connect to the remote access server using a pre-shared key. Unless you are using the Connection Manager profile to distribute the pre-shared key, each user must manually type the pre-shared key. If you change the pre-shared

key on a remote access server, clients with manually configured pre-shared keys will not be unable to connect to the server until the pre-shared key on the client is changed.



EXAM WARNING

Microsoft's recommendation is that pre-shared keys be used for authentication only for testing. It is recommended that you *not* use this authentication method on your production network. Pre-shared keys do not offer good security for sensitive communications, and if you did not need a high-security solution, you would not be implementing IPSec in the first place. Microsoft documentation emphasizes that Windows Server 2003 includes the pre-shared key option only for interoperability with computers that don't support Kerberos and in environments without a PKI.

Considerations when Choosing a Pre-shared Key


Remember that a pre-shared key is just a sequence of characters that is configured on both computers that are parties to an IPSec-secured communication. The pre-shared key can be any non-null string of any combination, up to 256 Unicode characters.

When you choose a pre-shared key, consider that users who use the New Connection Wizard to create a VPN client connection must type the pre-shared key manually. A key that is long and complex enough to provide adequate security might be difficult for the majority of your users to remember or type accurately. If the pre-shared key presented by one party to the communication deviates in any way from the pre-shared key configured on the other, IPSec authentication will fail.

Soft Associations

A *soft association* refers to an SA that was created with a computer that hasn't responded to main mode association attempts since the last time the IPSec service was started. If the IPSec policy is so configured, the communications will be allowed, even though there was no response to the main mode negotiation attempt. It's important to understand that a soft association is *not* protected by IPSec.

The soft association is just a communication that is not secured. This occurs when one of the two communicating computers doesn't support IPSec, and the IPSec policy allows unsecured communications in this situation.



Using RSoP for IPSec Planning

RSoP is a utility provided in Windows Server 2003 for gathering information to help you configure Group Policy in the way that best serves the needs of your network. It functions as a query engine that uses the Common Information Management Object Model (CIMOM) database to store this information.

RSoP is used to sort through the complexities of applying multiple policies and determine the totality of their effects. This is important, because it can be very difficult to predict the outcome when Group Policy is applied at several different levels (site, domain, and OU), and some of those policies conflict.

There are two modes in which RSoP can be used: logging mode and planning mode. Logging mode tells you the effects of the policy settings that are applied to the computer and currently logged-in user. Administrators can use RSoP in planning mode to check existing GPOs and search for all policy settings that can be applied. The results of this search can then be placed in a scenario-based simulation to view how the changes will affect the policies.



EXAM WARNING

The IPSec extension to the RSoP console is a new feature in Windows Server 2003, so you can expect to encounter one or more exam questions dealing with this topic.

Ideal situations for using the RSoP tool include the following:

- Simulating the effect of policy settings on a domain, site, OU, computer, or user
- Determining the effective policies for a newly created account in your Active Directory domain
- Testing policy precedence, such as the user or the computer in different OUs, the user or the computer in different security groups, and when the user or computer is moving

You can also simulate a slow network or create a network loopback situation. RSoP can provide network administrators with details such as security settings, scripts, Group Policy installation, folder redirection, templates, and Internet Explorer maintenance.



EXAM WARNING

If you need to use RSoP on a remote computer, you must be a member of the Domain Admins or Enterprise Admins security group, or be granted the Generate Resultant Set of Policy planning rights.

Using the RSoP Wizard

You can use the RSoP Wizard to create an RSoP query on your Windows Server 2003 server. You begin by adding the RSoP snap-in to an empty MMC console. You can also access RSoP through the Active Directory Users and Computers console and the Active Directory Sites and Services console.

To access RSoP planning through the Active Directory Users and Computers MMC and start the RSoP Wizard, do the following:

1. Select **Start | Programs | Administrative Tools | Active Directory Users and Computers**.
2. Right-click the name of the domain or OU and select **All Tasks**.
3. Choose **Resultant Set of Policy (Planning)**.

To access RSoP planning through the Active Directory Sites and Services MMC and start the RSoP Wizard, do the following:

1. Click **Start | Programs | Administrative Tools | Active Directory Sites and Services**.
2. Expand the **Sites** node in the left pane.
3. Right-click the name of a site and select **All Tasks**.
4. Select **Resultant Set of Policy (Planning)**.

To start the RSoP Wizard from a stand-alone RSoP MMC, right-click **Resultant Set of Policy** in the left pane and select **Generate RSoP Data** (or select it from the **Action** menu). The Wizard will display the query results in the RSoP snap-in. You can save, change, or refresh your RSoP queries. You can create more than one query by adding the RSoP snap-in to your console. The information that RSoP gathers comes from the CIMOM database through Windows Management Instrumentation (WMI).



NOTE

The RSoP Wizard differs depending on which method you use to open RSoP. When you open the RSoP Wizard through the Active Directory Users and Computers or Active Directory Sites and Services console (under Administrative Tools), you can use only planning mode. When you open the Wizard from the RSoP MMC, the first selection you make is whether to use logging or planning mode.

Security and RSoP

Administrators can use RSoP features to determine which particular security policies meet their organization's needs. You can use RSoP security templates to create and assign security

options for one or many computers. You can apply a template to a local computer, and then import that template into the GPO in the Active Directory. After the template has been imported, Group Policy will process the security template and apply the changes to the all members of that GPO. RSoP will also verify the changes that have been made by polling the system and then showing the resultant policy. RSoP can correct a security breach by taking the invalidly applied or overwritten policy setting or the priority policy setting. Group Policy filtering will report the scope of the GPO, based on the security group membership.

Through individual security settings, administrators can define a security policy in Active Directory that contains specific security settings for nearly all security areas. Security settings in a local GPO can establish a security policy on a local computer. When there are conflicts, security settings that are defined in Active Directory always override any security settings that are defined locally.

The RSoP console simplifies the task of determining which IPsec policy is being applied by displaying the following information for each GPO that contains an IPsec policy assignment:

- Name of the IPsec policy
- Name of the GPO that the IPsec policy is assigned to
- IPsec policy precedence (the lower the number, the higher the precedence)
- Name of the site, domain, and OU to which the GPO containing the IPsec policy applies (that is, the scope of management for the GPO)

The settings of the IPsec policy with the highest precedence apply in their entirety; they are not merged with the settings of IPsec policies that are applied at higher levels of the Active Directory hierarchy.

Selecting the RSoP Mode for IPsec-related Queries

As mentioned earlier, RSoP can be run in either of two modes: logging or planning. In the following sections, we will take a closer look at the differences between these two modes and help you determine when to use each for queries related to IPsec.

Logging Mode Queries

You can run an RSoP logging mode query to view all of the IPsec policies that are assigned to an IPsec client. The query results display the precedence of each IPsec policy assignment, so that you can quickly determine which IPsec policies are assigned but are not being applied and which IPsec policy is being applied. The RSoP console also displays detailed settings for the IPsec policy that is being applied, including the following:

- Filter rules
- Filter actions

- Authentication methods
- Tunnel endpoints
- Connection type

When you run a logging mode query, RSoP retrieves policy information from the WMI repository on the target computer, and then displays this information in the RSoP console. In this way, RSoP provides a view of the policy settings that are being applied to a computer at a given time.

Planning Mode Queries

You can run an RSoP planning mode query to view all of the IPsec policies that are assigned to members of a Group Policy container. RSoP will retrieve the names of the target user, computer, and domain controller from the WMI repository on the domain controller. WMI then uses the Group Policy Data Access Service (GPDAS) to create the policy settings that would be applied to the target computer, based on the RSoP query settings that you entered. RSoP reads the policy settings from the WMI repository on the domain controller, and then displays this information in the RSoP console user interface.

You can run an RSoP planning mode query only on a domain controller (when you run a planning mode query, you must explicitly specify the domain controller name). However, you can specify any IPsec client as the target for the query, provided that you have the appropriate permissions to do so.

Summary

In this chapter, we took a close look at Windows Server 2003's implementation of IPsec. We first provided an overview of the goals and purposes of IPsec, and then we discussed the features built into Microsoft's implementation, including the IPsec management console, IPsec integration with Active Directory, supported authentication methods, and backward compatibility with Windows 2000.

You learned some of the terminology and concepts used in discussing IPsec. Specifically, you learned about the two primary protocols used by IPsec: AH and ESP. You learned that AH provides for data authentication and integrity, and ESP also provides those services, and also adds data confidentiality. AH and ESP can be used separately or together.

You learned that an SA is an agreement between two IPsec-enabled computers as to the security settings that will be used for a communication session. The SA is negotiated according to the settings on each computer.

Then you learned about the key-management and key-exchange protocols associated with IPsec, including ISAKMP and IKE, and the Oakley key-determination protocol and the Diffie-Hellman key-generation protocols. You learned about the DES and 3DES encryption algorithms and the MD-5 and SHA hashing algorithms.

We covered the basics of how SAs function, and you learned that IKE uses a bidirectional SA called a main mode SA. However, the SAs used by IPsec itself are unidirectional, and there are two per communication: one for outbound and one for inbound traffic.

We discussed the purposes of security—authentication, integrity, and confidentiality—along with the related concept of nonrepudiation. You learned that authentication deals with verification of identity, integrity ensures that data has not been changed, and confidentiality “scrambles” the data so it cannot be read by unauthorized persons. Nonrepudiation is a way to ensure that the sender of a message will not be able to later deny sending it.

You learned about the two modes in which IPsec can operate: tunnel mode and transport mode. We examined how tunnel mode is used primarily between gateways or between a server and a gateway. You learned that transport mode, on the other hand, provides end-to-end security (from the originating computer to the destination).

We examined the role of the IPsec driver, and you learned that it is used to match packets against the filter list and applies specified filter actions.

You learned how to plan an IPsec deployment, and how to use the IPsec extensions for the new Windows Server 2003 tool, RSoP, to learn what the effects of IPsec policies will be. We took a look at the default policies and how you can use the IPsec management console to enable or modify them. You learned that there are three default policies: Client (Respond Only), Server (Request Security), and Server (Require Security). You also learned about creating custom policies.

We also discussed how to use the command-line tool `netsh` with the `ipsec` context that is new to Windows Server 2003, and you learned that this context operates in one of two modes: static mode, which can be used to perform the same basic functions as the IP

Security Policy Management MMC, and dynamic mode, which is used to display the current state of IPSec and immediately affect the configuration of IPSec policies.

Finally, you learned about troubleshooting problems with IPSec, using handy tools such as the IP Security Monitor console and the Network Monitor.

Exam Objectives Fast Track

Understanding IP Security (IPSec)

- ☑ The IETF designed the IPSec specifications. The IP Security Working Group of the IETF developed IPSec as an industry standard for encrypting TCP/IP traffic within networking environments.
- ☑ Before secure data can be exchanged, a security agreement between the two communicating computers must be established. This security agreement is called an SA.
- ☑ An SA is a combination of three things: security protocols, a negotiated key, and an SPI.
- ☑ IPSec uses cryptography to provide three basic services: authentication, data integrity, and data confidentiality.
- ☑ IPSec in Windows Server 2003 has two different modes: tunnel mode and transport mode.
- ☑ The two primary IPSec protocols are AH and ESP. They can be used separately or together.
- ☑ In addition to the protocols that operate within the IPSec framework, there are a number of operating system components involved in Microsoft's implementation of IPSec. The most important of these are the IPSec Policy Agent service and the IPSec driver.

Deploying IPSec

- ☑ The first step in deploying IPSec is to determine your organizational needs in regard to the security of data traveling over the network.
- ☑ When you begin to consider security levels within your organization, you must take into account the type of data each computer will typically be processing.
- ☑ There are three general types of environments: minimal security, standard security, and high security.

Managing IPSec

- ☑ Windows Server 2003 comes with several handy tools to enable administrators to manage IPSec. These include the IP Security Policy Management MMC and the netsh command-line utility.
- ☑ IPSec policies are used to apply security at various levels within a network.
- ☑ IPSec has three default policies defined: Client (Respond Only), Server (Request Security), and Server (Require Security).
- ☑ To create your own custom policies with the IP Security Policy Management MMC, open the MMC and select the policy you wish to customize.

Addressing IPSec Security Considerations

- ☑ There are two encryption algorithms supported by IPSec for data encryption: DES and 3DES. The 3DES algorithm is the strongest of these.
- ☑ Specific ports and protocols that can be used for firewall filtering include: IP and port 50, IP and port 51, and UDP port 500.
- ☑ Diffie-Hellman groups are used to define the length of the base prime numbers that are used during the key-exchange process.
- ☑ A pre-shared key is a string of Unicode characters. Pre-shared keys are stored as plaintext. This means the key can be compromised if a hacker is able to access the file on the computer. Thus, the pre-shared key is the weakest of the three IPSec authentication methods.

Using RSoP for IPSec Planning

- ☑ RSoP is used to sort through the complexities of multiple policy application and determine the totality of their effects.
- ☑ There are two modes in which RSoP can be used: logging mode and planning mode.
- ☑ RSoP can provide network administrators with details such as security settings, scripts, group policy installation, folder redirection, templates, and Internet Explorer maintenance.
- ☑ Administrators can use RSoP features to determine which particular security policies meet their organization's needs. RSoP security templates can be used to create and assign security options for one or many computers.

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 IPsec AH tunnel mode?

A: The AH tunnel mode is used by IPsec to ensure packet integrity and authentication by encapsulating an IP packet with an Authentication Header (AH) and an IP packet. AH does *not* provide encryption of data.

Q: What is the ESP tunnel mode?

A: The ESP tunnel mode is used by IPsec for data confidentiality. The mode works by encapsulating the packet with an Encapsulating Security Payload (ESP) and IP header as well as an ESP authentication trailer.

Q: On what Microsoft platforms does IPsec work?

A: Native support for IPsec is provided in Windows 2000, Windows XP Professional, and Windows Server 2003 products.

Q: What is the strongest encryption method for key-exchange settings available when implementing IPsec in Windows Server 2003?

A: Triple Data Encryption Standard (3DES), newly supported in Windows Server 2003, uses three 56-bit key exchanges to provide an effective key length of 168 bits.

Q: I am using NAT on my firewall. Can I pass IPsec traffic through my firewall?

A: Yes, if the firewall or NAT device is configured properly to allow for UDP traffic. Unlike Windows 2000, Windows Server 2003 includes support for *NAT traversal*, a method of allowing IPsec and NAT to work together.

Q: How can I manage my IPsec policies in Windows Server 2003?

A: You can use the netsh commands in ipsec context, or you can use the IP Security Policy Management MMC snap-in.

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 IP Security (IPSec)

1. You have decided to deploy IPSec in your organization because you have several departments that are doing sensitive work and communicating across the Internet and other networks with a variety of persons in various organizations. There have been a few incidents where messages were sent instructing lower-level employees to perform certain tasks, purporting to be from their managers. However, investigation revealed that the managers did not send the messages; rather, they were sent by someone else, pretending to be the manager, who was attempting to sabotage the project. This experience has pointed out the need to provide authentication for the data packets that travel across the network so that the receiver of a message can be assured that it is genuine. It is equally important to ensure that the data in these messages doesn't get changed during transmission. Finally, you want to be sure that nobody other than the authorized recipient is able to read the message itself. You want the entire packet to be digitally signed, so that it will have maximum protection. Which of the following IPSec configuration choices will provide this?
 - A. Use AH alone.
 - B. Use ESP alone.
 - C. Use AH and ESP in combination.
 - D. IPSec cannot provide authentication, integrity, and confidentiality simultaneously.
2. You have been hired as a consultant to help deploy IPSec for the network of a medium-size manufacturing firm that is developing a number of new products and must share sensitive data about its products over the network. As part of the planning process, you must determine the best authentication method to use with IPSec. What are the authentication methods that can be used with IPSec? (Select all that apply.)
 - A. Kerberos v5
 - B. Perfect Forward Secrecy (PFS)
 - C. Shared secret
 - D. Diffie-Hellman groups

Deploying IPsec

3. You are the network administrator for a company that has recently migrated some of its servers to Windows Server 2003 from Windows 2000. However, there are still a number of Windows 2000 servers and clients on the network. You want to use the enhanced security available on your network, and you have some interoperability issues you are concerned with pertaining to Windows Server 2003 and your Windows 2000 servers and clients. Which key method should you implement?
 - A. Rivest-Shamir-Adleman (RSA)
 - B. Diffie-Hellman group 1
 - C. Diffie-Hellman group 2
 - D. Diffie-Hellman group 2048

4. You are a network administrator for a medium-sized medical office and you have recently deployed IPsec on the network in response to the physician/owner's concerns about confidentiality of patient information. However, it appears that IPsec might not be working correctly on a particular client computer. You need to view the local routes assigned to this particular client on the network using the IPsec Policy Agent. How does the IPsec Policy Agent function in IPsec? (Select all that apply.)
 - A. Surveys the policy for configuration changes
 - B. Routes the assigned IPsec policy information to the IPsec driver
 - C. Uses the IP Security Policy Agent console to manage IPsec policies
 - D. For nondomain member clients, retrieves local IPsec policy information from the Registry

Managing IPsec

5. You are the network administrator for a large law firm. You have been tasked with the duty of deploying IP security for all network communications in the departments and divisions that handle sensitive data. You have delegated individual departments to your junior administrators. You now need to verify that IPsec has been deployed and configured properly on your Human Resources and Payroll computers. Which tools can be used to perform this function? (Select all that apply.)
 - A. IPsec Security Policy Monitor console
 - B. netsh command
 - C. Certificates snap-in
 - D. Resultant Set of Policy (RSOP)

6. You have deployed IPSec on your company's network and it has been working well, except for one thing. You've tried modifying some of the IPSec policy rules using netsh commands in the ipsec context, but each time you do so, the rules work only until you reboot the server, and then they seem to disappear. You want to make changes to the IPSec policy rules that are permanent and do not change when the server is rebooted. Which netsh command could you use?
- A. netsh ipsec dynamic set config
 - B. netsh ipsec dynamic
 - C. netsh interface ip
 - D. netsh interface ipv6 isatap

Addressing IPSec Security Considerations

7. You are the network administrator for a medium-sized company that provides accounting services to a number of different clients. To avoid having clients' financial information disclosed to the wrong parties, you are planning to implement IPSec on your network. You want your employees to be able to communicate securely both within the company and across the WAN with employees in your branch offices. You have recently hired a junior administrator who has his MCSE in Windows NT and 2000. You give him the task of implementing IPSec in your organization. The first thing he tells you is that because your smaller branch office uses NAT, that site will not be able to use IPSec. What is your response?
- A. You already knew this, and intend to change that site from a NAT connection to a routed connection to accommodate this.
 - B. He is mistaken; IPSec has been able to work with NAT since Windows 2000.
 - C. He is mistaken; IPSec did not work with NAT in Windows 2000 but it does in Windows Server 2003.
 - D. You know IPSec is not compatible with NAT "out of the box," but you can install a third-party program that will make it compatible.

8. You have been hired as network security specialist for a new startup company that has recently installed a new Windows Server 2003 network. The network was originally set up by a group of consultants, and they implemented IPSec for network communications so that communications with their secure servers could be protected. You are reviewing and evaluating the IPSec policies. Although several policies have been created, none of them seem to be effective. What do you conclude the consultants forgot to do after creating the policy?
- A. Authorize the policy in Active Directory
 - B. Assign the policy in the IP Security Policy Management console
 - C. Edit the policy after creating it
 - D. Enable the policy in the IP Security Monitor console
9. You have been tasked with the duty of implementing IPSec on your new Windows Server 2003 network to increase security. You have never worked with IPSec before and you have been reading up on it. You've decided that you want to use PFS, but you are concerned about the resource usage on the domain controller due to reauthentication. Which of the following types of PFS can you implement without putting an undue burden on the authenticating server?
- A. You can use master key PFS.
 - B. You can use session key PFS.
 - C. You can use either or both because PFS doesn't use any resources on the domain controller.
 - D. You can use neither because both types of PFS use considerable resources on the domain controller.
10. You are creating a project to implement IPSec using the IPv6 protocol. Part of your security plan states that you must maintain data confidentiality as part of your IPSec implementation. When developing your plan further, what must you remember about Microsoft's implementation of IPv6 that is included in Windows Server 2003?
- A. IPv6 does not support data encryption.
 - B. IPv6 does not support authentication.
 - C. IPv6 does not support integrity.
 - D. IPv6 does not support IPSec.
11. You have been hired as a consultant to evaluate the IPSec deployment in a small music publishing company. Management is concerned that copyrighted material might be intercepted as it passes over the network and be stolen. You discover that the former network administrator who initially set up IPSec configured it to use the AH

protocol only. You explain to the company manager that one of the things you recommend changing is to configure IPSec to use ESP. Why would you implement ESP in this situation? (Select all that apply.)

- A. ESP ensures data integrity and authentication.
 - B. ESP prevents capture of packets.
 - C. ESP provides confidentiality.
 - D. ESP encrypts the packets.
12. You are on an IT team that is planning the deployment of IPSec throughout a large enterprise network. You have been advised that cost-effectiveness and efficient use of personnel are two priorities, because the company does not want to hire additional IT staff to support the deployment. Of the authentication methods available, which has the lowest administrative overhead and is the most efficient if you wish to support the implementation on 10,000 client machines?
- A. Diffie-Hellman group 2048
 - B. Kerberos v5
 - C. Pre-shared keys
 - D. Digital certificates

Using RSoP for IPSec Planning

13. You have been hired to manage security for a medium-sized network. Your first project is to implement IPSec on the network to protect communications that travel across it. You have just assigned an IPSec policy to a client, and you need to view the precedence of IPSec policy assignments and which policies have been assigned to the client. Which logging mode would you use in RSoP?
- A. IPSec mode
 - B. RSoP mode
 - C. Logging mode
 - D. Planning mode

14. You have IPSec configured and running on your network. You want to capture some IPSec packets to ensure that the data inside cannot be viewed. You want to capture packets being sent from a remote client to a remote server, using a server in the server room. Which of the following tools will you need to use in order to capture these packets?
- A. Network Monitor in Windows Server 2003
 - B. netsh commands in the ipsec context
 - C. The IP Security Monitor console
 - D. Systems Management Server (SMS)
15. You want to use the RSoP tool in logging mode to build some reports on the existing policy settings of one of your client computers. You have used RSoP before in planning mode, but never in logging mode. You open the RSoP Wizard from the Active Directory Users and Computers console, as you've done before, but you notice that there is no mechanism for selecting the mode, and only planning mode seems to be available. What is the problem?
- A. The RSoP Wizard runs only in planning mode.
 - B. You should open the RSoP Wizard from Active Directory Sites and Services instead.
 - C. You should open the RSoP Wizard from the RSoP MMC instead.
 - D. You can select logging mode when you open the RSoP in Active Directory Users and Computers. You must have overlooked the option.

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

MCSE 70-293

Planning, Implementing, and Maintaining a Security Framework

Exam Objectives in this Chapter:

- 5 Planning and Maintaining Network Security
 - 5.4 Plan secure network administration methods.
- 6 Planning, Implementing, and Maintaining Security Infrastructure.
 - 6.3 Plan a framework for planning and implementing security.
 - 5.5 Plan security for wireless networks.
 - 6.3.1 Plan for security monitoring.
 - 6.3.2 Plan a change and configuration management framework for security.
 - 6.4 Plan a security update infrastructure. Tools might include Microsoft Baseline Security Analyzer and Microsoft Software Update Services.
- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

Security is one of the most important issues facing network administrators today. Windows Server 2003 contains many features and technologies that can be used to create a more secure networking environment. In this chapter, we'll look at several aspects of creating an effective security framework for an organization's network. First, we'll look at how to plan and implement Active Directory (AD) security. This includes such measures as physically securing domain controllers, securing the schema, managing cross-forest security relationships, account security, and implementing AD access controls.

Next, we'll discuss the issues and procedures involved in planning and implementing wireless security. We'll provide an overview of the terminology and concepts relating to 802.11 wireless technologies, and you'll learn about authenticators and supplicants, as well as how wireless networking works "under the hood." We'll discuss authentication methods for wireless networks, including authentication subtypes such as open system and shared key. You'll learn about the protocols generally used for wireless authentication, including the Extensible Authentication Protocol (EAP), EAP-Transport Layer Security (EAP-TLS), EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (EAP-MS-CHAPv2), and the Protected Extensible Authentication Protocol (PEAP). We'll also talk about using Internet Authentication Service (IAS) with wireless. We'll address wireless security issues such as common insecure default settings (administrative password, service set identifier, and Wired Equivalent Privacy settings) and the weaknesses of Wired Equivalent Privacy (WEP) protocol encryption, as well as how WEP can be made more secure.

We'll then move onto discuss security monitoring, and we'll address object-based access control and security policies, including password policies, Kerberos policies, account lockout policies, user rights, and security templates. We'll also talk about security auditing, and you'll learn how to set the auditing policy, modify the security log settings, and audit objects such as files or folders. In the next section, you'll learn about planning a change and configuration management framework. We'll walk you through the steps of using the Security Configuration Manager tool, as well as command-line tools included with Windows Server 2003. We'll also discuss security analysis and configuration best practices.

Finally, we'll take you through the process of planning a security update infrastructure. You'll understand the importance of regular security updates, and you'll learn how to use the Microsoft Baseline Security Analyzer (MBSA) and Microsoft Software Update Services (SUS) tools to ensure that your Windows Server 2003's security features are always current.

Planning and Implementing Active Directory Security

Windows Server 2003 supports statically assigned authorization to resources. This information is used to determine whether access to a resource is granted or denied. This is also referred to as *static access control*. Administrators can control access to AD objects by assigning

EXAM
70-293

OBJECTIVE

5

5.4

6

6.3

them *security descriptors*. The security descriptor consists of information regarding the object's ownership, access control lists (ACLs), and auditing.

Understanding Static versus Dynamic Access Control

As you might guess, in addition to static access control, there is another type of access control called *dynamic access control* (not currently supported by Windows Server 2003). With static access control, the information used to grant or deny access is preconfigured. At logon, a user is assigned an access control token according to the user's account information (such as group memberships). If that information is changed (for example, the user is added to a new group or removed from a group), the change does not become effective until the user logs off and then logs on to receive a new access token.

With dynamic access control, access information can be changed dynamically, and the system determines whether to grant access at the time the request is made, based on information at that time, instead of a token issued at logon.

The ACL holds the static access control information. There are two parts to the ACL in the security descriptor:

- **The discretionary access control list (DACL)** The DACL contains the information about which users and groups are allowed (or denied) permission to access the object, and the level of access granted. The security descriptor can even be configured to control access to a particular *attribute* of an object.
- **The system access control list (SACL)** The SACL specifies the events that should be audited (if auditing is enabled).

DACLs and SACLs are associated with each type of AD object in Windows Server 2003.

There are three types of standard permissions supported with AD: Full Control, Write, and Read. Each AD object has these three standard permissions available for use. In addition to these three standard permissions, there are a number of additional permissions that can be used to control access more granularly. These depend on the object type. For example, additional permissions that can be assigned for user objects include Create All Child Objects, Delete All Child Objects, Allowed to Authenticate, Change Password, Receive As, Reset Password, Send As, Read Account Restrictions, Write Account Restrictions, Read General Information, Write General Information, Read Group Membership, Write Group Membership, Read Logon Information, Write Logon Information, Read Personal Information, Write Personal Information, Read Phone and Mail Options, Write Phone and Mail Options, Read Public Information, Write Public Information, Read Remote Access Information, Write Remote Access Information, Read Web Information, Write Web Information, and Special Permissions (configured through the Advanced settings).

So how are all these additional permissions used? For example, you could use the Read Group Membership permission to specify which users and groups are allowed to read the group membership information for a particular user account, as follows:

1. Select **Start | All Programs | Administrative Tools | Active Directory Users and Computers**.
2. Click **View | Advanced Features** (this will place a check mark by it in the menu).
3. In the left pane of the Active Directory Users and Computers console, expand the domain node and click the **Users** container.
4. In the right pane of the console, right-click the user account for which you want to set access controls and select **Properties**.
5. Click the **Security** tab. View the list of group and usernames in the top pane. When you select one, you will see the permissions assigned to it in the bottom pane. To add a group or user to this list, click the **Add** button.
6. Click **Allow** or **Deny** for each permission to granularly configure the permissions on the user account object for each user or group.

Another important aspect of access control is *ownership*. When a user creates an object, that user is designated as the owner of the object. An owner/creator has full access to the objects that user owns. Ownership can be delegated by the current owner to someone else.

To view the DACL, SACL, and ownership information for an object (such as a user account), click the **Advanced** button on the **Security** tab and do the following:

- Click the **Permissions** tab to view the DACL.
- Click the **Auditing** tab to view the SACL.
- Click the **Owner** tab to view the ownership information.

NOTE



To get a detailed view of a particular access control entry (ACE), select the appropriate ACE on the **Permissions** or **Auditing** tab and click the **Edit** button.

When implementing Windows Server 2003's AD, you need to provide security for your entire organization. AD allows for that by permitting you to put into action *user authorization* and built-in *logon authentication*. User authorization is used to provide specified clients access to objects such as folders. Built-in logon authentication is used for system permissions on the network.

You can also use trust relationships and Group Policy to provide security solutions for your AD network. Table 11.1 shows some AD security scenarios, along with solutions or tools you can use to plan and implement security.

Table 11.1 Scenarios and Solutions for a Stronger and More Secure Active Directory

Scenario	Solutions
You need to support and manage two forests in your AD security framework, and authentication across forests needs to be simplified.	Use a forest trust, which is a trust between two Windows Server 2003 forests. This trust will create trust relationships between every domain in the two forests. They can be created only on the forest root domains in each forest. These forest trusts are transitive. They can be one-way or two-way trusts. Unlike a parent-child trust, which is automatically established (implicit trust), administrators must manually establish a forest trust (explicit trust).
You need to enforce strong password policies because you have seen the word <i>password</i> used as an actual client password.	Use Group Policy to enforce strong password policies. Strong passwords are at minimum eight characters long and contain at least three or four characteristics, such as uppercase characters, lowercase characters, numeric digits, and symbols found on the keyboard (for example, !, @, \$,#). They do not contain any part of the client's user account name, words in dictionaries, or other easily guessed information.
You need to check event logs as part of your daily routine.	Use Group Policy to enable the Audit Policy function. Checking event logs as a daily routine can allow you to quickly recognize a security risk.
You need to make sure no one is trying to guess at user's passwords in an attempt to compromise security.	Use Group Policy to use the Account Lockout Policy function on user accounts. This will limit the possibility of an attacker compromising the domain through frequent logon attempts.
You need to minimize the possibility of an attacker trying to crack a user's password.	Use Group Policy to enforce minimum and maximum password age policies on user accounts to decrease the possibility of an attacker compromising your domain. As a rule, it is best practice to have passwords expire every 30 to 90 days. The default password age is 42 days.
You need to authenticate and verify the validity of each client.	Use public key cryptography to authenticate and verify each client's identity.
You need to administer servers in your domain, but you do not want to be logged on as Administrator for long periods of time.	Use the Run as command to perform administrative tasks on the servers without needing to log on with your administrative credentials.

Continued

Table 11.1 Scenarios and Solutions for a Stronger and More Secure Active Directory

Scenario	Solutions
You need to thwart attacks from people who might try to grant elevated user rights to another user account.	Use security identifier (SID) filtering to stop the elevation of privilege attacks.
You have trade secrets and other confidential information on your network and you need to secure them.	Implement smart card authentication to provide two-factor authentication, encrypt data on the disk with Encrypted File System (EFS), and protect data traveling across the network with IPSec.
You need to secure account passwords on domain controllers, member servers, and local computers.	Implement the System Key Utility (<i>syskey</i>), which will provide strong encryption techniques to secure account password information.

The following are some basic security guidelines:

- Avoid granting Full Control permissions over objects or Organizational Units (OUs) except when absolutely necessary, because this allows someone to take ownership of an object and also modify permissions on the object. The user will also have full control over all objects in that container unless inheritance of permissions is blocked.
- Avoid changing the default permissions on AD objects, because this can cause unexpected results by creating access problems or reducing security.
- Reduce the number of access control entries (ACEs) that apply to child objects. When using the Apply Onto option to control inheritance, not only do the specified objects inherit that access control, but also *all* child objects receive a copy of that ACE. Too many copies of this ACE on the network could significantly reduce network performance.
- In the Windows Server 2003 family, ACLs feature single-instancing. Single-instancing works by storing only one instance of the ACL, even if multiple objects have identical ACLs.
- If possible, assign permissions to groups rather than users. This makes management easier.
- Generally, allow Read All Properties or Write All Properties permissions, rather than setting controls on individual properties, unless there are compelling reasons to do so.
- Allow Read or Write access to property sets, rather than to individual properties.



NOTE

Property sets are also called *attribute sets*. These are defined sets of attributes that represent the entire set in an ACL. Microsoft defines 10 attribute sets; and you can define custom attribute sets, but each attribute can be a member of only one set.

Understanding Permission Types

It's important to differentiate between the different types of permissions that can be assigned to user and group accounts in Windows Server 2003 networks. There are three types of permissions:

- AD object permissions
- NTFS file permissions
- Shared folder (share) permissions

The following sections describe the AD, NTFS, and share permission types, as well as the permissions available for each.

Active Directory Permissions

AD permissions are set on any AD object, as follows:

1. Select **Start | Administrative Tools | Active Directory Users and Computers**.
2. Select **View | Advanced Features**.
3. Right-click the object you wish to set permission on and select the **Properties** option.
4. Select the **Security** tab and choose **Advanced**. You will see all of the available permissions for this object.
5. Click **Add** to add permissions and type the name of the user, computer, or group you wish to add. Then click **OK**.
6. In the **Permission Entry for Objectname** option, select the **Allow** or **Deny** options from the **Object** and **Properties** tab. The Objectname would be whatever you choose to set the permission on in step 3 above.

Object permissions in AD have many rights, such as the following:

- **Extended Rights** Used for special operations within AD that are not related to either Read or Write access, such as Change Password, which allows the ability to change a password if the original password is known.

- **Validated Writes** Includes value checking, which makes certain that the changed value matches specific requirements. An example of this is Validated Write Add/Remove Self As a Member. This applies to a group and allows members to remove or add themselves to a specific group for membership.
- **Property Set** Allows a group of properties that have a specific set of rights, rather than individual rights. An example is Domain Password, which is managed through the user interface using the Domain Security Policy Group Policy Object (GPO).

Most of the time, when setting permissions on AD objects (and for assigning permissions in general) you should use groups rather than individual user accounts to control access. If one set of users needs Read permissions, and another set of users needs Change permissions, then create one group for each set of users and assign the permissions to the group. If multiple global groups need the same access, create a local group containing the global groups and assign permissions to the local group.



EXAM WARNING

Any object that has an explicit Allow permission entry cannot be overridden by the inherited Deny permission. Explicit permissions will always take precedence over inherited permissions. However, the explicit Deny permission always takes precedence over all other permissions.

You should try to minimize the total number of individual permissions that are published to child objects in AD. This can become a real performance issue if the total size of all the AD permissions approaches the limits of the disk storage space or the memory and processing capacity of your domain controller.

NTFS Permissions

NTFS permissions are applied at the file or folder level, and they are effective both when accessed across the network and when accessed at the local machine. NTFS permissions cannot be applied to files and folders located on FAT or FAT32 volumes. NTFS permissions can be applied in conjunction with share permissions, which give you more control of the shared resource. When the cumulative NTFS permission conflicts with the cumulative share permission, the more restrictive permission will always override the less restrictive permission. NTFS permissions allow you to set permissions for a group or user by using the Allow or Deny option.

You can set NTFS permissions as follows:

1. Open **Windows Explorer** and right-click the file or folder on which you wish to set permissions.
2. Select the **Security** tab.

3. To add a user or group, click the **Add** button. By default, the Administrators group has Full Control and the Everyone group has Read, Read & Execute, and List Folder Contents permissions.

NTFS permissions include Fill Control, Modify, Read and Execute, List Folder Contents, Read, Write, and Special Permissions. To use the Special Permissions option, you must select the **Advanced** option, then edit the **Permission Entry** for which you wish to grant special permissions. The Special Permissions option lists Fill Control, Traverse Folder/Execute File, List Folder/Read Data, Read Attributes, Read Extended Attributes, Create Files, Write Data, Create Folders/Append Data, Write Attributes, Write Extended Attributes, Delete Subfolders and Files, Delete, Read Permissions, Change Ownership, and Take Ownership.



NOTE

In the Windows Server 2003 family, the Anonymous group is not a member of the Everyone group. This includes Windows XP machines. If you have older applications that require anonymous access, you will probably need to go in and tweak the group permissions. If you need to grant access to the Anonymous logon group, you will need to explicitly add the Anonymous Logon security group and its permissions.

Share Permissions

Share permissions apply only to folders that are shared, and they only restrict access across the network. Share permissions on a folder are irrelevant when the folder is being accessed on the local machine.

To set share permissions, follow these steps:

1. Open **Windows Explorer**, right-click the folder on which you wish to set permissions, and select **Sharing and Security** in the context menu.
2. Select the **Sharing** tab.
3. To set permissions for a user or group, click the **Permissions** button.
4. Click the **Add** button to add a user or group. By default, the Everyone group has Read permission.

When setting share permissions, you can assign Read, Change, or Full Control permissions to a user or group.

**NOTE**

There are times when assigning individual permissions is appropriate. For example, if a help desk operator needs Write access to two properties of a user object, it makes sense to use one ACL for each of them, rather than trying use only one ACE by granting Write access to the entire user object.

Physically Securing Domain Controllers

A good security policy and framework starts with the basics: physical security. Implementing firewalls, access controls, permissions, and other software security measures on the network, but keeping the door to your server room unlocked could give malicious users access to everything they need to take control of your systems and access your data. Best security practice is to keep all domain controllers locked in a secure room that has limited public access.

You should also strictly limit which users you place in the following groups:

- Enterprise Admins
- Domain Admins
- Server Operators
- Account Operators
- Print Operators
- Backup Operators

Membership in these groups confers a lot of power (and the opportunity to misuse that power) and so should be limited to trusted personnel within your organization.

Securing the Schema

ACLs are used to protect schema objects from unauthorized use in AD. Members of the Schema Admins group are the only members permitted to have write access to the schema. The only default member of the Schema Admins group is the Administrator account in the root domain of the forest.

You should restrict membership in the Schema Admins group, because extending the schema improperly can have serious consequences to your network. For example, an improper change to the schema can cause existing objects in the directory to become invalid. If you disable a particular attribute in an object class and there are existing objects in that class that contain that attribute, these objects will become invalid because they contain an attribute that is not allowed in the class definition.

Managing Cross-domain and Cross-forest Security Relationships

Security gets more complicated when users in one domain need to access resources in another domain. The domains might be within the same forest or in different forests. When the domains are in the same forest, it is called a *cross-domain relationship*. When they are in different forests, this is called a *cross-forest relationship*.

Cross-domain Relationships

You might sometimes need to allow additional domains inside the same forest to gain security access to a trusted domain. This is done by using the Kerberos or NTLM protocol for user authentication. Kerberos works by allowing a domain controller called the *ticket granting authority* to issue a ticket to the client machine that made the request. The trusting domain then is presented with ticket request for authentication. After the user has been authenticated in either Windows 2000 or Windows Server 2003, AD will allow users to have access from one domain to another only after they are authenticated from their original domain. This is because two domains that belong in the same forest are joined by an implicit trust relationship.

All of the root domains for all domain trees within a forest have a two-way transitive trust with one another. This means that any domain in a forest can access the resources in any other domain in the forest (given the proper permissions), because the trust path flows up the tree to the root domain, across to the root of the other domain, and down the tree to the other domain. However, to shorten this path and make for faster authentication, you can create *shortcut trusts* directly between the two domains.

Because all domains within a forest trust one another, access control between domains is accomplished in the same way as within a domain: by using security groups and setting permissions and user rights. When using groups to assign permissions and rights, keep in mind a few concepts that relate to security groups:

- Group nesting
- Group scope
- Functionality level of the domain

We will discuss each of these in a more detail.

Group Nesting

Nesting refers to the ability to place groups inside one another; that is, a group can be a member of another group. You cannot just nest any type of group anytime, however. Nesting is allowed based on group scope and the functionality level of the domain.

Group Scope

The *scope* of a group represents the area that a group covers within a domain tree or a forest. There are three possible group scopes:

- **Universal (only in Windows 2000 native and Windows Server 2003 domains)** Can contain accounts, global groups, and universal groups from any domain in the forest.
- **Global** Can contain accounts from the same domain. If the functional level is Windows 2000 native or Windows Server 2003, this scope can also contain other global groups from the same domain.
- **Domain local** Can contain accounts and global groups from any domain. If the functional level is Windows 2000 native or Windows Server 2003, this scope can also contain universal groups from any domain and domain local groups from the same domain.

Functionality Level of the Domain

Both domains and forests can be set to different functionality levels, depending on the types of domain controllers that are present. Domain functionality affects how groups can be nested and what types of members groups of each scope can contain. There are four functionality levels available for domains:

- **Windows 2000 mixed** Can contain Windows NT 4, Windows 2000, and Windows Server 2003 domain controllers. This is the default functionality level when you create a new domain.
- **Windows 2000 native** Can contain Windows 2000 and Windows Server 2003 domain controllers.
- **Windows Server 2003 Interim** Can contain Windows NT 4 and Windows Server 2003 domain controllers.
- **Windows Server 2003** Can contain only Windows Server 2003 domain controllers.



NOTE

You can raise the functionality level of a domain using the Active Directory Domains and Trusts tool, but you cannot lower it.

Cross-forest Relationships

By default, domains in different forests do not have a trust relationship with one another. However, Windows Server 2003 gives you two ways to provide for such relationships: you can join two forests together in a *forest trust*, or you can create an *external trust* between two domains in two different forests.

External trusts were also available in Windows 2000. Forest trusts are new to Windows Server 2003. In Windows Server 2003 networks, security identifier (SID) filtering is enabled by default when you create a new external or forest trust. This is a feature that protects against attackers who seek to elevate their user rights or privileges. SID filtering ensures that only the SIDs of security principals in the trusted domain are contained in any authentication requests coming from that trusted domain. This prevents the misuse of the SIDHistory attribute to grant higher privileges than an account should have.

External Trusts

An external trust creates a one-way or two-way *nontransitive* trust between domains in different forests. These are used if you do not want the transitivity provided by forest trusts; that is, you want Domain B in Forest 1 to have a trust relationship with Domain E in Forest 2, but you don't want the trust relationship to extend to any other domains in other forests.



NOTE

Forest trusts are also used to give users access to resources that are located in a Windows NT 4 domain, since such domains are not members of forests.

When you create an external trust, AD handles the cross-forest relationship by creating a *foreign security principal* object in the trusting domain, also called the *internal* domain (the one where the resources are located) to represent each of the security principals (the users who want to access those resources) from the trusted, or *external*, domain. The foreign security principals can be put into Domain Local groups in the trusting domain, because Domain Local groups are allowed to contain members from domains that are in different forests.

You can create an external trust either by using the Active Directory Domains and Trusts tool or by using the **netdom** trust command. To use Active Directory Domains and Trusts, do the following:

1. Log on as a Domain Admin, Enterprise Admin, or user who has been delegated the proper authority, or use the **Run as** command to enter your credentials.
2. Select **Start | All Programs | Administrative Tools | Active Directory Domains and Trusts**.
3. In the left pane of the console, right-click the domain node with which you want to create the trust relationship and select **Properties**.

4. Click the **Trusts** tab.
5. Select **New Trust** and click **Next**.
6. Enter the Domain Name System (DNS) or NetBIOS name of the domain on the **Trust Name** page. Click **Next**.
7. Select **External trust** on the **Trust Type** page. Click **Next**.
8. The next page asks you to select the direction of the trust (**one way incoming**, **one way outgoing**, or **two-way**). Select the appropriate trust type and follow the Wizard's directions (which depends on the trust type you selected).

You can create both sides of a two-way trust *if* you have the proper credentials for both domains. You can either allow users from the other domain to access all resources in the domain or you can allow access to only selected resources.

The syntax for using the **netdom trust** command to create a two-way trust is as follows:

```
netdom trust <trustingdomainname> /d:<trusteddomainname> /add /two-way
```

Forest Trusts

Forest trusts reduce the number of external trusts that need to be created. Forest trusts are created between the root domains of two forests. A forest trust can be either a one-way or two-way transitive trust. If you create a two-way trust relationship, this will effectively provide a trust relationship between every pair of domains within the two forests.

The forest functional level in both forests must be set to Windows Server 2003 before you can create a forest trust, and DNS must be properly configured. Both Kerberos v5 and NTLM are supported for authentication between the forests, and user principal names (UPNs) can be authenticated across the forests.



EXAM WARNING

A forest trust affects only the two forests between which the trust is explicitly created. It does not extend to other forests. So, although transitivity exists *within* the trust, there is no transitivity in terms of a third trust. In other words, if you create a trust between Forest 1 and Forest 2, and you create another trust between Forest 2 and Forest 3, this does *not* result in any kind of trust relationship between Forest 1 and Forest 3.

A forest trust is created using the Active Directory Domains and Trusts tool. To create a trust between two Windows Server 2003 forests, perform these steps:

1. Select **Start | All Programs | Administrative Tools | Active Directory Domains and Trusts**.

2. In the left console pane, right-click the domain that is the forest root and select **Properties**.
3. Click the **Trust** tab and select **New**. Click **Next**.
4. Enter the DNS or NetBIOS name of the forest with which you want to create a trust on the **Trust Name** page. Click **Next**.
5. Select **Forest trust** on the **Trust Type** page. Click **Next**.
6. Select the direction (**two-way**, **one way incoming**, or **one way outgoing**) on the **Direction of Trust** page.
7. Follow the Wizard's directions (which vary depending on the direction of trust selected).



NOTE

There is a special group called the Incoming Forest Trust Builders group. Members of this group are allowed to create one-way incoming trusts to the forest root domain. By default, there are no members in this group, but you can use it to delegate this authority.

You can also synchronize certain types of data across forests. This includes Global Address Lists (GALs) used by Microsoft Exchange Server, public folders, and directory objects.

Account Security

Windows Server 2003 uses *accounts* to represent security principals (entities that access resources on the computer or on the network) of the following types:

- Users
- Groups
- Computers



NOTE

Services can also have accounts. These are special user accounts that are used as dedicated service accounts so applications can access resources.

These accounts automatically have a SID assigned to them, that will be used by the account to access resources within the domain. As mentioned earlier, if a security principal from an outside domain requests resources from a local AD domain, AD will create a foreign security principal for that request.

The security information that is associated with an account is used in the authentication process, which is when a user or computer's identity is verified. After authentication, it is used for authorization, which is the determination of which resources the user or computer can access and the level of access that is allowed. Additionally, the account is used for auditing any activities that are performed on the computer.

AD contains three built-in accounts that are created by default when the domain is created:

- **Guest account** This account does not require a password (although you can assign one) and is for users who do not have an account in the domain. It is a member of the built-in Guest group. This account is disabled by default, and for security reasons, it should remain disabled.
- **Administrator account** This account has complete and total control of the domain. It belongs to the Administrator, Domain Admins, Enterprise Admins, Schema Admins, and Group Policy Creator Owners groups. For obvious reasons, this account should have a strong password and be used only for tasks that require this level of privileges.
- **HelpAssistant account** This ad hoc account is created when you request a remote assistance session and is used to establish a remote assistance session. If no remote assistance requests are pending, the account is deleted.



NOTE

In earlier versions of Windows, the Administrator account could not be disabled. With Windows Server 2003, you can disable the built in Administrator account to prevent it from being compromised, and give administrative privileges to other individual accounts. If you choose not to disable it, you should at least consider renaming it. If you don't, attackers already have half of what they need to know (username and password) to gain control of the computer or network. If you disable the Administrator account, it can still be used when the computer is started in Safe Mode, to perform troubleshooting and repair activities.

User Account Security

Since all of the built-in accounts have the same permissions by default, it is a good idea to either rename or disable them. For instance, all Windows Server 2003 servers have an Administrator account by default. This is a security risk because it is no secret that built-in accounts are part of the operating system. You can rename the account to something that sounds innocuous. For example, hackers will not immediately know that the jrjones account is actually the built-in Administrator account. If you rename the Administrator account, it will still retain its properties, because it keeps the same SID. The SID will enable

the account to retain its password, group membership, profile, assigned permissions, user rights, and account information.

Another way to increase security is to use the Account Lockout Policy. This policy will look at how many times a user has tried to log on to the domain and will deny access if a specified threshold is reached, locking out the account so that no more attempts can be made for a specified period. The threshold can be set at any number from 0 to 999. For example, if you set the account lockout threshold to 3, after the account credentials have been entered incorrectly three times the user will receive a message that the account has been locked out. By default, the lockout duration is 30 minutes, but you can set it to any number of minutes from 0 to 99,999. If you set the lockout duration to 0, the account will remain locked until an administrator explicitly unlocks it.

You can also set a time after which the lockout counter is to be reset. If a user tries to log on and is unsuccessful, that counts as one attempt. If the user then tries again before the reset period, that will count as a second attempt. However, if the reset period has passed and the user tries to log on again, it will count as the first attempt.

All Account Lockout Policy options are set in Group Policy. The Local Security Policy is used to set lockout parameters for logging on to the computer, and the Domain Policy is used to set lockout parameters for logging on to the domain. To set the lockout policies, open the appropriate (local or domain) GPO and follow these steps:

1. In the left console pane of the GPO Editor, expand the **Computer Configuration** node, expand **Windows Settings**, and expand **Security Settings**.
2. Click **Account Lockout Policy**.
3. In the right console pane, you will see three policy entries: **Account lockout duration**, **Account lockout threshold**, and **Reset account lockout counter after**.
4. Click the policy you want to configure, check the **Define this policy setting** check box, and then enter the desired value (number of attempts before lockout, duration in minutes of lockout, or minutes to pass before resetting the lockout counter).
5. Click **OK**.

The Account Lockout Policy options will be discussed in more detail in the “Security Policies” section later in this chapter.

Computer Accounts

All machines that join a domain have a computer account created in AD, and each computer account is unique. This includes Windows NT, Windows 2000, Windows XP, and Windows Server 2003 machines. These are the only machines that can have domain accounts. Clients that are running any of the Windows 9x operating systems (including Windows ME) are not assigned computer accounts. A user who has a user account in the

domain can use a Windows 9x machine to log on, but the machine itself does not belong to the domain and cannot be centrally managed as computers running more secure operating systems can.

You can create a computer account in advance for a computer that is going to join the domain using Active Directory Users and Computers, or the account can be created during the process of joining the domain as long as the user who is joining the computer to the domain has domain administrative credentials or has been granted the right to join computers to the domain. Table 11.2 outlines some common situations that you might encounter and the options you can set in the user account properties to provide a solution for each.

Table 11.2 User Account Scenarios

Scenario	Option to Use	What It Does
A user gave his password to a consultant because the user did not wish to contact the tech support help desk to request a user account and password for the consultant.	Force the user to change his password the next time he logs on. Set the User must change password at next logon option in the user's account properties.	This will force the user to change his password the next time the user logs on to the network. This will ensure that only this user knows his new password.
You have a temporary worker who will be with the corporation for only eight weeks. You want to give her access to network resources, but do not want the user to change her password.	Set the User cannot change password option in the user's account properties.	This will allow you to have total control over this account by preventing the temporary worker from changing her password.
You need to create new service accounts for your backup software. You don't want to deal with changing its password every 45 days in accordance with your current password policies.	Set the Password never expires option in the account's properties.	This password will never expire. By using this option for your service accounts, you will not need to worry about passwords expiring (which could cause software that uses these service accounts to stop functioning). This setting overrides the password expiration policies set in Group Policy.
You have a client that needs to log on to the network from an Apple Macintosh computer.	Set the Store passwords using reversible encryption option in the user's account properties.	This option will let Apple computer clients log on to a Windows network. This option should be used only for this purpose.

Continued

Table 11.2 User Account Scenarios

Scenario	Option to Use	What It Does
You have received a call stating that a person has left the company and you need to stop the user from gaining access to the network.	Set the Account is disabled option in the user's account properties.	This will stop the user from accessing any network accounts.
You have a remote client who has a laptop with a smart card enabled for logging on to the network. This user will be in town for a conference and needs to log on to the domain.	Set the Smart card is required for interactive logon option in the user's account properties.	The user can use his or her smart card and smart card reader with the valid personal identification number (PIN) for logon access. The password will be set to a random and complex value, and the password will never expire option is set automatically.
You need to set a service account to behave as if it were a user account on the network.	On the Delegation tab of the user's account properties, set the Account is trusted for delegation option.	This will allow a service running under this account to perform tasks as if it were a user account.
You need to make sure no one assigns a temporary or guest account to a service.	Set the Account is sensitive and cannot be delegated option in the account properties.	This will prevent the delegation of the specified account to a service.
You have clients who need to use 3DES encryption for logging on to the domain.	Set the Use DES encryption types for this account option in the user's account properties.	This option will allow you to support the Data Encryption Standard (DES). This encryption level supports the following standards: IPsec 56-bit DES, IPsec Triple DES (3DES), MPPE Standard (40-bit), MPPE Standard (56-bit), MPPE Strong (128-bit), and IPsec DES (40-bit).
You need to allow for implementation of Kerberos other than the Windows 2000 and Windows Server 2003 implementations.	Set the Do not require Kerberos pre-authentication option in the user's account properties.	This will provide support for alternate implementations of the Kerberos protocol. Domain controllers running Windows 2000 or Windows Server 2003 can use other mechanisms to synchronize time. Note that pre-authentication provides additional security.

You can configure all of the settings noted in Table 11.2 by right-clicking the user account in **Active Directory Users and Computers** and selecting **Properties**. All of these settings are on the **Account** tab, except when otherwise specified. Note that the

Delegation tab might not appear in the user account properties. If not, you need to do one or both of the following:

- Raise the domain functional level to Windows Server 2003.
- Register a service principal name (SPN) for the user account.

To register an SPN for a user account, follow these steps:

1. If you haven't done so previously, install the Support Tools from the Windows Server 2003 installation CD (navigate to the Support Tools folder on the CD and click the .msi file).
2. Select **Start | Run**, type **cmd**, and click **OK** to open a command prompt window.
3. At the command prompt, type **setspn -A http/<SPNNAME> accountname**. For example, if the account name is joejones, you should type **setspn -A http/joe joejones** to register the SPN "joe" for this account.
4. You will receive the output message that the system is registering the ServicePrincipalName.
5. Open the **Properties** sheet for the account, and you will see a **Delegation** tab.



TEST DAY TIP

Constrained delegation is new to Windows Server 2003. Administrators can specify which service an account can be delegated to, and the trust delegated for a service can be limited to a select group of services that are defined by the administrator. For more detailed information about this new feature, see the article titled "How New Delegation of Authentication Options Improve Security in Windows Server 2003," by Debra Littlejohn Shinder (http://www.windowsecurity.com/Deb_Shinder/).

User Authentication

When a user's identity is established by the authentication process, the *Local Security Authority* (LSA) will process the use Kerberos v5 or NTLM authentication requests. Then LSA will generate an access token that contains the username and SID or SIDS (depending on whether the user belonged to more than one group) for the groups to which the user belongs. The SID will be permanently associated with this user account. If the account is renamed, it will still have the same SID. If the user is added to or removed from a group after the access token has been issued at logon, however, the user must log off and log back on so that the access token can be updated before the new group membership information will take effect. At that point, the user will have the permissions and rights associated with the new group membership status.

EXAM 70-293
OBJECTIVE
5
5.4
5.5
6
6.3

Planning and Implementing Wireless Security

Wireless networking is a new technology that has become more widely used since the Institute of Electrical and Electronics Engineers (IEEE) 802.11 specifications have been implemented. These include the following specifications:

- 802.11b, the original wireless networking standard
- 802.11a, the newer wireless standard that provides faster transfer speeds but a shorter distance range
- 802.11g, the newest wireless standard that provides higher speeds like 802.11a but greater distances like 802.11b
- 802.1x, the wireless security standard

The IEEE 802.11 standard provides for all the necessary definitions and constructs for wireless networks. Everything from the physical transmission specifications to the authentication negotiation is defined by the standard. Wireless traffic, like its wired counterpart, consists of frames transmitted from one station to another. The primary feature that sets wireless networks apart from wired networks is that at least one end of the communication pair is either a wireless client or a wireless access point.

Connecting to a wireless network is often transparent to users, and from their perspective is apparently no different from connecting to a copper-based or fiber-based Ethernet network, with the exception that no wires are involved. With Windows XP, Microsoft introduced automatic configuration and seamless roaming from one wireless network to another through its Wireless Zero Configuration service.

The obvious and primary difference between wired and wireless networks is that wireless networks use a special type of electric current, commonly known as radio frequency (RF), which is created by applying alternating current (AC) to an antenna to produce an electromagnetic (EM) field. The resulting RF field is used by devices for broadcast and reception. In the case of wireless networks, the medium for communications is the EM field, the region of space that is influenced by the EM radiation (unlike audio waves, radio waves do not require a medium such as air or water to propagate).

The 802.11 standard provides for two modes for wireless clients to communicate: ad hoc and infrastructure. The ad hoc mode is geared for a network of stations within communication range of each other. Ad hoc networks are created spontaneously between the network participants. In infrastructure mode, wireless access points (WAPs) provide for a more permanent structure for the network.

To distinguish different wireless networks from one another, the 802.11 standard defines the service set identifier (SSID). The SSID can be considered the identity element that “glues” various components of a wireless local area network (LAN) together. Traffic from wireless clients that use one SSID can be distinguished from other wireless traffic using a

different SSID. Using the SSID, an AP can determine which traffic is meant for it and which is meant for other wireless networks.

There are a number of different mechanisms for providing security for wireless networks. These include WEP, Remote Authentication Dial In Services (RADIUS) authentication, and the 802.1x standard.

Because of the nature of the 802.11 wireless LANs, the IEEE working group implemented a mechanism to protect the privacy of the individual transmissions. The intent was to mirror the privacy found on the wired LAN (WLAN), and the mechanism became known as WEP. Because WEP uses a cryptographic security countermeasure for the fulfillment of its stated goal of privacy, it has the added benefit of becoming an authentication mechanism. This benefit is realized through a shared-key authentication that allows the encryption and decryption of the wireless transmissions. Up to four keys can be defined on an AP or a client, and they can be rotated to add complexity for a higher security standard in the WLAN policy.

WEP was never intended to be the absolute authority in wireless security. The IEEE 802.11 standard states that WEP provides for protection from “casual eavesdropping.” Instead, the driving force behind WEP was privacy. In cases that require a high degree of security, other mechanisms should be used such as RADIUS authentication, access control, password protection, and virtual private networks (VPNs).

WEP provides for several implementations: no encryption, 40-bit encryption, and 128-bit encryption. Clearly, no encryption means no privacy. When WEP is set to no encryption, transmissions are sent in the clear, and they can be viewed by any wireless sniffing application that has access to the RF signal propagated in the WLAN (unless some other encryption mechanism, such as IPsec, is being used). In the case of the 40- and 128-bit varieties (just as with password length), the greater the number of characters (bits), the stronger the encryption. The initial configuration of the AP will include the setup of the shared key. This shared key can be in the form of either alphanumeric or hexadecimal strings, and must be matched on the client. WEP uses the RC4 encryption algorithm, a stream cipher developed by Ron Rivest (the *R* in RSA).

The Windows Server 2003 family, including Windows Server 2003 Standard Edition, Enterprise Edition, and Datacenter Edition (but excluding Web Edition), comes with the Internet Authentication Service (IAS), which acts as a centralized connection service for wireless security, VPN, and remote-access connections. IAS uses RADIUS to perform centralized user authentication for wireless clients. IAS can be implemented to provide for wireless authentication using the Extensible Authentication Protocol (EAP) for better wireless security. The current IEEE 802.11b standard is severely limited because it is available only for the current open and shared-key authentication scheme, which is nonextensible. To address the weaknesses in these authentication mechanisms, several vendors (including Cisco and Microsoft) adopted the IEEE 802.1x authentication mechanism for wireless networks. The IEEE 802.1x standard was created for the purpose of providing a security framework for port-based access control that resides in the upper layers of the protocol stack. The most

common method for port-based access control is to enable new authentication and key-management methods without changing current network devices.

The 802.1x standards specify the use of EAP over wireless. 802.1x and EAP provide for a mutual authentication capability. This will make the clients and the authentication servers mutually authenticating endpoints and will assist in the mitigation of attacks from man-in-the-middle types of devices.

With the addition of the 802.1x standard, clients are identified by usernames, not by the Media Access Control (MAC) addresses of the devices. This design not only enhances security, but it also streamlines the process for authentication, authorization, and accountability for the network. 802.1x was designed so that it could support extended forms of authentication, using password methods (such as one-time passwords, or GSS_API mechanisms such as Kerberos) and non-password methods (such as biometrics, Internet Key Exchange, and smart cards).

Understanding Wireless Networking

Wireless networks are becoming more and common these days. If you have the right type of adapter card for your laptop, you can sit in a coffee shop and drink your espresso while you check your e-mail. Many universities, airports, and commercial establishments (such as the aforementioned coffee shops) have adopted wireless technologies.

You need to make certain that the equipment you purchase for this technology meets the standards set by one or more of these organizations:

- Institute of Electrical and Electronics Engineers (IEEE)
- Internet Engineering Task Force (IETF)
- Wireless Ethernet Compatibility Alliance (WECA)
- International Telecommunication Union (ITU)

This will help to ensure that your equipment is compatible with that of the network to which you are connecting.

Wireless Network Types

There are a number of different categories of wireless networks, based on the distance data can be transmitted. Three common types are shown in Table 11.3.

Table 11.3 Wireless Network Types

Network Name	Description	Standard
Wireless personal area network (WPAN) <i>aka</i> Bluetooth 1.0 and Bluetooth.	Used for devices such as PDAs, cell phones, and laptops. For extremely short distances, it can also use infrared to connect.	802.15

Continued

Table 11.3 Wireless Network Types

Network Name	Description	Standard
Wireless local area network (WLAN)	Used to connect devices that are in either a corporate or campus environment.	802.11, 802.11b, 802.11a, and 802.11g
Wireless wide area network (WWAN) aka second-generation or 2G system	Used to provide wireless connections over remote public and private networks. These connections can go between cities and countries via antenna sites or satellite systems.	Not yet available
Wireless metropolitan area network (WMAN)	Used to allow for wireless connections between metropolitan areas, such as libraries, without the use of copper or fiber lines. These can also be used as a backup to wired networks.	802.16



EXAM WARNING

802.1x and 802.11x are not the same thing. Do not get these two confused. 802.11x is a wireless standard, and 802.1x is an authentication standard. 802.1x handles authentication using EAP.

EAP Authentication

When EAP authentication is used, there are two types of roles that a LAN port plays during network access on a wireless network: authenticator or supplicant. The EAP supplicant (in this case, the wireless client) communicates with the WAP over an “uncontrolled port.” The WAP sends an EAP-Request/Identity packet to the supplicant as well as a RADIUS-Access-Request packet to the RADIUS access server. The supplicant then responds with an Identity packet, and the RADIUS server sends a challenge based on the Identity packets sent from the supplicant. The supplicant provides its credentials in the EAP-Response packet that the AP forwards to the RADIUS server. If the response is valid and the credentials are validated, the RADIUS server sends a RADIUS-Access-Accept packet to the WAP, which then allows the supplicant to communicate over a “controlled” port. This is communicated by the WAP to the supplicant in the EAP-Success packet.



NOTE

The supplicant is usually the client software, and the authenticator is usually the wireless access point.

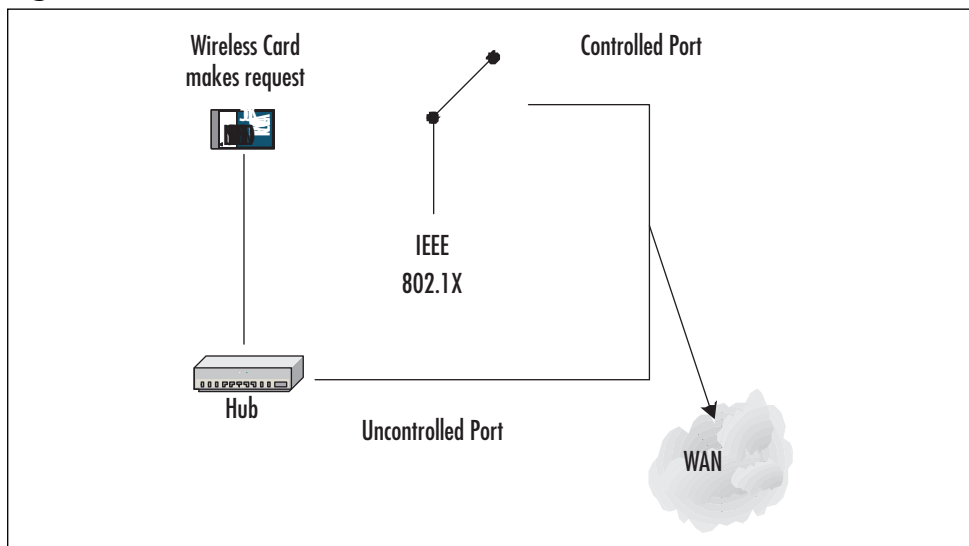
The first data path, the uncontrolled port, allows data exchange between the authenticator and a computing device on the LAN, regardless of the authentication state of that device. This is the path that EAP over LAN (EAPOL) messages will take. EAPOL will encapsulate the EAP messages so that wireless and Ethernet LANs can transport the data securely. The second data path, the controlled port, allows data exchange between an authenticated LAN user and the authenticator. This is the path that all other network traffic will take after the computing device is authenticated. Usually, a RADIUS server is in place to authenticate the supplicant's credentials. However, the 802.1x specification does not require a RADIUS server.

The procedure is as follows:

1. The wireless supplicant receives a challenge from the wireless authenticator.
2. The supplicant forwards its identity to the authenticator.
3. The authenticator sends the supplicant's identification to the RADIUS server.
4. The RADIUS server requests the credential of the supplicant.
5. The uncontrolled port on the authenticator processes the requests that are sent between the supplicant and the RADIUS server, because the supplicant's credentials have not been verified yet.
6. RADIUS receives the supplicant's credentials and then verifies the information.
7. If the verification is successful, the RADIUS server sends the authentication key to the supplicant.

Figure 11.1 shows the process between LAN port roles.

Figure 11.1 Process between LAN Port Roles





EXAM WARNING

Understand the differences between the authenticator and the supplicant. Know the function that each of these roles play on a wireless network.

How Wireless Networking Works

Windows Server 2003 makes it relatively simple to set up and use wireless networking. The operating system has a built-in automatic wireless network configuration so that, with the proper equipment, users can roam from building to building without changing network settings. There are a number of wireless networking technologies in use today, such as the following:

- **WLAN** Also known as WI-FI, RF technology that includes 802.11b, which operates in the 2.4GHz range. 802.11a and 802.11g technologies operate at 5 GHz. These differ in terms of data transfer speed and range. Normal range for 802.11b is about 100 to 300 meters, but with a high-gain directional antenna, this can be extended up to several miles.
- **IrDA** Uses infrared (IR) signals to transmit data for IrDA 1.1. Normal distance is 3 to 6 meters, although some IR technologies have a maximum distance of 1.5 miles. Because IR signals are used to transmit data, distance limits for long-range IR depend on weather conditions (such as humidity). Additionally, IR is a line-of-sight technology that requires a clear path between the transmitter and receiver.
- **Bluetooth** Uses short-range radio waves to transmit data used mainly in PDAs and has a maximum distance range of up to 10 meters.
- **HomeRF** Also known as home radio frequency, works up to only about 150 feet.

Authentication for Wireless Networks

There are two authentication methods in the 802.11 standard: open authentication and shared-key authentication. Open authentication is more precisely described as device-oriented authentication and can be considered as a null authentication—all requests are granted. Without WEP, open authentication leaves the WLAN wide open to any client who knows the SSID. With WEP enabled, the WEP secret key becomes the indirect authenticator.



NOTE

Open authentication can also require the use of a WEP key. Do not assume that just open authentication is used and that a WEP key does not need to be set.

Shared-key authentication is a four-step process that begins when the AP receives the validated request for association. After the AP receives the request, a series of management frames are transmitted between the stations to produce the authentication. This includes the use of the cryptographic mechanisms employed by WEP as a validation. The four steps in the process are as follows:

1. The requester (the client) sends a request for association.
2. The authenticator (the AP) receives the request and responds by producing a random challenge text and transmitting it back to the requester.
3. The requester receives the transmission, encrypts the challenge with the secret key, and transmits the encrypted challenge back to the authenticator.
4. The authenticator decrypts the challenge text and compares the values against the original. If they match, the requester is authenticated. On the other hand, if the requester does not have the shared key, the cipher stream cannot be reproduced. At this point, the plaintext cannot be discovered, and theoretically, the transmission is secured.

One of the greatest weaknesses in shared-key authentication is that it provides an attacker with enough information to try to crack the WEP secret key. The challenge, which is sent from authenticator to requester, is sent in the clear. The requesting client then transmits the same challenge, encrypted using the WEP secret key, back to the authenticator. An attacker who captures both of these packets now has two pieces to a three-piece puzzle: the cleartext challenge and the encrypted ciphertext of that challenge. The algorithm is also known—it's RC4. All that is missing is the secret key.

To determine the key, the attacker simply tries a brute-force search of the potential key space using a dictionary attack. At each step, the attacker tries to decrypt the encrypted challenge with a dictionary word as the secret key. The result is then compared against the authenticator's challenge. If the two match, the secret key has been determined. In cryptography, this attack is termed a *known-plaintext* attack and is the primary reason why shared-key authentication is actually considered slightly weaker than open authentication.

You can use the Wireless Monitor console to determine if the wireless network to which you are connecting has multiple APs. To access the Wireless Monitor console, add the Wireless Monitor snap-in to a custom Microsoft Management Console (MMC).



NOTE

When using 802.1x for enhanced security, authentication is available only to Windows XP Service Pack 1 clients and Windows Server 2003 systems.

EXERCISE 11.01

SETTING UP A WINDOWS XP CLIENT FOR WIRELESS NETWORKING

Installing and configuring wireless networking on a Windows XP client is simple. Make sure you have an 802.11b wireless network interface card (NIC) installed. After the wireless NIC has been installed, the **Automatic Wireless Wizard Configuration** window appears. Windows XP will automatically search the network for a WAP. If a WAP is found, it will attempt to make a connection.

To manually configure the wireless network connection on the Windows XP machine, use these instructions:

1. Select **Start | Control Panel | Network Connections**.
2. Right-click the **Wireless Connection** and click **Properties**.
3. Select the **Wireless Networks** tab.
4. Look in the **Available Networks** box and choose a WAP to add under the **Preferred networks** option.
5. In the **Wireless Network Properties** dialog box, enter the name of the WAP under the **Network Name (SSID)** box and place check marks in the appropriate boxes that apply to your network settings.
6. Click the **OK** button. Wireless network access should now be available to you from your Windows XP machine.



TEST DAY TIP

Understand how open system and shared key work as authentication subtypes. Know the weaknesses and strengths of each, as well as how they are configured on a domain controller and client machine.

Defining a Subtype on a Domain Controller

You can configure the subtype you wish to use on a domain controller by configuring Group Policy. To configure the subtype in Group Policy, follow these steps:

1. Select **Start | All Programs | Administrative Tools | Active Directory Users and Computers**.

2. Find the domain for which you wish to configure the subtype, right-click the domain node, and choose **Properties**.
3. Select the **Group Policy** tab.
4. In the GPO Editor, expand **Computer Configuration**, and then expand **Windows Settings**.
5. Select the **Wireless Network IEEE 802.11 Policies** node.
6. Double-click the policy in the right-console pane.
7. On the **Preferred Networks** tab, click **Add** to add a new wireless network.
8. Double-click the name of the new wireless network and choose the **Network Properties** tab.
9. Type in a name for the network. You can choose to enter a description if you wish. You can also select the **Network Authentication (Shared Mode)** box under **Wireless network key (WEP)**. If you leave this blank, open-system authentication will be used instead.
10. To allow the key to be provided automatically for client machines, check the **The key is provided automatically** box.



NOTE

To ensure that 802.1x will be used for wireless access control, click the **IEEE 802.1x** tab and check the **Enable network access control using IEEE 802.1x** check box.

Defining a Subtype on a Client Computer

Next, you can enter the subtype on a client-computer. To do this, perform the following steps:

1. Open **Network Connections**. Right-click **Wireless Network Connection**.
2. Click **Properties** and click **Add** on the **Wireless Networks** tab to add the wireless network connection.
3. On the **Association** tab, for **Network name (SSID) service set identifier**, enter a unique name. Configure the setting for the use of a network key for data encryption if needed. Next, choose the **Data encryption WEP enabled** option if it is not already enabled.
4. To specify that a network key be used for authentication to the wireless network, select the **Network Authentication (Shared mode)** check box.

5. If the network key is automatically provided for dynamically, leave the **Network Key** option blank. If the network key is not automatically provided, type the key in this field, and then type it again in the **Confirm network key field**.
6. Click the **Authentication** tab to specify that 802.1x authentication is being used for the wireless network connection.

Authentication Protocols

Wireless clients and servers can use different types of authentication protocols.

Authentication is the method used to verify a user's identity when the user is trying to access network resources. Windows Server 2003 offers several types of authentication protocols that allow users to use the *single sign-on* method to access the network.

EAP

EAP is used by 802.1x to take care of authentication. EAP handles conversations between wireless clients and servers. In order for the conversation to work properly, both the client and server must use the same authentication method. EAP is an IEEE standard method for authentication.

EAP-TLS

The strongest authentication method is EAP-Transport Layer Security (EAP-TLS), which is used in certificate-based wireless networks. If you are using smart cards or certificates for client authentication, this is the only protocol that supports those types of authentication. EAP-TLS is the strongest method of authentication, because authentication is mutual between the client and server machine. Also, EAP-TLS is not dependent on the user account password, and the client does not need to intervene when using the authentication method.

EAP-MS-CHAPv2

EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (EAP-MS-CHAPv2) can be implemented when you do not wish to use certificates to authenticate your wireless users. Instead, it uses passwords to authenticate wireless clients. This type of authentication can be installed only on RADIUS servers. After the authentication, users can be allowed to change their passwords. EAP-MS-CHAPv2 is a mutual authentication method that supports password-based user or computer authentication and is available only with PEAP.

PEAP

Protected Extensible Authentication Protocol (PEAP) is a new addition to the EAP extensions. PEAP works by creating an encrypted channel from the wireless client to the authenticator of the wireless session. To create the encrypted channel, PEAP uses TLS. Because TLS creates a secure channel between the client and authenticator, it protects against attacks

such as denial of service (DoS). PEAP does not support the use of VPN clients or remote-access clients, because it does not have an encryption method; rather, it allows other authentication protocols to use its TLS channel for much improved security.

To configure PEAP, use the following steps:

1. Select **Start | All Programs | Administrative Tools | Internet Authentication Services**.
2. Click **Remote Access Policies**.
3. Open the policy you wish to configure by double-clicking the policy.
4. Click **Edit Profile** and choose the **Authentication** tab.
5. In the **Authentication** tab, select **EAP methods**.
6. From the **Select EAP providers** option, click the **Add** button and select the **Protected EAP (PEAP)** option.
7. Click the **OK** button.
8. Select **OK**.
9. Click **Apply**, and then click **OK**.



EXAM WARNING

Understand all of the authentication protocols that are available and remember which protocols work best for scenario-based use.

Using IAS with Wireless

IAS can be used to authenticate wireless connections. Before you begin to implement IAS with your wireless network, you need to plan the deployment carefully. Use the following guidelines as you step through the IAS and wireless configuration process to optimize your setup:

- Before you start to configure IAS with wireless networking, make sure that all of the equipment being used meets the IEEE specifications. If you do not have hardware that meets these requirements, your deployment might not work properly.
- Determine whether or not you will be using certificates for authentication. If you will be using certificates, install the certificates on the IAS server and client computers.
- Install IAS on the server.

To install IAS, follow these steps:

1. Select **Start | Control Panel | Add/Remove Windows Components**.
2. Select **Networking Services**, and then select **Detail**.
3. Scroll down to the **Internet Authentication Service** option and put a check mark in the box. Click **OK**.

To open IAS after it has been installed follow these steps:

4. Open IAS by selecting **Start | Programs | Internet Authentication Services**.
5. To add RADIUS clients, open the **Internet Authentication Service** console.
6. Select the **Add Radius Client** option
7. Choose the **New Radius Client** option.
8. Enter the **Friendly Name** of the RADIUS client.
9. Enter the **IP address** of the client that can be verified by selecting the **Verify** button.
10. When the information has been entered, click the **Next** button and choose the **Client-Vendor** attribute of the RADIUS client.
11. Enter the **Shared Secret password** and enter it again under the **Confirm Shared Secret**.
12. If you want to require the Message Authenticator attribute in the request, click the **Request must contain the Message Authenticator Attribute** box.
13. Click **Finish** to complete the process.

Wireless Security Issues

There are many threats to networks today. With the advent of wireless networking, the risks have grown. This is especially true if you do not secure the wireless network. Hackers can download and run tools such as NetStumbler, AirSnort, and WEPCrack to locate and access your network without any type of advanced skills. Even if no one with malevolent intent is within range of your wireless network, it is not uncommon for persons with wireless equipment to accidentally access networks they are not authorized to access.

The following are some additional wireless security issues:

- **Rogue WLANs** Someone inside your organization can purchase the inexpensive equipment and a wireless LAN card and install a wireless network, connected to your wired network, for their personal use.
- **Spoofing** An outside user can send harmful files or viruses to internal users by pretending to be an insider.

- **Drive-by and freeloading** Many hackers, who call themselves “war drivers” go out with wireless-equipped laptops and high-gain Yagi antennas, looking for open wireless networks in an attempt to gain free access to Internet bandwidth or access to network data.
- **Eavesdropping** This is when a person uses a protocol analyzer (sniffer) to capture unprotected network traffic and open the packets to read the data inside.
- **DoS attacks** Attackers can flood your network with unauthorized traffic and bring it to a halt.

Because many businesses do not understand the security and encryption methods involved with a wireless network rollout, these security measures are not implemented. Use the proper encryption and security methods described in this chapter if you implement wireless technologies in your organization.



NOTE

AirSnort is used to capture WEP keys. It is sometimes called an auditing tool, but it can also be used to sniff out data packets on WEP. AirSnort works by capturing “interesting packets” with weak keys. If you are using 128-bit encryption, WEP cards can generate roughly 16 million keys, and of those some 9000 are weak. Windows operating system users should be aware that AirSnort runs only on Linux boxes. It is assumed that many passwords can be guessed after about 2000 or so interesting packets.

Default Settings

Some manufacturers ship wireless network devices with default settings that are well known in the industry, so any malicious user who has a bit of knowledge (experience with the particular device model or the ability to do a little Internet research) already has a head start on the information needed to break into your network.

The default settings that you should be especially concerned with (and should immediately change when you set up your wireless network) include the following:

- Administrative password
- SSID
- SSID broadcasting
- WEP settings

Administrative Password

When setting up WAPs, the first thing you should do is change the password on the device. Many devices come preset with the Administrator password set to something like admin. Do not leave these devices in their default state. Anyone who knows that you have this hardware could easily change your WAP configuration with this information.

As with any administrative-level password, you should set a strong password on the WAP administrator account. Most WAP vendors provide Web-based administration for performing configuration tasks such as changing the password.

SSID

The SSID acts like a network name for your wireless access device. It has a 32-character identifier attached to the header of data packets that are sent over the WLAN. When this information is sent from the mobile device to the WAP, it acts like a password, in that devices that cannot provide correct SSID information to the WAP will not gain access to the network.

However, the SSID does *not* provide wireless security for a couple of reasons. First, the SSID is sent over the network in plaintext and is not encrypted. This means that an outsider can use sniffer software to capture the packets being sent over the network and obtain the SSID from them. Additionally, by default, most broadcast their SSIDs, bringing us to the next section.

SSID Broadcasting

SSID broadcasting makes it easy for wireless clients to find the WAP. You don't even need to know the SSID of the network in order to configure your client to connect; if the WAP broadcasts its SSID, your client computer will automatically intercept it and provide it in the list of wireless networks that are available for you to connect to.

WEP Settings

Another issue to be aware of is that most WAPs do not have WEP enabled by default. Although WEP has some weaknesses (which we'll discuss in the next section) and is considered by many to provide weak security, it does provide *some* security and thus should be enabled as a best practice.

Remember that when you enable WEP on the AP, you also must set the client to use it. On a Windows XP client, follow these steps:

1. Double-click your wireless connection icon in the system tray or in **Control Panel** and select **Network Connections** to open the Connection Status dialog box.
2. Click **Properties**.
3. In the Properties box, click the **Wireless Networks** tab.

4. Under **Available networks**, select the SSID of the network you want to configure.
5. Click the **Configure** button.
6. Under the section labeled **Wireless network key (WEP)**, check the check box that says **Data encryption (WEP enabled)**.
7. If the key is provided automatically, check the check box that says **The key is provided for me automatically**.
8. If the key is not provided automatically, enter the key in the field labeled **Network key**.
9. Select the **Key format** (ASCII characters or hexadecimal digits) and the **Key length** (40 bits for a 5-character key; 104 bits for a 13-character key).

WEP Weaknesses

One factor that makes WEP vulnerable is the lack of a defined key-management process. As discussed previously, WEP uses the RC4 encryption algorithm. This is a symmetric (secret key) algorithm, which means the same key is used both for encrypting and for decrypting the data that travels between the WAP and the wireless client. Thus, all clients must share this same key. This is obviously a weak system because the key must be disseminated so widely, and users must manually enter the key into their wireless configuration.

If someone steals, or even gets temporary access to, one of the organization's wireless clients, he or she can easily discover the shared key. Many organizations don't update the keys frequently. Manually rekeying every wireless device can take a lot of time and effort. If the WAP provides the key automatically, this makes changing the key easier.

The second problem with WEP is weak encryption. WEP can use either a 40-bit or 104-bit key. A 40-bit key is fairly easy to break. One would assume that the 104-bit key would be significantly more secure. However, it's not quite that simple. The problem is the 24-bit initialization vector (IV) that is appended to the WEP key (this creates a 64-bit or 128-bit encryption key that is used to encrypt the data and its checksum). Regardless of the key length, the IV is still only 24 bits. This short IV ensures that the same key stream will be reused. If attackers capture multiple packets that use the same key stream, they can use analysis software to break the encryption.

Attackers can sniff the network and capture packets to determine the IV. By performing an XOR operation on two captured packets, the key can be extrapolated. A hacker needs no special talents to break the WEP key—software such as WEPCrack can do it.

Making Wireless More Secure

Wireless security experts recommend that highly sensitive data should never be sent over a wireless network. If it must be, you should use higher level security mechanisms rather than rely on WEP. If your security requirements are low, you can make WEP more secure by

implementing measures such as broadcast key rotation. There are new key-management schemes that use “dynamic WEP.” Instead of a static base key, they provide an updated base key at frequent intervals. Unfortunately, these systems are proprietary and do not really offer a high level of protection. Even with updates every half hour or less, it generally takes only a few hours for a skillful hacker to capture enough packets to analyze and use that information to breach the network.

Another possible solution is hardware that supports the Advanced Encryption Standard (AES). In addition, organizations are investigating the use of IPSec to secure wireless communications. It is possible to incorporate IPSec VPNs into WAPs for high security, but there are some issues with firewalls that do not allow IPSec traffic through. There are also compatibility problems using IPSec with some implementations of network Address translation (NAT), although NAT traversal is a new technology that makes it possible for IPSec and NAT to work together.

A new IEEE subcommittee is working on the WEP problems, under the standard 802.11i. One of their proposed solutions is called the Temporal Key Integrity Protocol (TKIP). This combines a 128-bit “temporal” key with the MAC address of the wireless client and adds an IV. After 10,000 packets, the temporal key changes. Another new technology that uses this TKIP process, WiFi Protected Access (WPA), has been developed by a cooperative effort of the WiFi Alliance and the IEEE. It uses a stronger IV (48 bits), so that the number of possible keys generated is much greater. WPA also uses per-packet key mixing and a hierarchical key structure that can prevent attacks such as those that use AirSnort. A second version of WPA, WPA2, is expected out in 2004 and will use the Counter with Cipher Block Chaining Message Authentication Code (CCMP) protocol, which is based on AES.

If you do rely on WEP as your primary security for your WLAN, it is important to change the key often. You can also combine other protective mechanisms, such as MAC filtering, with WEP. MAC filtering allows you to specify only particular client MAC addresses that can connect to your WAP. This is set up on the WAP, and the procedure differs depending on the manufacturer. Be aware, though, that knowledgeable hackers can capture packets sent by legitimate wireless clients, discover their MAC addresses, and then spoof the address. So like WEP, MAC filtering offers only a small measure of security. Centralized authentication (using RADIUS/IAS with EAP/PEAP) is Microsoft’s most important solution to overcoming wireless security vulnerabilities.



TEST DAY TIP

As a network administrator, you need understand the weaknesses of the wireless network and why WEP is not the best protocol for security.

EXAM 70-293
OBJECTIVE
5
6
6.3
6.3.1

Monitoring and Optimizing Security

Monitoring security in Windows Server 2003 is an important task that you should not take lightly. Windows Server 2003 has added functionality that will enable you to monitor for security breaches on a daily basis and move on with your other administrative duties. The results of your monitoring should be used as a basis for evaluating and optimizing the security measures in place on your network. You can use the following tools to help monitor and optimize security in your organization:

- Wireless Monitor
- Object-based access control
- Auditing
- Security policies

Each of these tools is discussed in the following sections.

Wireless Monitor

If you have implemented wireless networking in your organization, you will need to monitor your wireless network for day-to-day activity. Good security begins with being aware of what is happening on your network.

The Wireless Monitor MMC snap-in allows you to view wireless network activity such as MAC addresses, network names, strength of the signal, and the data rate that the wireless network will support, just to name a few available monitoring details.

You can view information about both the WAP and the clients. There are two nodes in the left pane of the console: **Access Point Information** and **Wireless Client Information**. To enable or disable logging of client information, right-click the **Wireless Client Information** node and make the appropriate selection.

The information shown in the Wireless Monitor console is collected from the Wireless Configuration service. You will be able to view configuration changes and track the reaction of the Wireless Configuration service to events generated outside the network. The client events you can log include the following:

- The addition of a NIC (wireless or otherwise)
- Association to a network using a random WEP key
- Failure to detect the presence of an ad hoc network
- Failure to associate to an infrastructure (WAP) network
- Failure to associate to any wireless network
- Successful association to a wireless network
- Driver errors

- Resetting of the network interface
- Beginning or completing a scan for wireless networks
- Configuration changes needed
- Processing of commands
- Device arrival notifications
- Media connect and disconnect notifications
- Timeout notifications



NOTE

You do not need to have administrative credentials to add the Wireless Monitor snap-in and view wireless information.

Object-based Access Control

Object-based access control can be added to a GPO to allow administrators to apply or modify auditing policy settings for items such as the following:

- Files and folders (object auditing must be enabled in Group Policy before these items can be audited)
- Printers
- Services
- Registry keys

Auditing

Administrators should set up auditing on items that have high sensitivity (for example, folders and files containing confidential documents). You should *not* attempt to audit all items, because this would create a system and work overload for the following reasons:

- Auditing uses memory and processor resources, resulting in performance issues.
- Audit logs use hard disk space (and can use a considerable amount if you audit many items).
- Sorting through hundreds or thousands of entries can make it difficult for you to find the audit events that are important.

Auditing must be turned on before it will work. To turn on auditing for a local computer (not a domain controller), perform the following steps:

1. Select **Start | All Programs | Administrative Tools**.
2. Select **Local Security Policy**.
3. In the left pane of the GPO Editor, expand **Local Policies**.
4. Click **Audit Policy**.
5. In the right console pane, double-click **Audit object access**.
6. In the **Local Security Setting** tab, shown in Figure 11.2, under **Audit these attempts**, check the **Success** and/or **Failure** check box depending on the type of attempts you want to audit.
7. Click **OK**.



NOTE

If the computer is a member of a domain, the domain policy will override local settings. If the Success and Failure check boxes are grayed out in the **Security Policy Setting** dialog box and you cannot change them, this means that a domain policy is in effect.

Figure 11.2 Enabling Object Access Auditing



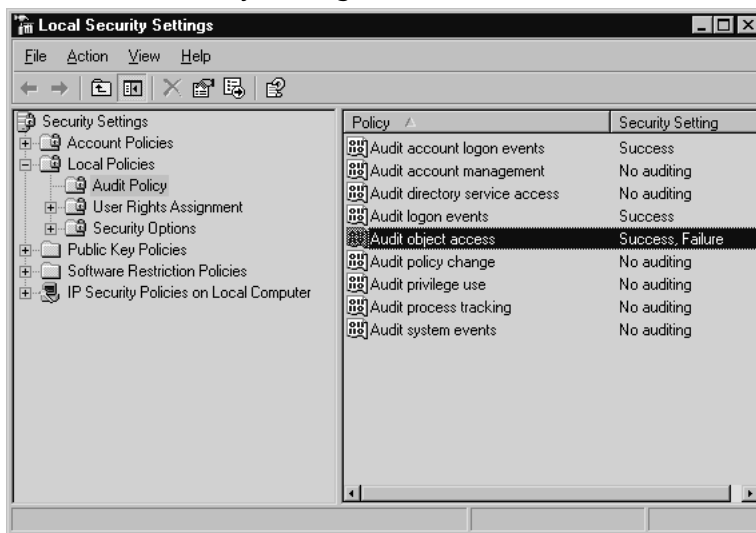
To turn on auditing for a domain controller, follow these steps:

1. Select **Start | All Programs | Administrative Tools | Active Directory Users and Computers**.
2. Right-click the domain name in the left console pane and select **Properties**.
3. Select the **Group Policy** tab.
4. Select the policy you want to edit (**Default Domain Policy**) and click **Edit**.

5. In the left pane of the GPO Editor, expand **Computer Configuration**, then **Windows Settings**, then **Security Settings**, then **Audit Policy**, then **Local Policies**.
6. In the right console pane, double-click **Audit object access**.
7. In the **Security Policy Setting** dialog box, check **Define these policy settings**. Check the **Success** and/or **Failure** check box depending on the type of attempts you want to audit. For example, if you want to know only when someone attempts to open a folder on which you enable auditing and is unable to do so, check only the **Failure** check box. If you want to know when anyone tries to access the folder whether successful or not, check both check boxes.
8. Click **OK**.

After you have enabled object auditing, the security setting appears in the right console pane, as shown in Figure 11.3.

Figure 11.3 Turning on Auditing for Object Access Using the Local Security Settings Console



NOTE

After you make the change to the audit policy, you might need to run **gpupdate** to refresh the policy before it will take effect.

Now that object auditing is enabled, you must pick the objects you want to audit and set auditing on their properties. To set auditing on a folder, perform the following steps:

1. In **Windows Explorer**, navigate to the folder, right-click it, and select **Properties**.
2. Click the **Security** tab, and then click the **Advanced** button.
3. In the **Advanced Security Settings** dialog box, click the **Auditing** tab.
4. To add an auditing entry, click the **Add** button.
5. In the **Select User or Group** dialog box, enter the name of a user or group whose access to the object you want to audit. If you want to audit all access, select the Everyone group.
6. In the **Auditing Entry for <foldername>** dialog box, select the access types you want to audit. For example, if you only want to know when someone reads the data, select **List Folder / Read Data**. Check the **Successful** and/or **Failed** check box to indicate which type of access you want to audit. For example, if you only want to know when someone tries unsuccessfully to read data, check the **Failed** check box.
7. Click **OK**.

On the **Auditing** tab, you can also select whether you want to allow inheritable auditing entries from the parent object to propagate to this object and its child objects. If you check this box (which is checked by default), any auditing properties set on a parent object will flow down through the tree to all objects under it. Additionally, you can choose to replace the auditing entries that are on this object's child objects with the ones you are setting here.

Auditing Registry Keys

To audit Registry keys, use the following steps:

1. Drilling down from **Computer Configuration**, select **Windows Settings**, **Security Settings**, and then **Registry**.
2. Right-click the **Registry** option and choose **Add Key**.
3. Browse to the key that you wish to audit and click the **OK** button.
4. To change settings on a file or folder that has been added to this GPO, in the details pane, right-click the Registry key and click the **Properties** button.
5. Click **Edit Security**, choose **Advanced**, and click the **Auditing** tab.
6. Click **Add** and type in the name or the group that you wish to audit.
7. From the correct entry, choose the **Apply Onto** list.
8. From the **Access** box, choose the actions that you wish to audit. Select **Successful** to audit successful events, **Failed** to audit events that failed, or deselect these actions to stop auditing those events. To stop all auditing, click the **Clear All** box.

- To stop any subfolders or files in the tree from inheriting these audit properties, choose the **Apply these auditing entries to objects and/or containers within this container only** box.

Auditing Files or Folders

To audit files or folders, use the following steps:

- Drilling down from **Computer Configuration**, select **Windows Settings**, then **Security Settings**, and then **File System**.
- Right-click the **File System** option and choose **Add File**.
- Browse to the file that you wish to audit and click the **OK** button.
- To change settings on a file or folder that has been added to this GPO, from the details pane, right-click the file or folder and click the **Properties** button.
- Click **Edit Security**, choose **Advanced**, and click the **Auditing** tab.
- Click **Add** and type in the name or the group that you wish to audit.
- From the correct entry, choose the **Apply Onto** list.
- From the **Access** box, choose the actions that you wish to audit. Select **Successful** to audit successful events, **Failed** to audit events that failed, or deselect these actions to stop auditing those events. To stop all auditing, click the **Clear All** box.
- To stop any subfolders or files in the tree from inheriting these audit properties, choose the **Apply these auditing entries to objects and /or containers within this container only** box.



TEST DAY TIP

If check boxes are not available in the **Advanced Security Settings** dialog box for a particular file or folder, this means that auditing has been inherited from the parent folder. This is also the case if in the **Auditing Entry** dialog box for the file or folder, the **Remove** option is not available.

Viewing the Results of Auditing

After you have enabled object auditing in Group Policy and set auditing on at least one object, you can view the results in the Security log in Event Viewer.

Security Log Settings

Auditing security events can take up a great deal of system resources and disk space. This is especially true when you audit items such as logon and logoff times, folder access, OUs, and users. Make sure that you either have a large amount of hard drive space available on the server or you set your Security log to overwrite itself at a certain point, such as if the log reaches 50MB of space. To configure your event log settings, use the following steps:

1. Select **Start | Administrative Tools | Event Viewer**.
2. Choose **Security Log**, right-click the log, and select **Properties**.
3. In the log's property page, you can change the maximum log size by entering the size in the **Maximum Log Size** field.
4. To set the log to overwrite itself as needed (when the maximum size has been reached) or when it is a certain number of days old, click **Overwrite events older than** and selecting the number of days in the drop-down box. You can also specify that events should not be overwritten if you want to always clear the log manually.
5. To clear the log, click the **Clear log** option. You will be prompted to save the Security log file before the clear is complete. Enter a name for the Security log file, and it will be saved with an .evt extension in a specified location. It can then be opened in Event Viewer for later viewing.



TEST DAY TIP

Understand how to audit users for folder and object access. Also, remember that auditing must be enabled in Group Policy before it can function and log audited events.

Security Policies

Windows Server 2003 makes it easy to set security policies on local computers or for a domain, using Group Policy. To set security policies on a local computer, open the Local Security Policy GPO by selecting **Start | All Programs | Administrative Tools** and selecting **Local Security Policy** (you will not find this option on domain controllers). To set security policies in a domain, edit the default domain policy as follows:

1. Select **Start | All Programs | Administrative Tools | Active Directory Users and Computers**.
2. Right-click the domain node in the left pane and click **Properties**.
3. Choose the **Group Policy** tab.

4. Select the **Default Domain Policy** and click **Edit**.
5. In the left pane of the GPO Editor, expand **Computer Configuration**, then **Windows Settings**, then **Security Settings**.

In either case, you will see the following folders under Security Settings:

- **Account Policies** Password, Account Lockout and Kerberos policy settings.
- **Local Policies** Audit, User rights assignment and Security options, Guest account names, CD-Rom access, driver installation and logon prompts.
- **Public Key Policies** Certificate submission, certificate requests and installations and create then distribute certificate trust lists.
- **Software Restriction Policies** Used to create hash rules, certificate rules. File identity through a specified path and the ability to create an internet zone rule.
- **IP Security Policies** Used to create and manage IPSec security policies.

In the case of the domain policy, you will also see other entries under Security Settings, including Restricted Groups, System Services, Registry, File System, and Wireless Networks.

Some of the most important aspects of your security strategy include the configuration of password policies, Kerberos policies, account lockout policies, and user rights policies. In the following sections, we will discuss each of these in more detail.

Password Policies

Password policies allow administrators to enforce password history, age, and complexity and also use reverse encryption. Some options you can enforce include:

- **Enforce password history** This will allow or disallow the availability of password histories. Changing this setting to 1, for example, would make it impossible for the user to use the last password he or she had used and force the user to create a new one. However, the user could reuse the password he or she had prior to the current one, because only one previous password would be remembered. The default setting is 24.
- **Maximum password age** This is used to set the age for the password. After the number of days specified has been met, the password will expire and need to be reset. If password history is enforced, the user will need to change the password to a new one. If not, the user could reenter the same password. It makes little sense to set a maximum password age unless you also enforce password history. The default setting is 42 days.
- **Minimum password age** This setting allows you to specify a time that a password must be in effect before it can be changed again. You can use this to keep users from continually changing their passwords. The default setting is one day.

- **Minimum password length** This setting allows you to set a minimum length for passwords. If a user tries to set a password that is less than the minimum specified length, he or she will receive a message that the password is unacceptable. The default setting is seven characters.
- **Password must meet complexity requirement** You can set passwords to meet a complexity requirement, which means they must contain both uppercase and lowercase letters and numeric characters to make them more difficult to guess. If a user tries to set a password that doesn't meet the complexity requirements, he or she will receive a message that the password is unacceptable. This would stop users from using dictionary or easy-to-guess passwords. This policy is set to either Enabled or Disabled (you cannot define the specific required characteristics for the password within this policy; Microsoft has preset them).
- **Store password using reversible encryption for all users in the domain** This setting allows you to use reversible encryption to store user passwords. Reversible encryption is not secure because it is the same as storing password in plaintext. This setting should be used only when necessary for compatibility purposes. By default, it is disabled.

Kerberos Policies

Kerberos policies are used for domain user accounts only. They determine Kerberos-related settings, such as ticket lifetimes and enforcement. Kerberos policies do not exist in Local Computer Policy. By right-clicking the policy, you can change the following options:

- **Enforce User logon restrictions** Enabling this could slow your network performance. This is used to specify whether the Kerberos v5 Key Distribution Center (KDC) validates each request it receives for a session ticket against the target computer's user rights policy.
- **Maximum lifetime for service ticket** This setting must be greater than 10 minutes and less than or equal to the setting for the **Maximum lifetime for user ticket** setting. This is used to specify the time in minutes that a granted session ticket can be used to access a specified service.
- **Maximum lifetime for user ticket** This setting is used to specify the time in hours that a user's ticket granting ticket (TGT) can be used. A new TGT must be requested when the old one expires. By default, this is set to 10 hours.
- **Maximum lifetime for user ticket renewal** This is used to determine the time in days during which a user's TGT can be renewed. The default is seven days.
- **Maximum tolerance for computer clock synchronization** This can be used to prevent replay attacks. This setting will determine the maximum time in

minutes that can differ between the time on a server clock and the time on a user clock.

Account Lockout Policies

Account lockout policies are used by administrators to lock out an account when someone tries to log on unsuccessfully several times in a row. We can usually assume that a legitimate user might type his or her password incorrectly once or twice, but not numerous times. Thus, numerous failed logons can indicate that someone is trying a brute-force password attack (trying to keep guessing the password until he or she gets it right). There are three options:

- **Account lockout duration** You can specify the time in minutes that the account can be locked out. For example, if the account locks out for two hours, the user can try again after that time. The default is no lockout. When you define the policy, the default time is 30 minutes. The setting can be from 0 to 99,999. When set to 0, the account will remain locked out until an administrator manually unlocks it.
- **Account lockout threshold** This specifies the number of failed attempts at logon a user is allowed before the account is locked out (for example, three). After the threshold has been reached, the account will be locked out. If this value is set to 0, the account will not lock out. This setting can be from 0 to 999.
- **Reset account lockout counter after** You can choose to have the account lockout counter reset after a number of minutes. At that time, the count will start over at one.

User Rights

When configuring user rights, it is important to remember how user rights differ from permissions:

- Permissions are attached to a specific object such as a file, folder, or printer and determine who can access that object and what level of access is granted.
- User rights are attached to user (or group) accounts and determine what a user can do on the computer in general, not in regard to a particular object.

User rights can be assigned to group accounts or individual accounts. These assigned rights will give the user the ability to do specific tasks, such as backing up and restoring system files, shutting down the server, and installing programs. It is easier and more efficient to assign user rights to groups rather than to specific users. For instance, if you work at a large company with more than 10,000 users, it is much easier to manage a user group with

1000 users than to manage the 1000 users individually. If you need to remove user rights for a user, you can just remove that user from the group.

Users are granted two types of rights: *logon rights* and *privileges*. Logon rights will give the user the right to log on to the system locally. Privileges are set to give the user access to perform specific tasks on the computer.



TEST DAY TIP

Understand the difference between user rights and permissions. Also be able to distinguish between logon rights and privileges. Know what each of these allows in regard to security pertaining to a specific user.

Security Templates

You can use the Security Configuration and Analysis MMC snap-in to apply specific security templates to a local machine. The templates can also be imported into Group Policy and applied throughout the domain. There are two different types of security templates available in the Windows Server 2003 family:

- Local Security Policy, which is used to maintain local security settings on a computer
- Security Settings extension to Group Policy, which is used to maintain security settings in AD for a domain, a site, or an OU

With the Security Configuration and Analysis console, you can compare the current security settings to a database that uses one or more security templates, or directly configure security by importing templates and applying them to the local computer. To access and use this console, follow these steps:

1. Select **Start | Run**, type **mmc**, and click **OK**.
2. In the console, **File | Add/Remove Snap-in**.
3. Select **Add**, choose **Security Configuration and Analysis** from the list, and click **Add**.
4. Click **Close**, and then click **OK** to open the console.
5. To create a new analysis database, right-click the **Security Configuration and Analysis** item and choose **Open Database**.
6. Type in a new database name and click **Open**.
7. Choose a security configuration file to import and click **Open**. Seven templates are available:

- COMPATWS, which is compatible with the default permissions for workstations and servers granted to administrators, power users, and users (should not be applied to domain controllers)
- HISECDC, for a high-security domain controller
- HISECWS, for a high-security workstation
- ROOTSEC, for system root security, specifies root permissions (can be used to reapply root directory permissions if they are changed)
- SECURED, for a secure domain controller
- SECUREWS, for a secure workstation
- Setup Security, for default security, created during the installation of each computer and contains the default security settings (can be used on clients and servers but not domain controllers)



NOTE

Never modify the default `security.inf` file because it is the default security template. If you make a modification and save the file as `security.inf`, you will have changed the default template. Then, if you have a problem, you will not be able to set security back to the default level.

An additional tool that you can use to automate security configuration is the `Secedit` command-line utility. This utility can be handy for administrators who like working at the command prompt instead of dealing with GUI. Type **`secedit`** at the command line to see the command's syntax, as shown in Figure 11.4. The following switches are available:

- **`secedit /configure`** Used to configure security information that will be stored inside a database. Type **`secedit /configure`** at the prompt to see a detailed list of this command's multiple parameters.
- **`secedit /import`** Used to import security configuration settings into a computer for use or to be analyzed. This command also has many parameters which can be viewed by typing the command at the prompt.
- **`secedit /export`** Allows administrators to export security settings stored in the database. The syntax is:

```
secedit /export /DB filename /mergedpolicy /CFG filename /areas area1
        area2 /log filename and /quiet
```

- **`secedit /analyze`** Allows administrators to analyze the settings on a computer using a baseline stored in a database. The syntax is:

```
secedit /analyze /db filename.sdb /cfg /overwrite /log filename
/quiet
```

- **secedit /validate** Validates the syntax of a security template that you are trying to import into a database. The syntax is:

```
secedit /validate /cfg filename
```


- **secedit /generaterollback** Used when you wish to configure a rollback template. This will reset the security settings to values that were applied before the newer security template. The syntax is:

```
secedit /generaterollback /CFG
```

The CFG is the security template that you are creating a rollback template of:

- filename.inf /RBK (this is the name of the rollback security template that is being created)
- /SecurityTemplatefilename.inf
- /log RollbackFileName.inf /quiet

Figure 11.4 The Secedit Command Syntax



EXAM WARNING

Know how and why you would use the secedit.exe syntax and switches. Understand that you also need to know the security database name in order to use this tool at the command prompt.

EXAM
70-293OBJECTIVE
5
5.4
6
6.3
6.3.1

Planning a Change and Configuration Management Framework

It is an important part of the network administrator's job to plan a change and configuration management strategy. A network is not a static entity; it is continuously in need of updates and changes. Exercising proper control over this process will help you to increase the productivity of users within the organization and lower the total cost of operations of the network equipment and systems.

Security is a vital part of maintaining a productive network in today's business environment. You can use the tools provided by Microsoft with Windows Server 2003 to ensure that when changes are required in your network environment, they are first properly tested and then effectively and efficiently deployed without creating opportunities for attacks on the network, breach of data confidentiality, or unintentional introduction of malicious code and viruses.

In this section, you will learn how to plan a security update infrastructure and how to use two Windows Server 2003 utilities: Microsoft Baseline Security Analyzer (MBSA) and Microsoft Software Update Services (SUS).

EXAM
70-293OBJECTIVE
5
5.4
6
6.3
6.4

Planning a Security Update Infrastructure

Even the best security practices can prove ineffective if you do not perform routine security updates on your servers and client computers. Windows Server 2003 has valuable add-on tools that enable administrators to analyze security flaws, as well as to update servers and client machines from a single location. This can save time for administrators and provide a more secure environment. Continuous monitoring and updating of your network infrastructure will allow you to maintain continuity and strong security on your network and client machines. Some of the add-on tools available include Subinacl.exe and Permcopyp.exe.

Subinacl.exe allows administrators to gather information pertaining to files, Registry keys, and services. This information can then be transferred from user to user, local to local group, global to global group, and domain to domain. For example, you could use this tool to adjust the files for a user share after the user has been moved from the TEST1 domain to the TEST2 domain. The domains TEST1 and TEST2 have a trust relationship already established. The two domains must be trusted for this command to work properly. The server name is SERVER1 with the user share name USER1. To perform this task, you would click **Start** | **Run**, type **cmd**, and click **OK** to open a command prompt window. In the command prompt window, issue the following command:

```
subinacl /subdirec\\server1\user1\*. * /replace=TEST1\USER1=TEST2\USER1
```

Permcopyp.exe allows administrators to copy share-level permissions from one share to another. For example, you could use this tool to copy from a share named david on test-server1 to the testserver2 share called tom. To perform this task, open a command prompt window and issue the following command:

```
permcopyp \\testserver1 david \\testserver2 \tom
```

This will copy the share-level permissions from the testserver1 david share to the testserver2 tom server.

Understanding the Importance of Regular Security Updates

The machines on your network are always at risk, and no operating system is totally secure. New exploits are discovered every day, and hackers quickly pass the news. Hackers even create scripts and programs to allow others, who are less technically skilled, to exploit protocol, operating system, and application vulnerabilities.

Not maintaining updates on your client and server machines could be compared to purchasing a new car and never having routine maintenance done. On the other hand, it is important to note that not all security updates necessarily be applied to your computers. This is because, unfortunately, in the rush to fix problems that are brought to their attention, Microsoft and application vendors sometimes release patches that are buggy and can cause more problems than they solve. Even when updates work as intended, you should ensure that the intentions are what you want. Not long ago, for example, a Microsoft Outlook update was issued that did not allow any attachments to be delivered via e-mail. This did indeed address security concerns about macros and other malicious code embedded in attachments, but it also created a big problem for businesses in which employees relied on e-mailed attachments to perform their work.

It is wise for administrators to closely monitor updates that are released and to choose the ones that will patch machines without causing work stoppages. You should thoroughly test any patches, fixes, and updates in a safe (nonproduction) environment before deploying them on your working network.

To make security updating easier, Microsoft has implemented two useful tools for the Windows Server 2003 family: MBSA and SUS. These two utilities make it much easier to find security holes in your Microsoft software and help you to effectively roll out the patches to client machines from a Windows Server 2003 machine.

Using Microsoft Baseline Security Analyzer (MBSA)

MBSA 1.1.1 is available via download from the Microsoft Web site. This version replaces the stand-alone HFNetChk tool. This tool scans for security vulnerabilities in the operating system software. It also will scan Internet Information Server (IIS), SQL Server, Exchange Server, Windows Media Player, and Internet Explorer for improperly configured security settings. After a scan has completed, a report will be available for each machine that is scanned. To run this utility, you must have local administrative rights.

Installing the Microsoft Baseline Security Analyzer

The Microsoft Baseline Security Analyzer can be downloaded and installed by following these steps:

1. Download the Microsoft Baseline Security Analyzer from the following link:
<http://www.microsoft.com/downloads/details.aspx?FamilyID=9a88e63b-92e3-4f97-80e7-8bc9ff836742&DisplayLang=en>
2. Double click the **Download** button to the right and either **Open** or **Save** the program to your hard drive.
3. If you choose to download the program instead of using the Open command double click on the saved file and the installation will begin.
4. The Welcome Screen wizard will appear and you need to click **Next**.
5. The End User license agreement will appear and you can read the agreement and click on the **I Accept the License Agreement** option then select **Next**.
6. Enter the User Information in the **Full Name** and **Organization** boxes. You can also choose if you wish for the settings to be installed for the current user only or for Anyone who uses this computer by selecting either the **Anyone who uses this computer** option or the **Only for Me** option.
7. Select **Next** and then you have the option to change the destination folder by selecting the **Browse** option or select **Next** to leave as the default install folder.
8. You can also uncheck the check boxes listed if you do not want the following performed:
 - Place a Shortcut on the desktop
 - Show Readme file after installation
 - Launch the application after installation
11. Select **Next** after this has been done.
12. Now you have the ability to install features on your local harddrive and also install all features on the local harddrive by clicking the **Microsoft Baseline Security Analyzer** down arrow select **Next**.
13. Choose **Next** once this has been done and the install will begin. Once it is complete the Microsoft Baseline Security Analyzer will open (if you did not de-select the option in step 8).

If you wish, you can use this tool as a command-line utility. To do so, open a command prompt window (click **Start | Run**, type **cmd**, and click **OK**), and then type **mbsacli.exe** at the command prompt. Running this command will show the tool's syntax and switches. If you run the command without any switches, it will scan the local computer. Some of the switches include the following:

- `/d domainname` Scan a domain.
- `r /c domainname\computername` Scan the computer that is named.
- `/i 10.10.1.1` Scan a particular IP address.
- `/r /i 10.10.1.1- 10.10.1.10` Scan an IP address range.
- `/n OS` Skip operating system commands.
- `/n SQL` Skip SQL Server checks.
- `/n IIS` Skip IIS checks.
- `/n Updates` Skip checking for security updates.
- `/n Password` Skip checking passwords. This option will substantially increase the amount of time it takes for MBSA to run.
- `/e` List errors from the latest scan.
- `/l` List all available reports.

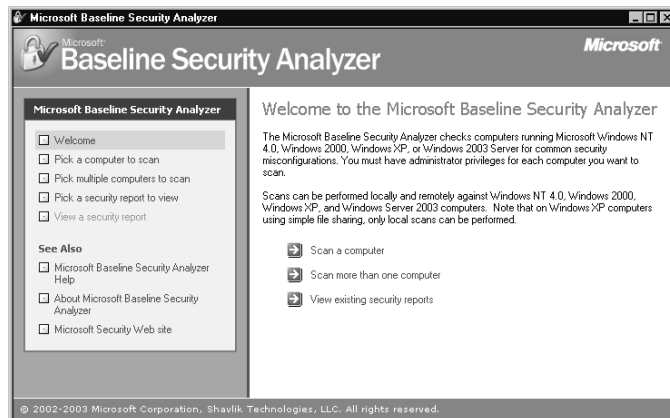
EXERCISE 11.02

SETTING UP A WINDOWS XP CLIENT FOR WIRELESS NETWORKING

Previous versions of MBSA were command-line tools only. Version 1.1.1 provides a graphical interface as well. To use the MBSA GUI after you've installed the program, follow these steps:


1. Select **Start | All Programs | Administrative Tools | Microsoft Baseline Security Analyzer**. You will see the opening window, as shown in Figure 11.5.

Figure 11.5 Starting MBSA




2. You can choose to scan a single computer, or you can scan a group of computers. (When you choose **Scan more than one** computer, you next enter the domain name or IP addresses of a beginning and ending address range.) For this example, we will scan the local computer. Click the **Scan a Computer** button, and you will see the window shown in Figure 11.6.

Figure 11.6 Select a Computer to Scan Using MBSA



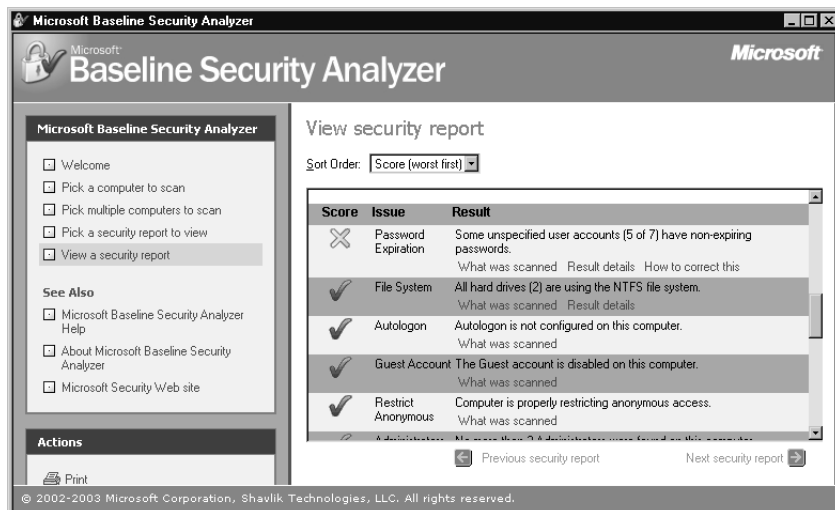
3. Enter the **Computer name** or the **IP address** of the machine you wish to check and select which options you wish to scan. Then click the **Start scan** button, and the scan will begin.
4. When the scan has completed, you will see the window shown in Figure 11.7. As you can see in the figure, the computer name, IP address, security report name, scan date and time, MBSA version, MBSA database version, security assessment, and security update scan results are listed. You can change the sort order of the security report by clicking the **Sort Order** box and selecting **Issue name**, **Score (worst first)**, or **Score (best first)**. **Score (worst first)** is the default setting.

Figure 11.7 The MBSA Output Report on a Local Computer



5. You can scroll down the right side of the page inside the report using the scroll bar to view all of the information. For this example, select the **Password Expiration** result and double-click. The Password Expiration problem (some user accounts have passwords that are set to never expire) is shown in Figure 11.8.

Figure 11.8 A Portion of an MBSA Report Showing the Password Expiration Result



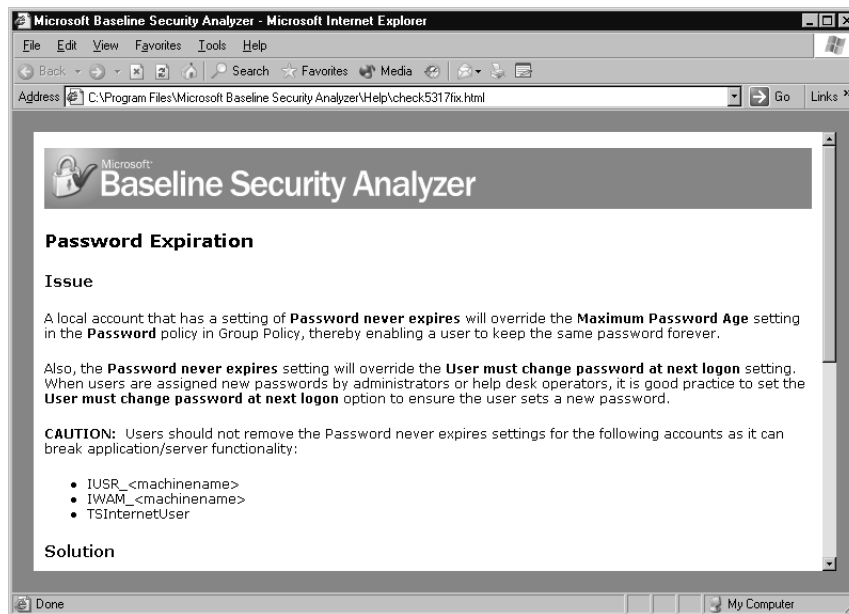
- When you open the item by clicking on the report, you will see details giving you a better understanding of why the problem it captured is a security risk. It will also list a solution to the problem to instruct you on how to correct the risk. Figure 11.9 shows the security issues and action to take to correct the vulnerabilities that were discovered.




NOTE

As always, be careful when you begin to apply security fixes. It is always a good idea to do these types of administrative functions on a test machine in a nonproduction environment

Figure 11.9 Security Issues and How to Correct Vulnerabilities



- Scroll down the window to see a step-by-step guide for correcting your security issues, as shown in Figure 11.10.

Figure 11.10 The MBSA Step-by-Step Solution**NOTE**

MBSA 1.1.1 cannot be used on Windows 9x or Windows ME machines. You can use it on Windows 2000 and Windows XP machines, in addition to Windows Server 2003. However, if a Windows XP machine is using simple file sharing, you can run only local scans. You must have Internet Explorer 5.01 or later installed, and the Workstation and Server services must be enabled. You can remotely scan computers running Windows NT 4.0 Service Pack 4 or above.

Using Microsoft Software Update Services (SUS)

Microsoft SUS is a great add-on product that administrators can use to install software updates from a central location to manage up to 15,000 clients on a single server.

**NOTE**

SUS with Service Pack 1 can be installed on an AD domain controller or a server running Microsoft Small Business Server. The original version of SUS 1.0 could not be installed on these machines.

Microsoft recommends that you use the following guidelines and hardware as a minimum when installing SUS on a machine:

- Pentium III 700 MHz CPU
- 512MB of RAM
- NTFS partition with a minimum of 100MB available for the SUS installation folder
- A minimum of 6GB of disk space available for holding the SUS updates
- NIC and Internet connection
- Windows Server 2003, Windows 2000 Server, Windows 2000 Server Advanced, or Microsoft Windows 2000 Datacenter Server with Service Pack 2 or above
- IIS 5.0
- Internet Explorer 5.5 or newer

SUS works by retrieving updates from Microsoft and storing these updates on a server that has the SUS tool installed. Clients then can be configured to connect to SUS and retrieve approved hotfixes and patches from the SUS server. Administrators have flexibility over the retrieval of the hotfixes and updates because they can choose which languages can be downloaded. Administrators also can approve the hotfixes.

There are two parts required to implement the service:

- **Software Update Service (SUS) server component** Installed on a Windows 2000 or Windows Server 2003 server.
- **Automatic updates** Installed on Windows XP Professional, Windows 2000 Professional and Server with Service Pack 2 or above, and Windows Server 2003 servers; allows them to receive updates from the server running SUS.

The SUS service on the server connects to the Microsoft Windows Update site to download the latest updates. Then the clients connect to the server and receive the updates from it. This prevents numerous clients from needing to use Internet bandwidth to download critical updates, and allows them to receive the update programs faster because they get it over the fast LAN connection instead of a relatively slower Internet connection. The updates need to be downloaded from the Internet only once by the SUS server. Both components are available on the Microsoft Web site.

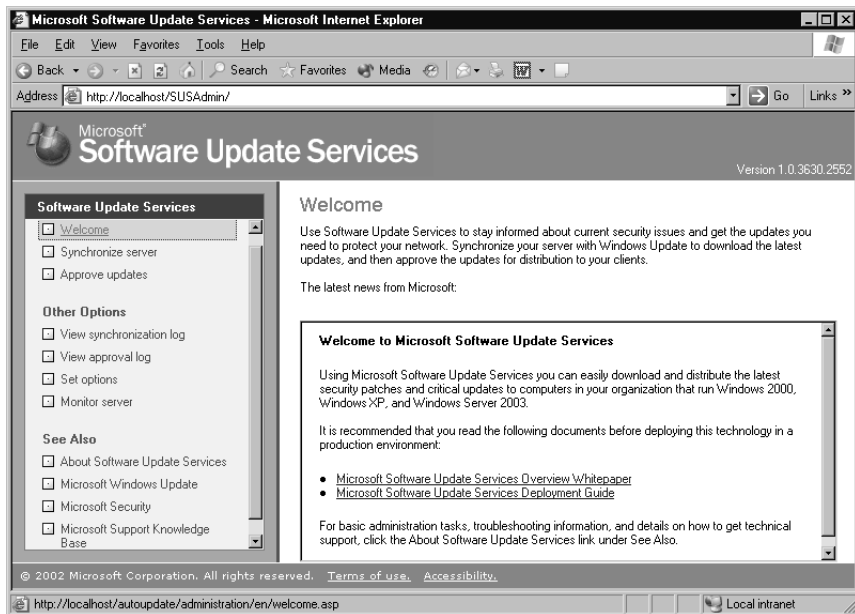
The updates that are managed through SUS include Windows critical updates, critical security updates, and security rollups. Administrators can sign up for e-mail notification of critical updates on the Microsoft SUS Web site.

If you have a server that meets the hardware and software requirements, you can install the SUS software. It is imperative that you make sure you have a virus-free machine, because you must turn off the antivirus software during the installation of SUS. If you leave

your antivirus software running, it might mistake the SUS software installation for a virus. After the installation has completed, you can run SUS as follows:

1. Type **http://computername/susadmin/default.asp** in the address box on your Web browser. This will display the Welcome window, as shown in Figure 11.11.

Figure 11.11 The SUS Welcome Window



2. To get started, you can synchronize your SUS server with the Microsoft Windows Update site by clicking the **Synchronize server** option in the left pane.
3. You can choose to synchronize the server now or you can schedule the server to synchronize at a later time, as shown in Figure 11.12. When you choose to synchronize later (by clicking the Synchronization Schedule button), you see the Schedule Synchronization dialog box, as shown in Figure 11.13. For this example, click the **Synchronize Now** button.

Figure 11.12 The Options for Synchronizing the SUS Server

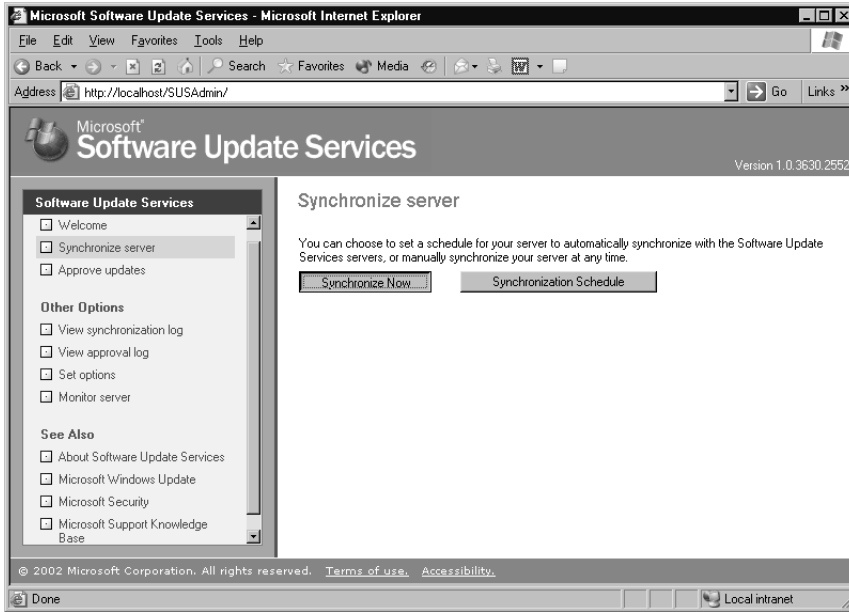
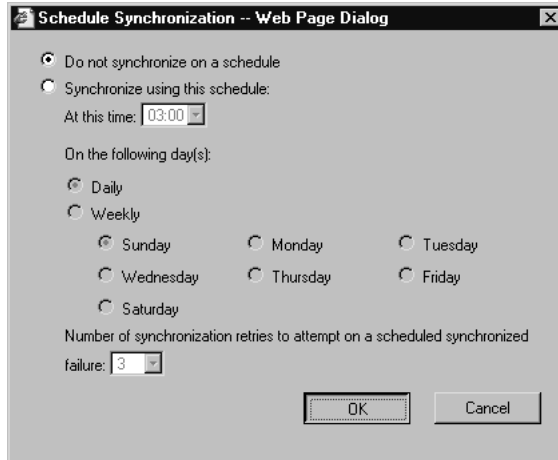
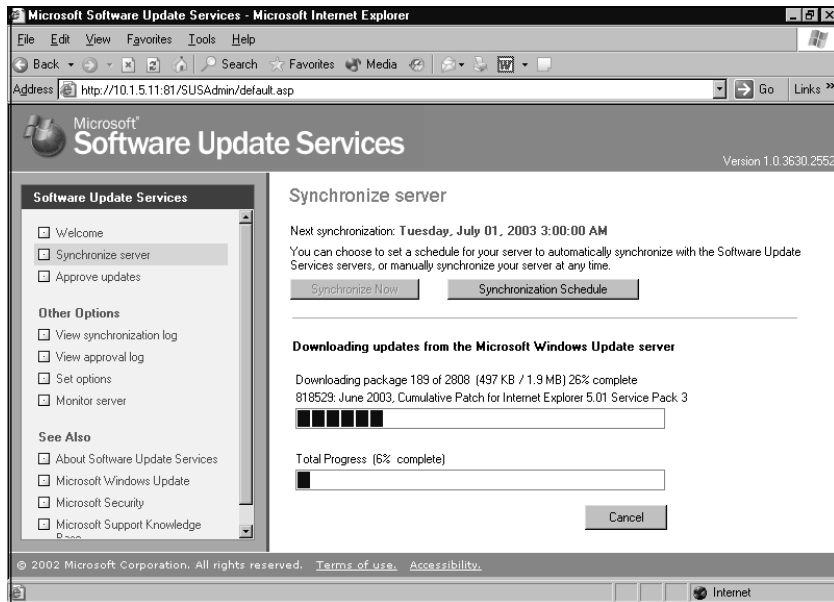


Figure 11.13 Schedule Synchronization for the SUS Server



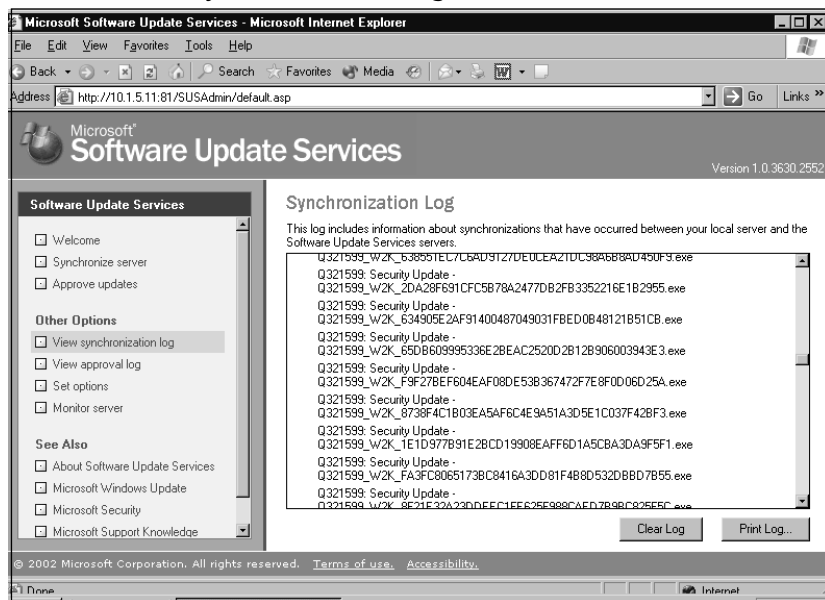
4. After you click **Synchronize Now**, the catalog progress bar will appear, and the download of updates will begin as shown in Figure 11.14. Note that this process can take a great deal of time to complete, since all of the updates are being downloaded onto the server for the first time.

Figure 11.14 Catalog Download Progress Bar



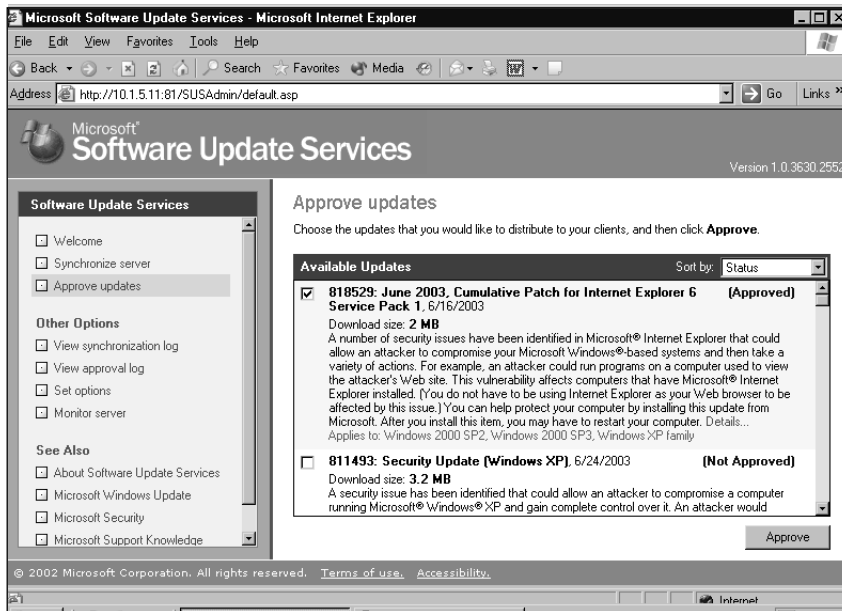
5. After the download has completed, the Synchronization Log will appear as shown in Figure 11.15. You can scroll down the Synchronization Log window and view all of the downloaded updates.

Figure 11.15 The Synchronization Log




- To approve the updates, you can click the **Approve Updates** option in the left panel. Using the **Sort by:** drop-down box, you can sort the updates by status, date, title, or platform. Before approving the updates, ensure that you understand the risk involved with implementing each security fix. Sometimes a fix can cause other problems. As always, make certain that you have a current emergency repair disk (ERD), as well as a backup before you make changes on any machine. You can individually approve the update items you wish to distribute to the clients, as shown in Figure 11.16.

Figure 11.16 Approving SUS Available Updates



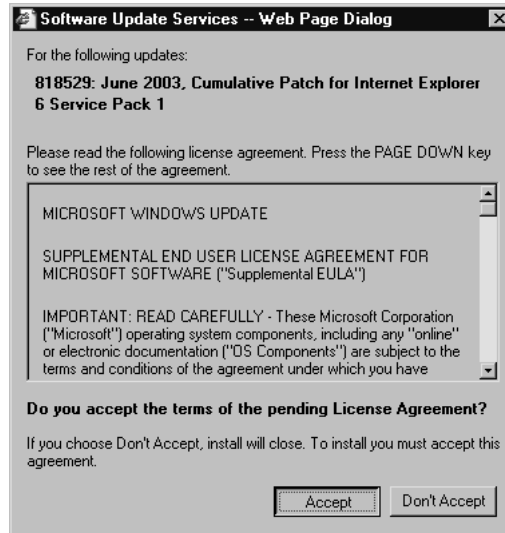
- For this example, we will approve the Internet Explorer 6 Service Pack 1 update. Just place a check mark beside the patch or hotfix and click the **Approve** button. A dialog box will appear, asking if you wish to continue, as shown in Figure 11.17. Click the **Yes** button to continue, and the process of replacing previous updates will begin.

Figure 11.17 SUS Approval Confirmation



- Figure 11.18 shows the End User License Agreement (EULA) that appears before the fix is applied. Click the **Accept** button if you agree to the License Agreement (you must accept the agreement if you wish to distribute the update).

Figure 11.18 SUS License Agreement



- This dialog box shown in Figure 11.19 informs you that your updates have been successfully approved and that they are now available to client machines. Click the **OK** button to close the dialog box.

Figure 11.19 Successful Updates Ready for Client Distribution



In order to set up your clients to use SUS, you need to install the Automatic Updates software, available at <http://www.microsoft.com/windows2000/downloads/recommended/susclient/default.asp>. This software replaces the Critical Updates feature. When the installation of the client software has been completed, you can distribute the updates to client machines as necessary. Clients using Windows XP Professional Service Pack 1, servers running Windows Server 2003, and computers running Windows 2000 Service Pack 3 can be set to automatically receive their SUS updates.

You have three ways to configure clients to retrieve updates from the SUS Server: use the Local Security Policy on a computer, use Group Policy in AD, or edit Registry settings. To use the Local Security Policy on a computer that is not a member of an AD domain, follow these steps:

1. Select **Start | Run**, type **gpedit.msc**, and click **OK**.
2. In the left pane of the GPO Editor, expand **Computer Configuration**, then **Administrative Templates**.
3. Right-click **Administrative Templates** and choose **Add/Remove Templates**.
4. Choose **Add** and select **Wuau.adm** in the folder.
5. You might be prompted to confirm and replace your existing Wuau.adm file; select **Yes** to overwrite this file.
6. Choose **Close** to complete the process.
7. Under **Computer Configuration**, expand **Administrative Templates**, then **Windows Components**, then **Windows Update**.
8. Double-click **Configure Automatic Updates**, and the **Configure Automatic Updates Properties** dialog box will open.
9. Choose **Enabled** from the list of options, and then select any of the following three options:
 - **Notify for download and install** This will provide clients with an icon in the taskbar that will notify users that updates are ready to be installed. To begin the installation, select the icon and the option used to install updates will appear.
 - **Auto Download and notify for install** This is the default method for installing updates. The user will not be notified during the download process. The client machine will automatically find the updates and install them.
 - **Auto download and schedule the install** If this option is chosen and no time is selected, the updates will be automatically installed by the clients at 3:00 A.M. If a reboot is necessary after the installation, the client machine will be rebooted. If a user is working on the machine at 3:00 A.M, the user will be prompted to reboot the machine so the updates can be applied.

You can also choose whether Automatic Updates should restart at a specific time if it was missed, and you can choose whether you want the computer to be restarted by any user using that specific computer.

To configure automatic updates for computers belonging to AD by using Group Policy, follow these steps:

1. Select **Start | All Programs | Administrative Tools | Active Directory Users and Computers**.
2. Right-click the domain or OU to which you wish to apply this setting and select **Properties**.
3. Click the **Group Policy** tab, and then click **New**.
4. Enter a name for the new policy and click the **Edit** button.
5. Under **Computer Settings** or **User Settings**, right-click **Administrative Templates** and choose **Add/Remove Templates**, and then select **Add**.
6. Enter the name of the automatic update file, **wuau.adm**, and click **Open**.

Figure 11.20 shows the Approval Log that can be accessed via the **View approval log** option in the left pane of the SUS administration page. You can choose to clear the log or print the log.

Figure 11.20 Viewing the SUS Approval Log





Figure 11.21 shows the **Set options** page, which can be accessed via by choosing **Set options** in the left pane of the SUS administration page.

Figure 11.21 Setting SUS Options



This screen will allow you to set the following options on your SUS server:

- Proxy configuration
- Name of SUS Web server that you wish to have clients locate
- Server used to synchronize content
- Whether the content should come from Microsoft Windows Update servers or the local SUS server
- How to handle previously approved updates
- Where the updates should be stored
- Language preferences for the downloaded patches



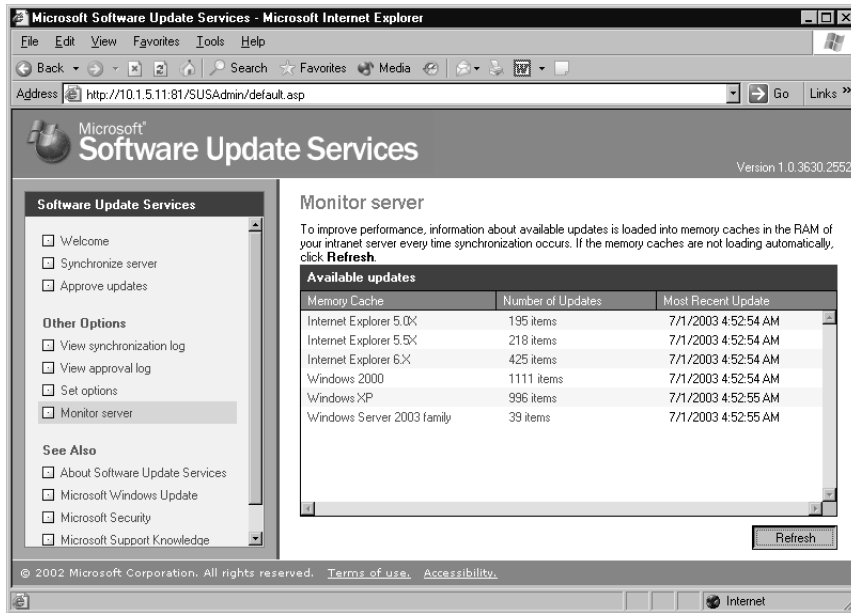
NOTE

You might wish to change the language option if you use only one language on your network, because you do not need all of the other languages to be downloaded. This can greatly reduce your synchronization time.

The SUS **Monitor server** option allows you to view the available SUS updates that are loaded into memory caches of the server each time a synchronization event occurs.

Figure 11.22 shows the **Monitor server** page. You can refresh this page if the cached information is not loading from memory correctly.

Figure 11.22 Monitoring Server Updates



Summary of Exam Objectives

Windows Server 2003 provides administrators with new technology and resources for system, domain, and wireless security. Secure cross-domain and cross-forest trust relationships can be used to allow access to resources between domains in the same or separate forests, using various types of available trusts.

Wireless networking is convenient for users, but it also presents security challenges for network administrators. You can use RADIUS, which controls user authentication for wireless clients with IAS to provide for enhanced connection capabilities in the Windows Server 2003 network and better security for your wireless connections.

Support is provided for numerous IEEE standards to allow wireless clients the ability to connect to the network for data access. Using various protocols such as EAP-TLS, PEAP, and EAP-MS-CHAPv2, security can be greatly enhanced for all types of clients including VPN clients, clients who use certificates, and clients who want to change their passwords after authentication has been successfully completed. Because of the lack of security inherent in the wireless standards, it is a good idea to use data-encryption technologies when connecting to a wireless network. WEP is the basic encryption protocol for wireless networks, but it has many flaws. Although better than no security, WEP should not be relied on as your sole security mechanism if you have any confidential data on the network. New standards and technologies such as 802.1x and 802.11i address some of the problems with WEP and provide solutions designed to make wireless networking more secure.

Monitoring is an important aspect of evaluating and improving your security strategy. You can use the Wireless Monitor to gather information about your wireless clients and WAPs. You can enable auditing of objects and use Group Policy to log object-based access control events on items such as files, folders, services, and Registry information. This can allow you more control over security issues within your organization. Security policies can also be enforced to allow you to control many aspects of network security, including account lockout, password policies, and user rights assignments. Windows Server 2003 includes a variety of built-in utilities that make it easier for you to manage security within your organization. The security configuration and management tools can be used to access and control security policies at the command prompt or by using a management console.

Microsoft has also made it easier for you to keep up with security vulnerabilities and installation of updates and hotfixes with the add-on programs MBSA and SUS. MBSA allows administrators to check for security on client and server machines. SUS allows administrators to provide a central internal location from which clients can retrieve security patches and hotfixes.

Exam Objectives Fast Track

Planning and Implementing Active Directory Security

- ☑ A forest trust is a trust between two Windows Server 2003 forests.
- ☑ Explicit Allow permissions cannot be overridden by inherited Deny permissions.
- ☑ ACLs are used to protect schema objects from unauthorized use in AD.
- ☑ External trusts can be set up using AD domains and trusts for authentication purposes when you do not want to create a transitive forest trust or you need access to Windows NT 4 domains.

Planning and Implementing Wireless Security

- ☑ RADIUS is managed through the Routing and Remote Access (RRAS) MMC.
- ☑ PEAP is used with wireless clients.
- ☑ EAP-MS-CHAP v2 is one of the strongest authentication methods and allows users to change their passwords.
- ☑ WLANs are the most common wireless networks used in corporate and school environments. The 802.11b standard can carry data up to 22 Mbps, and 802.11a and 802.11g can carry data up to 54 Mbps.
- ☑ A LAN port takes on one of two roles during network access, as either an authenticator or supplicant.

Monitoring and Optimizing Security

- ☑ Object-based access control allows administrators to apply audit settings for files, folders, services, and Registry keys.
- ☑ Security policies can be used to control password policies, lockout policies, Kerberos policies, and other aspects of security.
- ☑ Kerberos policies can be used only on domain user accounts.
- ☑ User rights and user permissions are different. Rights are attached to a user or group account and are not object-based. Permissions are attached to specific objects. User rights allow clients to perform specific tasks on their machines or on the network.

Planning a Change and Configuration Management Framework

- ☑ Secedit is used at the command prompt to automate security configuration tasks.
- ☑ Local Security Policy is used to configure security policies on a nondomain controller. These policies apply only to the local machine.
- ☑ Security templates are used to configure security policies according to preset definitions and can be imported into Group Policy.
- ☑ The Security Settings extension to Group Policy is used to configure security on an OU, a site, or a domain.

Planning a Security Update Infrastructure

- ☑ MBSA scans for security vulnerabilities in the operating system and other Microsoft components, including IIS, Exchange Server, SQL Server, Internet Explorer, and Windows Media Player.
- ☑ The command-line program for running MBSA is mbsacli.exe.
- ☑ MBSA gives administrators a report after a scan has been completed. This report explains what security issues were discovered and how to correct them.
- ☑ Microsoft SUS is used to apply security updates from a centralized location within the LAN, giving administrators more control and providing more efficient downloading of updates.

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 have a legacy application that requires anonymous access, and some users cannot access the application. What can I do?
- A:** It is possible that your application requires you to grant access to the Anonymous Users group, which is not part of the Everyone group. If you need to grant access to the Anonymous group, you must explicitly add the Anonymous Logon security group and its permissions.
- Q:** I have multiple domains that need access to resources located in other domains. How can this be set up?
- A:** If users in one domain need access to resources in another domain within the same forest, you do not need to do anything special. This is because, by default, a two-way transitive trust exists between the root domains of every domain tree in the forest so users in any domain in the forest can access resources in any other domain in that forest (if they have the proper permissions). However, to speed up the authentication process between domains, you can create a shortcut trust. If the users in one domain need access to resources in a domain that is in a different forest, you can either create a forest trust between the two forests (which is transitive and will allow all domains in each forest to access all domains in the other) or you can create an external nontransitive trust directly between the two domains.
- Q:** I want to keep my domain Administrator account under wraps for security reasons. What can I do to accomplish this?
- A:** You can disable the built-in Administrator account, since all hackers know the default account name and that is half the information they need to take control of your server. Then you can give administrative privileges to another account. When the Administrator account is disabled, it can still be used in Safe Mode for troubleshooting and repairing problems. Alternatively, you can rename the built-in Administrator account so hackers won't be able to recognize it so easily. You should not log on as Administrator for performing everyday tasks. Instead, use the Run as command when you need to perform administrative tasks.

- Q:** I am trying to audit folder access by a particular user, and I cannot see any information in the event log. What could be the problem?
- A:** Although you can set other types of auditing and they will start immediately, when you want to audit access to objects such as folders, object auditing must be enabled. Then you need to set auditing properties on the object you want to audit (in this case, the folder). To enable object auditing, edit Group Policy for the local computer or the domain policy. In the left pane of the GPO Editor, click Computer Configuration | Windows Settings | Security Settings | Local Policies | Audit Policies and in the right pane, double-click Enable object auditing. Then select to audit successes, failures, or both.
- Q:** I need to apply password policies to all clients. How can I do this?
- A:** Password policies are configured in the Security Settings | Account Policies node of Group Policy on a local or domain GPO. Password policies cannot be set at the site or OU level. You can configure Group Policy to enforce password history, set a maximum and minimum password age, set a minimum password length, enforce complexity requirements, or enable storage of passwords using reversible encryption. The latter should be done only if necessary for compatibility purposes, since it decreases security instead of increasing it.
- Q:** How can I centrally manage security and provide updates for my client machines?
- A:** If client computers are running Windows XP, Windows 2000 Professional or Server, or Windows Server 2003, you can use the Microsoft Baseline Security Analyzer (MBSA) to scan for security problems and use a Microsoft Software Update Services (SUS) server to apply security updates. Both of these tools can be downloaded from the Microsoft Web site. SUS consists of two parts: the SUS server component and the client Automatic Update feature. The SUS server component synchronizes with the Windows Update site and downloads critical updates, security updates, and security rollups to the SUS server. Client machines need the Automatic Update feature installed so they can connect to the SUS server and download the updates that you have approved for distribution.
- Q:** I've just installed a WAP on our company network so employees can roam with their laptops and stay connected to the network (for example, when they attend meetings in conference rooms). Is there anything I need to be aware of in regard to security issues?
- A:** Wireless networking is inherently less secure than traditional wired networks because data is transmitted via radio frequency (RF) signals, which are "out there in the air," vulnerable to capture by anyone who is within range and has the proper equipment. Although you might think "within range" means within the 300 feet or so that wireless manufacturers specify for their devices, a hacker with a high-gain Yagi antenna can connect to your network from much farther away. This situation is exacerbated by the

fact that default settings for most WAPs leave the network wide open, with SSID broadcasting enabled and WEP disabled. Even if you have turned off SSID broadcasting and enabled WEP, that doesn't mean you're safe. A hacker can still use commonly available tools to capture packets sent between legitimate users and determine the SSID from them. Then they can break WEP encryption, which has numerous vulnerabilities, using WEPCrack or other hacker tools. It is best to treat a wireless network as an untrusted network; however, you can make it more secure by using technologies such as 802.1x and 802.11i, by incorporating other mechanisms such as MAC filtering along with WEP, and by implementing secure authentication methods such as RADIUS/IAS and using higher-level protocols such as IPSec to protect wireless traffic.

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 Implementing Active Directory Security

1. You have instituted new security policies for the IT department. One important rule is to never log on as Administrator unless it is absolutely necessary. To enhance security, you want everyone to use their regular user accounts for everyday tasks so you can maintain security as much as possible. A junior administrator comes to you and says he does not wish to log on to the server with an administrative account, but he needs to use a program that requires administrative privileges. What can he do?
 - A. If running the program requires administrative privileges, he cannot run it unless he logs off and logs back on as Administrator.
 - B. He can open the Computer Management console and use the Set password option.
 - C. He can right-click the program he wants to run, select Properties, click the Advanced button, and configure the program to run without administrative privileges.
 - D. He can right-click the program, choose the Run as command, and enter the Administrator account name and password.

2. You have been hired as the network administrator for a small law firm. The first thing you want to do when you take over the job is increase the security on the network. You evaluate the current security level and find it lacking. You decide that you need to secure account passwords using strong encryption on domain controllers. Which utility should you use?
 - A. System Key Utility
 - B. Secedit
 - C. MBSA
 - D. SUS

3. You have recently hired a new junior administrator to assist you in running the network for a medium-sized manufacturing company. You are explaining to your new assistant that AD objects are assigned security descriptors to allow you to implement access control. You tell your assistant that the security descriptor contains several different components. Which of the following are contained in the security descriptor for an object? (Select all that apply.)
 - A. Discretionary access control list
 - B. System access control list
 - C. Dynamic access control list
 - D. Ownership information

4. You are attempting to troubleshoot some problems with access that you think can be traced back to membership in multiple groups. You want to ensure that all administrative accounts are able to perform the tasks they need to accomplish, but you want to remove the built-in accounts from all groups to which they've been added by another administrator, and give them only the access they had by default. You are a little confused because you know that the built-in accounts already belong to some groups at installation, and you don't want to remove them from groups they are supposed to belong to. To which groups does the Domain Administrator account belong in Windows Server 2003 by default? (Select all that apply.)
 - A. Schema Admins
 - B. Enterprise Admins
 - C. Group Policy Creator Owners
 - D. Backup Operators

Planning and Implementing Wireless Security

5. You want to allow wireless clients the ability to change their passwords after they authenticate on the network. Which method of authentication should you implement for these clients?
 - A. EAP-TLS
 - B. EAP
 - C. PEAP
 - D. EAP-MS-CHAP v2

6. You are implementing a new wireless network and need to change the default settings for the equipment on the WLAN. What information should you change? (Select all that apply.)
 - A. SSID password
 - B. SSID network name
 - C. Domain Administrator password
 - D. Domain Administrator account should be renamed

7. You have a number of users who need to be able to roam through the building with their laptop computers and still stay connected to the network. Because of the nature of their work, it is important that they have relatively fast access for transferring a lot of very large data files over the network. You need to implement a wireless network that can connect devices up to 54 Mbps and a minimum of 24 Mbps. Which IEEE standard should you choose?
 - A. 802.15
 - B. 802.11a
 - C. 802.11b
 - D. 802.1x

8. You have hired a consultant to help set up wireless access points on your network. He tells you that you should turn on WEP for the wireless network to help protect it from intruders. You tell him that you have heard that WEP has many flaws and you think additional security measures should be implemented. He assures you that WEP works fine. What do you tell him are some of the problems with WEP?
- A. WEP does not use encryption.
 - B. WEP uses a short (24 bit) initialization vector (IV).
 - C. WEP can use only a 40-bit key.
 - D. WEP uses a public key algorithm.

Monitoring and Optimizing Security

9. Your junior administrator wants to change the name of a user account, but he is worried that if he does so, the user will have problems accessing resources that she had previously been given permissions for. The administrator doesn't want to need to re-create all the group memberships for the newly named account. You tell him there is no need to worry; he can go ahead and change the name, and all the account properties will remain intact. What enables an account to retain its password, profile, group membership, user rights, and membership information?
- A. Group membership of the account
 - B. Domain the account belongs as a member
 - C. Password encryption method
 - D. Security identifier (SID)
10. You suspect that one of your users has been trying to access data in a folder to which he is not supposed to have permission. You are trying to set auditing on this folder so you can see if there are any failed events in the log indicating that the user did try to open the folder. You enable object auditing in the domain's Group Policy Object. However, when you go to add this user to be audited for access to the folder, you find that the folder's property pages do not contain a Security tab. What could be the problem?
- A. Auditing is not set via the Security tab for folders because they don't have such a tab.
 - B. You cannot audit folder access for a particular user.
 - C. The folder is not on an NTFS partition.
 - D. You must share the folder before you can audit it.

Planning a Change and Configuration Management Framework

11. You need to configure Kerberos policies because you want to force user logon restrictions. You go to the computer of the user on whom you want to enforce these policies and access the Local Security Policy. However, in the GPO Editor, you cannot find Kerberos policies in the Security Settings node under Computer Configuration, under Windows Settings. What is the problem?
 - A. You are looking in the wrong section; Kerberos policies are located in the User Configuration node.
 - B. You cannot set Kerberos policies through the Local Security Policy console.
 - C. You must first raise the domain functional level before Kerberos can be used and this option will appear in the GPO.
 - D. Another administrator has deleted the Kerberos policies node from the GPO.

12. You have been analyzing all of your security configuration information as part of a new project that requires you to provide a detailed report on your network's security to management. Toward that end, you need to evaluate the security database test.sdb at the command prompt. What command can you use to do this?
 - A. `secedit /validate test.sdb`
 - B. `secedit /analyze test.sdb`
 - C. `secedit /configure test.sdb`
 - D. `secedit /export test.sdb`

13. You want to set up auditing on several folders that contain important and sensitive information. There are other folders within the specified folders that contain less sensitive information, so you don't want to audit them, because you want to put as little overhead burden on the network as you can. What happens to subfolders and files within a parent folder if auditing has been enabled?
 - A. Subfolders only are audited
 - B. Files only are audited; special access must be turned on for the folders to be audited
 - C. Subfolders and files are audited
 - D. No auditing is performed

14. A parent folder has auditing enabled. Two folders, Applications and Phone Listings, are listed under this parent folder. You need to have the Phone Listings folder audited but not the Applications folder. How can this be accomplished?
- A. It cannot; all subfolders are audited when the parent folder has auditing enabled.
 - B. Right-click the Applications folder, and click the **Properties** tab, select the **Security** tab, and click **Advanced**. Then select the **Auditing** tab and clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Include these with entries explicitly defined here**.
 - C. Right-click the **Phone Listings** folder, click the **Properties** tab, select the **Security** tab, and click **Advanced**. Then select the Auditing tab and clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Audit entries defined here**.
 - D. Right-click the **Phone Listings** folder, click the **Security** tab, and click **Advanced**. Then select the Auditing tab and clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Include these with entries explicitly defined here option**.

Planning a Security Update Infrastructure

15. You need to install the Microsoft Software Update Services (SUS) within your domain to update security information on client computers. What are the minimum requirements that you should use for hardware for the server?
- A. Pentium III, 256MB RAM, NTFS with a minimum of 50MB for the installation folder and 6GB for SUS updates and Active Directory installed
 - B. Pentium III, 512MB RAM, NTFS with a minimum of 100MB for the installation folder and 6GB for SUS updates without Active Directory installed
 - C. Pentium III, 256MB RAM, NTFS with a minimum of 25MB for the installation folder and 6GB for SUS updates without Active Directory installed
 - D. Pentium III, 512MB RAM, NTFS with a minimum of 50MB for the installation folder and 5GB for SUS updates and Active Directory installed

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

MCSE 70-293

Planning, Implementing, and Maintaining a Public Key Infrastructure

Exam Objectives in this Chapter:

- 6 Planning, Implementing, and Maintaining Security Infrastructure.
- 6.2 Plan a public key infrastructure (PKI) that uses Certificate Services.
 - 6.2.1 Identify the appropriate type of certificate authority to support certificate issuance requirements.
 - 6.1 Configure Active Directory directory service for certificate publication.
 - 6.2.2 Plan the enrollment and distribution of certificates.
 - 6.2.3 Plan for the use of smart cards for authentication.
- Summary of Exam Objectives
- Exam Objectives Fast Track
- Exam Objectives Frequently Asked Questions
- Self Test
- Self Test Quick Answer Key

Introduction

Public Key Infrastructure (PKI) is the method of choice for handling authentication issues in large enterprise-level organizations today. Windows Server 2003 includes the tools you need to create a PKI for your company and issue digital certificates to users, computers, and applications. This chapter addresses the complex issues involved in planning a certificate-based PKI. We'll provide an overview of the basic terminology and concepts relating to the public key infrastructure, and you'll learn about public key cryptography and how it is used to authenticate the identity of users, computers, and applications and services. We'll discuss the role of digital certificates and the different types of certificates; user, machine, and application certificates.

You'll learn about certification authorities (CAs), the servers that issue certificates, including both public CAs and private CAs such as the ones you can implement on your own network using Windows Server 2003's certificate services. Next, we'll discuss the CA hierarchy and how root CAs and subordinate CAs act together to provide for your organization's certificate needs. You'll find out how the Microsoft certificate services work, and we'll walk you through the steps involved in implementing one or more certification authorities based on the needs of the organization. You'll learn to determine the appropriate CA type – enterprise or stand-alone CA – for a given situation and how to plan the CA hierarchy and provide for security of your CAs. We'll show you how to plan for enrollment and distribution of certificates, including the use of certificate requests, role-based administration, and auto-enrollment deployment.

Next, we'll discuss how to implement the use of smart cards for authentication within the PKI. You'll learn what smart cards are and how smart card authentication works, and we'll show you how to deploy smart card logon on your network. We'll discuss smart card readers and show you how to set up a smart card enrollment station. Finally, we'll discuss the procedures for using smart cards to log on to Windows, for remote access and VPNs, and to log on to a terminal server.

Planning a Windows Server 2003 Certificate-Based PKI

EXAM
70-293
OBJECTIVE
6
6.2

Computer networks have evolved in recent years to enable an unprecedented sharing of information between individuals, corporations, and even national governments. The need to protect this information has also evolved, and network security has consequently become an essential concern of most system administrators. Even in smaller organizations, the basic goal of preventing unauthorized access while still enabling legitimate information to flow smoothly requires the use of more and more advanced technology.

In the mid-1990s, Microsoft began developing what was to become a comprehensive security system of authentication protocols and technology based on already developed cryptography standards known as Public Key Infrastructure (PKI). With the release of Windows 2000 Server, Microsoft used various existing standards to create the first

Windows-proprietary PKI – one that could be implemented completely without using third-party companies. Windows Server 2003 expands and improves on that original design in several significant ways, which we'll discuss later in this chapter.

Understanding Public Key Infrastructure

To understand how a PKI works, you first need to understand what it is supposed to do. The goals of your infrastructure should include the following:

- Proper authentication
- Trust
- Confidentiality
- Integrity
- Non-repudiation

By using the core PKI elements of public key cryptography, digital signatures, and certificates, all of these equally important goals can be met successfully. The good news is that the majority of the work involved in implementing these elements under Windows Server 2003 is taken care of automatically by the operating system and is done behind the scenes.

The first goal, proper *authentication*, means that you can be highly certain that an entity such as a user or a computer is indeed the entity that he, she, or it is claiming to be. Think of a bank. If you wanted to cash a large check, the teller will more than likely ask for some identification. If you present the teller with a driver's license and the picture on it matches your face, the teller can be highly certain that you are that person – that is, if the teller trusts the validity of the license itself. Because the driver's license is issued by a government agency – a trusted third party – the teller is more likely to accept it as valid proof of your identity than if you presented an employee ID card issued by a small company that the teller has never heard of. As you can see, trust and authentication work hand in hand.

When transferring data across a network, *confidentiality* ensures that the data cannot be viewed and understood by any third party. The data might be anything from an e-mail message to a database of social security numbers. In the past twenty years, more effort has been spent trying to achieve the goal of data confidentiality than perhaps all the others combined. In fact, the entire scientific field of cryptology is devoted to ensuring confidentiality (as well as all the other PKI goals). Cryptology has even claimed a place in Hollywood – the movie *Sneakers* is just one example.



NOTE

Cryptography refers to the process of encrypting data; *cryptanalysis* is the process of decrypting, or "cracking," cryptographic code. Together, the two make up the science of *cryptology*.

As important as confidentiality is, however, the importance of network data *integrity* should not be underestimated. Consider the extreme implications of a patient's medical records being intercepted during transmission and then maliciously or accidentally altered before being sent on to their destination. Integrity gives confidence to a recipient that data has arrived in its original form and hasn't been changed or edited.

Finally, we come to *non-repudiation*. A bit more obscure than the other goals, non-repudiation enables you to prove that a particular entity sent a particular piece of data. It is impossible for the entity to deny having sent it. It becomes extremely difficult for an attacker to masquerade as a legitimate user and then send malevolent data across the network. Non-repudiation is related to, but separate from, authentication.

Public Key Cryptography

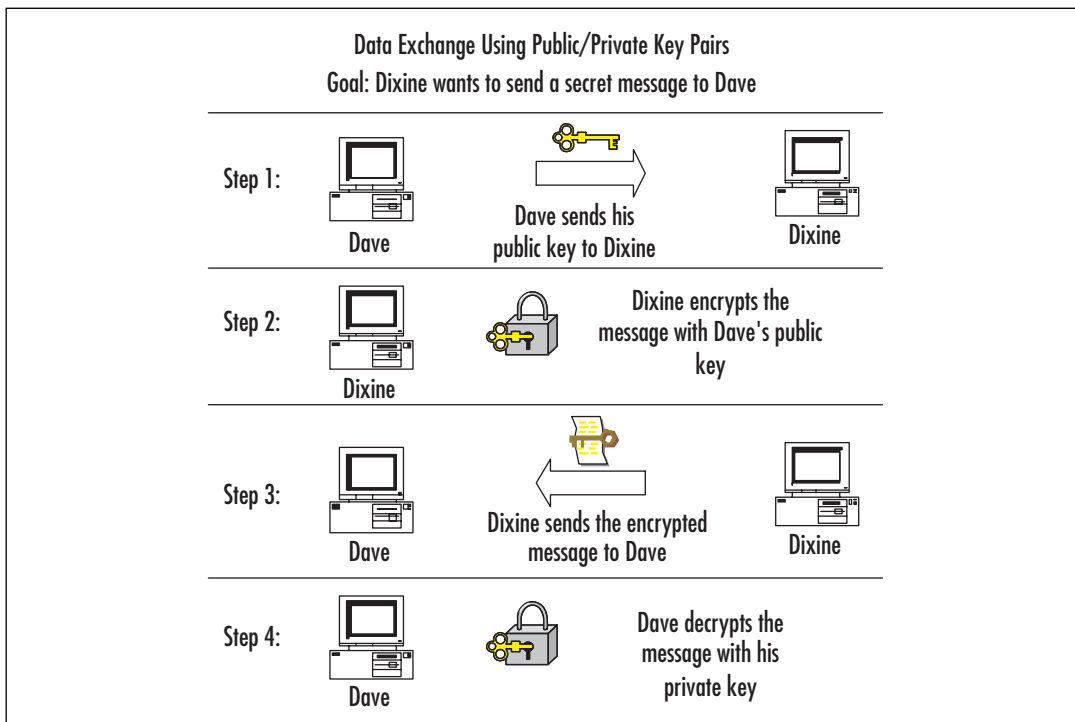
The history of general cryptography almost certainly dates back to almost 2000 B.C. when Roman and Greek statesmen used simple alphabet-shifting algorithms to keep government communication private. Although complexity increased, not much changed until the 1970s, when the National Security Agency (NSA) worked with Dr. Horst Feistel to establish the Data Encryption Standard (DES) and Whitfield Diffie and Martin Hellman introduced the first Public Key Cryptography Standard (PKCS). Windows Server 2003 still uses Diffie-Hellman (DH) algorithms for Secure Sockets Layer (SSL), Transport Layer Security (TLS), and IP Security (IPSec).

DH algorithms are known collectively as *shared secret key* cryptographies, also known as symmetric key encryption. Say you have two users, Greg and Matt, who want to communicate privately. With DH, Greg and Matt each generate a random number. Each of these numbers is known only to the person who generated it. Part one of the DH function changes each secret number into a non-secret, or public, number. Greg and Matt now exchange the public numbers and then enter them into part two of the DH function. This results in a private key – one that is identical to both users. Using advanced mathematics, this shared secret key can be decrypted only by someone with access to one of the original random numbers. As long as Greg and Matt keep the original numbers hidden, the shared secret key cannot be reversed.

Another major force in modern cryptography came about in the late 1970s. RSA Labs, founded by Ronald Rivest, Adi Shamir, and Leonard Adleman, furthered the concept of key cryptography by developing a technology of key pairs, where plaintext that is encrypted by one key can only be decrypted by the other matching key. Windows Server 2003 uses RSA technology in its various forms extensively for such things as Kerberos authentication and S/MIME. The theory goes something like this: Two users, Dave and Dixine, wish to communicate privately. Dave and Dixine each own a key pair consisting of a public key and a private key. If Dave wants Dixine to send him an encrypted message, he first transmits his public key to Dixine. She then uses Dave's public key to encrypt the message. Fundamentally, since Dave's public key was used to encrypt, only Dave's private key can be used to decrypt. When he receives the message, only he is able to read it. Security is main-

tained because only public keys are transmitted – the private keys are kept secret and are known only to their owners. Figure 12.1 illustrates the process.

Figure 12.1 Public/Private Key Data Exchange



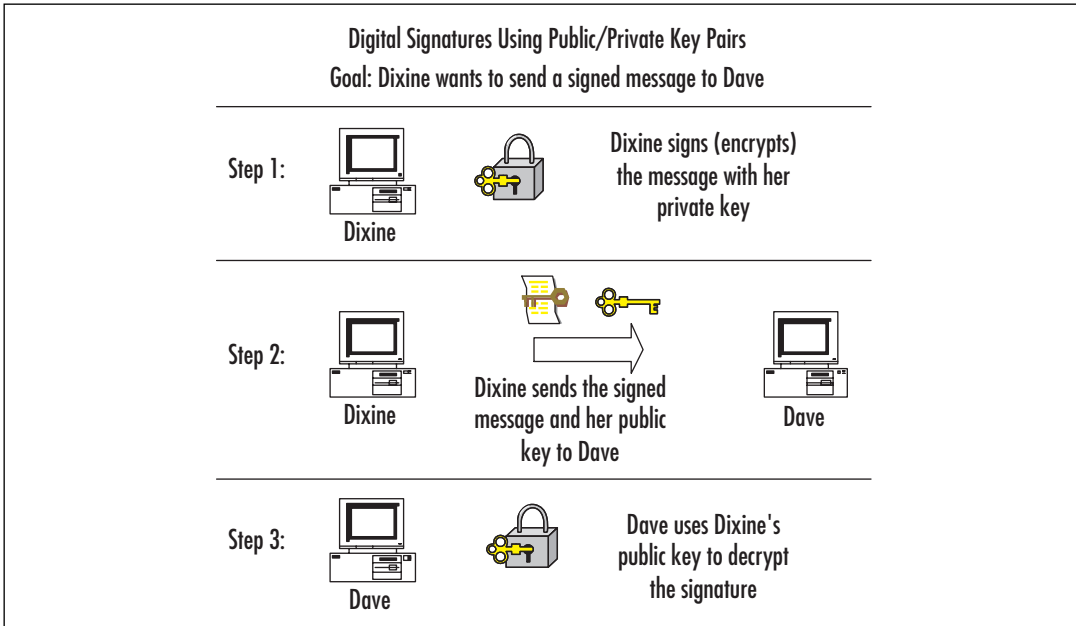
EXAM WARNING

In a Windows Server 2003 PKI, a user's public and private keys are stored under the user's profile. For the administrator, the public keys would be under *Documents and Settings\Administrator\System Certificates\MY\Certificates* and the private keys would be under *Documents and Settings\Administrator\Crypto\RSA* (where they are double encrypted by Microsoft's Data Protection API, or DPAPI). Although a copy of the public keys is kept in the registry, and can even be kept in Active Directory, the private keys are vulnerable to deletion. If you delete a user profile, the private keys will be lost!

RSA can also be used to create "digital signatures" (see Figure 12.2 below). In the communication described above, a public key was used to encrypt a message and the corresponding private key was used to decrypt. If you invert the process, a private key can be used to encrypt and the matching public key to decrypt. This is useful, for example, if you want people to know that a document you wrote is really yours. If you encrypt the docu-

ment using your private key, then only your public key can decrypt it. If people use your public key to read the document and they are successful, they can be certain that it was “signed” by your private key and is therefore authentic.

Figure 12.2 Digital Signatures



Head of the Class... Modern Cryptography 101

Thanks to two mathematical concepts, prime number theory and modulo algebra, most of today's cryptography encryption standards are considered intractable – that is, they are unbreakable with current technology in a reasonable amount of time. For example, it might take 300 linked computers more than 1000 years to decrypt a message. Of course, quantum computing is expected to someday change all that, making calculations exponentially faster and rendering all current cryptographic algorithms useless – but don't worry about that for now.

First, an explanation of the *modulo* operator. Think about elementary school where you first learned to do division. You learned that $19/5$ equals 3 with a remainder of 4. You also probably concentrated on the 3 as the important number. Now, however, you get to look at the remainder. When you take the modulus of two numbers, the result is the remainder; therefore, $19 \bmod 5$ equals 4. Similarly, $24 \bmod 5$ also equals 4 (can you see why?). Finally, you can conclude that 19 and 24 are congruent in modulo 4. So how does this relate to cryptography and prime numbers?

Continued

The idea is to take a message and represent it by using a sequence of numbers. Call the sequence x_i . What you need to do is find three numbers that make the following modulo equation possible: $(x^e)^d \bmod y = x$

The first two numbers, e and d , are a pair and are completely interchangeable. The third number, y , is a product of two very large prime numbers (the larger the primes, the more secure the encryption). Prime number theory is too complex for an in-depth discussion here, but in a nutshell, remember that a prime number is only divisible by the number 1 and itself. This gives each prime number “uniqueness.”

After you have found these numbers (although we won't go into how because this is the really deep mathematical part), the encryption key becomes the pair “ e, y ” and the decryption key becomes the pair “ d, y .” Now it doesn't matter which key you decide to make public and which key you make private, because they're interchangeable. It's a good thing that Windows Server 2003 does all the difficult work for us!

The Function of the PKI

The primary function of the PKI is to address the need for privacy throughout a network. For the administrator, there are many areas that need to be secured. Internal and external authentication, encryption of stored and transmitted files, and e-mail privacy are just a few examples. The infrastructure that Windows Server 2003 provides links many different public key technologies to give the IT administrator the power necessary to maintain a secure network.

Most of the functionality of a Windows Server 2003-based PKI comes from a few crucial components, which are described below. Although there are several third-party vendors, such as VeriSign (www.verisign.com) that offer similar technologies and components, using Windows Server 2003 can be a less-costly and easier-to-implement option – especially for small- and medium-sized companies.

Components of the PKI

Properly planning for and deploying a PKI requires familiarity with a number of components, including but not limited to the following:

- Digital Certificates
- Certification Authorities
- Certificate Enrollment
- Certificate Revocation
- Encryption/Cryptography Services

In the following sections, we will discuss each of these in more detail.

PKI Enhancements in Windows Server 2003

Windows Server 2003 introduces many new enhancements that allow for a more easily implemented PKI solution. The following list items include the major highlights:

- **Auto-enrollment for Users** Windows 2000 first introduced the concept of auto-enrollment for a PKI, but it was limited in scope to machine certificates. Windows Server 2003 now enables the automatic requesting and issuing of user certificates as well.
- **Key Archival and Recovery** Exchange Server 2000 was the first Microsoft product to employ the capability to recover lost keys, but Windows Server 2003 now enables the retrieval of encryption private keys. This eliminates the need to completely reconstruct a user's key pairs.
- **Delta Certificate Revocation Lists (Delta CRLs)** Delta lists enable new additions to a CRL to be published without the need to publish the entire CRL again. Much like an incremental backup in theory, this advancement helps optimize network speed and simplifies the distribution of CRLs.
- **Triple DES and Advanced Encryption Standard (AES) Support** With Windows Server 2003, Microsoft has adopted more components of the standard PKI endorsed by many organizations. The acceptance of 3-DES, or triple DES, in particular has been greatly anticipated by many cryptography experts. AES is still a relatively new standard, but possibly represents the future of encryption.
- **Qualified Subordination** When linking an outside organization's certification authority (CA) structure with your own, trust issues are paramount. New advancements enable the limiting of trust chains and enable the restriction of certificate types acceptable when issued by an external authority.
- **Version 2 Certificate Templates** Windows Server 2003 Enterprise Edition and Windows Server 2003 Datacenter Edition provide many enhancements to the certificate templates found in Windows Server 2003 Standard Edition. Delta CRLs, user certificate auto-enrollment, and key archival/recovery are just some of the important enhancements that version 2 templates have.

Understanding Digital Certificates

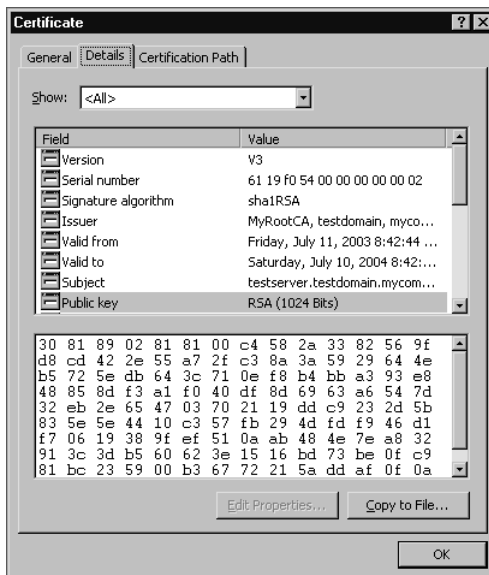
In our previous discussion of public and private key pairs, two users wanted to exchange confidential information and did so by having one user encrypt the data with the other user's public key. We then discussed digital signatures, where the sending user "signs" the

data by using his or her private key. Did you notice the security vulnerability in these methods?

In this type of scenario, there is nothing to prevent an attacker from intercepting the data mid-stream and replacing the original signature with his or her own using of course his or her own private key. The attacker would then forward the replacement public key to the unsuspecting party. In other words, even though the data is signed, how can you be sure of who signed it? The answer in the Windows PKI is the certificate.

Think of a certificate as a small and portable combination safe. The primary purpose of the safe is to hold a public key (although quite a bit of other information is also held there). Someone you trust must hold the combination to the safe – that trust is the basis for the entire PKI system. If I am a user and want to send you my public key so that you can encrypt some data to send back to me, I can just sign the data myself, but I am then vulnerable to the attack mentioned above. However, if I allow a trusted third-party entity to take my public key (which I don't mind because they're trustworthy), lock it away in the safe, and then send the safe to you, you can ask the trusted party for the combination. When you open the safe, you can be certain that the public key and all other information inside really belongs to me, because the safe came from a trustworthy source. The "safe" is really nothing more than a digital signature, except that the signature comes from a universally trusted third party and not from me. The main purpose of certificates, then, is to facilitate the secure transfer of keys across an insecure network. Figure 12.3 shows the properties of a Windows certificate. Notice that the highlighted public key is only part of the certificate.

Figure 12.3 A Windows Server 2003 Certificate





TEST DAY TIP

Certificates are at the very core of the Windows PKI. Make certain that you understand what certificates are, and why they are needed when using public keys. Also, be familiar with the types of certificates listed in this section and the differences between them.

User Certificates

Of the three general types of certificates found in a Windows PKI, the *user certificate* is perhaps the most common. User certificates are certificates that enable the user to do something that would not otherwise be allowed. The Enrollment Agent certificate is one example. Without it, even an administrator is not able to enroll smart cards and configure them properly at an enrollment station. Under Windows Server 2003, required user certificates can be requested automatically by the client and subsequently issued by a certification authority (discussed below) with no user intervention necessary.

Machine Certificates

Also known as computer certificates, *machine certificates* (as the name implies) give the system – instead of the user – the capability to do something out of the ordinary. The main purpose for machine certificates is authentication, both client-side and server-side. As stated earlier, certificates are the main vehicle by which public keys are exchanged in a PKI. Machine certificates are mainly involved with these behind-the-scenes exchanges and are normally overseen by the operating system. Machine certificates have been able to take advantage of Windows' auto-enrollment feature since Windows 2000 Server was introduced. We will discuss auto-enrollment later in this chapter.

Application Certificates

The term *application certificate* refers to any certificate that is used with a specific PKI-enabled application. Examples include IPSec and S/MIME encryption for e-mail. Applications that need certificates are generally configured to automatically request them and are then placed in a waiting status until the required certificate arrives. Depending upon the application, the network administrator or even the user might have the capability to change or even delete certificate requests issued by the application.

Understanding Certification Authorities

Certificates are a way to transfer keys securely across an insecure network. If any arbitrary user were allowed to issue certificates, it would be no different from that user simply signing the data. For a certificate to be of any use, it must be issued by a trusted entity – an entity that both the sender and receiver trust. Such a trusted entity is known as a *certification authority* (CA). Third-party CAs such as VeriSign or Entrust can be trusted because they are

EXAM
70-293
OBJECTIVE
6.2.1

highly visible and their public keys are well known to the IT community. When you are confident that you hold a true public key for a CA, and that public key properly decrypts a certificate, you are then certain that the certificate was digitally signed by the CA and no one else. Only then can you be positive that the public key contained inside the certificate is valid and safe.

In a third-party, or external PKI, it is up to the third-party CA to positively verify the identity of anyone requesting a certificate from it. Beginning with Windows 2000, Microsoft has allowed the creation of a trusted *internal* CA – possibly eliminating the need for an external third party. With a Windows Server 2003 CA, the CA verifies the identity of the user requesting a certificate by checking that user's authentication credentials (using Kerberos or NTLM). If the credentials of the requesting user check out, a certificate is issued to the user. When the user needs to transmit his or her public key to another user or application, the certificate is used to prove to the receiver that the public key inside can be used safely.

In the analogy we used earlier, the state driver's licensing agency is trusted because it is known that the agency requires proof of identity before issuing a driver's license. In the same way, users can trust the certification authority because they know it verifies the authentication credentials before issuing a certificate.

EXAM
70-293
OBJECTIVE
6.2.1

CA Hierarchy

For a very small organization, it might be possible under Windows Server 2003 for you to use only one CA for all PKI functions. However, for larger groups, Microsoft outlines a three-tier hierarchical structure starting at the top with a root CA, moving downward to a mid-level CA, and finally an issuing-level CA. Both the mid-level CA and issuing-level CA are known as subordinate CAs.



EXAM WARNING

Although there are certain advantages to using both external and internal CAs when planning an organization's PKI, you should know that it is possible for a Windows Server 2003 root CA to trust an external root CA, but it is nearly impossible to get the external root CA to trust yours.

The reason is that external CAs are established and highly visible, and therefore easily verifiable to the outside world. Your internal CA is most definitely not. To prove your identity to the external authority, you must jump through a most rigorous set of hoops, and you must also justify the business need for such a relationship. If you go to Microsoft's home Web site at www.microsoft.com and search for the words *CA cross trust*, you will find a white paper entitled *Public Key Interoperability*. This is a good place to start learning more about this complex topic.

Root CAs

When you first set up an internal PKI, no CA exists. The first CA created is known as the root CA, and it can be used to issue certificates to users or to other CAs. As mentioned earlier, in a large organization there usually is a hierarchy where the root CA is not the only certification authority. In this case, the sole purpose of the root CA is to issue certificates to other CAs to establish their authority.

The question then becomes: who issues the root CA a certificate? The answer is that a root CA issues its own certificate (this is called a *self-signed* certificate). Security is not compromised for two reasons. First, you will only implement one root CA in your organization and second, configuring a root CA requires administrative rights on the server. The root CA should be kept highly secured because it has so much authority.

Subordinate CAs

Any certification authority that is established after the root CA is a subordinate CA. Subordinate CAs gain their authority by requesting a certificate from either the root CA or a higher-level subordinate CA. After the subordinate CA receives the certificate, it can control CA policies and/or issue certificates itself, depending on your PKI structure and policies.



TEST DAY TIP

Remember that if a root or subordinate CA becomes compromised (e.g., the server's hard drive is damaged), all CAs subordinate to it will lose their trust relationship and therefore their authority. Always keep current backups of your CAs.

Worse still is the scenario in which a CA's private key is obtained by an attacker. If the CA in question is your root CA, your entire PKI will be compromised.

How Microsoft Certificate Services Works

The Windows Server 2003 PKI does many things behind the scenes. Thanks in part to auto enrollment (discussed later in this chapter) and certificate stores (places where certificates are kept after their creation), some PKI-enabled features such as EFS work with no user intervention at all. Others, such as IPSec, require significantly less work than would be required without an advanced operating system.

Even though a majority of the PKI is handled by Windows Server 2003, it is still instructive to have an overview of how certificate services work.

1. First, a system or user generates a public/private key pair and then a certificate request.

2. The certificate request, which contains the public key and other identifying information such as user name, is forwarded to a CA.
3. The CA verifies the validity of the public key. If it is verified, the CA issues the certificate.
4. After it is issued, the certificate is ready for use and is kept in the certificate store, which can reside in Active Directory. Applications that require a certificate use this central repository when necessary.

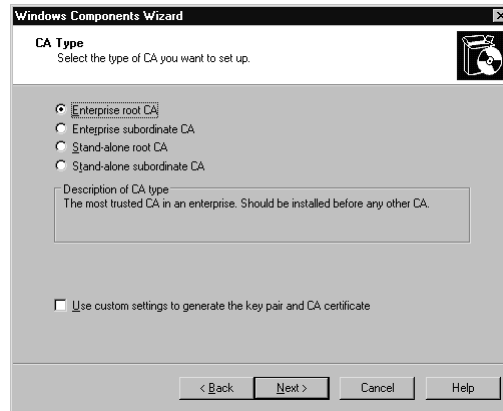
In practice, it isn't terribly difficult to implement certificate services, as the following exercise shows. Configuring the CA requires a bit more effort, as does planning the structure and hierarchy of the PKI – especially if you are designing an enterprise-wide solution. We'll cover these topics later in this chapter.

EXERCISE 12.01

INSTALLING CERTIFICATE SERVICES


1. After logging on with administrative privileges, click **Start | Control Panel**, and then click **Add/Remove Programs**.
2. Click **Add/Remove Windows Components**, and then check **Certificate Services**. This selects both sub-components of certificate services, which are Certificate Services CA and Certificate Services Web Enrollment Support. If Web enrollment support is *not* checked, you will not be able to complete Exercise 12.03.
3. A warning dialog box appear telling you that after certificate services have been installed you will not be able to change the machine's domain membership or change its computer name. Click **Yes** to continue.
4. You now must choose the type of CA to establish, as seen in Figure 12.4. You have two decisions to make – that of root vs. subordinate and enterprise vs. standalone (discussed later in this chapter). For this exercise, click **Enterprise root CA** and click **Next**. If you checked the **Use custom settings to generate the key pair and CA certificate**, you would be prompted to choose a custom cryptographic service provider (CSP), a hash algorithm, and a key length. You could also elect to use an existing key or to use an imported one.

Figure 12.4 Choosing the CA Type



5. The next dialog box presented is the **CA Identifying Information** box. See Figure 12.5. Enter a common name for the CA. For this exercise, type **My Root CA**. The distinguished name suffix is provided by the operating system and is used along with the common name you just typed in to form the distinguished name. Note that you can also change the default five-year validity period of the CA. You can set the validity period as a number of days, weeks, months, or years. Click **Next** to continue.

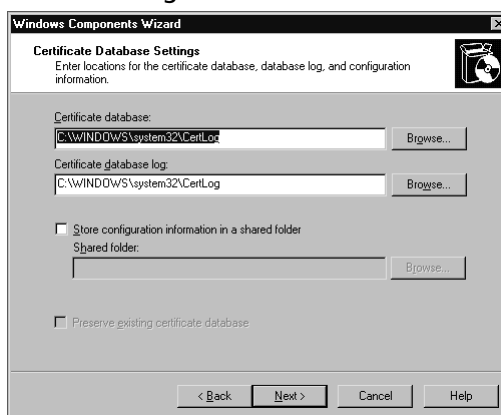
Figure 12.5 Naming the CA



6. After the key pair is generated, the Certificate Database Settings dialog box appears. As in Figure 12.6, you will notice that both the certificate database and certificate database log textboxes are already filled with default values. You may elect to **Store configuration information in a**

shared folder, but do not check it for purposes of this exercise. Click **Next** to complete the installation. After Windows Server 2003 has completed its work (you might be notified during this process that the Internet Information Service (IIS) will stop if you have IIS running on this machine), click **Finish**. During the configuration process, you might be prompted to insert your Windows Server 2003 installation CD or enter the path to the installation files on the hard disk or on a network share. You will also be notified that Active Server Pages (ASP) must be enabled in IIS to provide Web enrollment services. Click **Yes** to enable ASP.

Figure 12.6 Selecting the Certificate Database Location



EXAM WARNING

Pay special attention to the warning given in step 3 in the above exercise. Because the distinguished name of the CA is a part of the certificates it issues, renaming the server or removing it from the domain is not allowed. Windows Server 2003 uses the X.500 standard for distinguished names.

EXAM
70-293
OBJECTIVE
6
6.1
6.2.1

Implementing Certification Authorities

Planning a PKI structure that includes multiple CAs in a hierarchy with proper security can be a test in patience and fortitude. The actual implementation, however, is relatively simple. In Exercise 12.01, you installed certificate services and chose to create an enterprise root

CA. That's pretty much it for the implementation of a CA, but there is much more involved in the configuration of process. Before we talk about the differences between enterprise and stand-alone CAs, and the security concerns involved, we'll go over the many options you have control over in the following exercise.

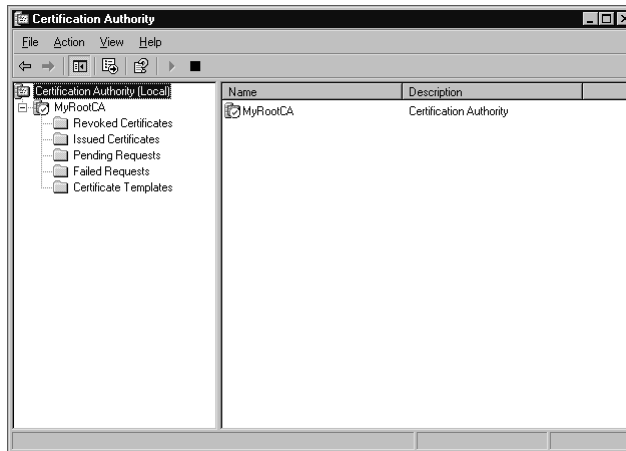
EXERCISE 12.02

CONFIGURING A CERTIFICATION AUTHORITY


In this exercise, we'll explore the different properties you have control over when configuring a CA. We won't go over all the options now, but we will cover them all in this chapter.

1. Click **Start | Administrative Tools | Certification Authority** (note that certificate services must be installed before this step – see Exercise 12.01; otherwise, this choice will not appear in the **Administrative Tools** menu).
2. In the left pane of the **Certification Authority** snap-in, click **My Root CA** (or whichever CA name you have listed) and expand it. As Figure 12.7 shows, this is where you can view revoked and issued certificates, pending and failed certificates, and certificate templates (discussed later in this chapter).

Figure 12.7 The Certification Authority Snap-In



3. Highlight **My Root CA** and right-click it. Click **Properties**. Figure 12.8 shows the **General** tab. Here, all installed CA certificates are listed as well as the CSP and hash algorithms used. Click **View Certificate** if you want to see the certificate itself.

Figure 12.8 General Tab of the CA Property Sheet


4. Click the **Policy Module** tab. A policy module defines how the CA handles incoming certificate requests. Notice in Figure 12.9 that the Windows default policy is listed. The **Select** button is used to choose a different policy module, usually a customized version. Click the **Properties** button (see Figure 12.10). The default setting tells the CA to follow the settings in the certificate template if applicable and to automatically issue the certificate otherwise. The other setting tags all incoming requests to *pending* status, forcing the administrator to manually approve or deny each certificate request. Keep the default setting and click **OK** to return to the CA property sheet.


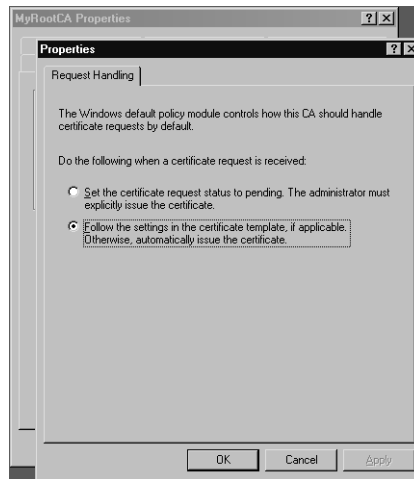
Figure 12.9 Policy Module Tab of the CA Property Sheet


Figure 12.10 Request Handling Tab of the Default Policy Module



5. Click the **Exit Module** tab. Whereas a policy module defines how a CA handles incoming certificate requests, an exit module defines what a CA does with certificates that it issues. Figure 12.11 shows that the Windows default policy is listed. In addition to the **Add** and **Remove** buttons, there is a **Properties** button. Click the **Properties** button. Figure 12.12 shows that the only setting is to allow certificates to be published to the file system if a certificate template dictates, which the default policy does not allow. Again, keep the default setting and click **OK** to return to the CA properties sheet. Skip the **Extensions** tab for now; we'll discuss it when we talk about certificate revocations later in the chapter.

Figure 12.11 Exit Module Tab of the CA Property Sheet

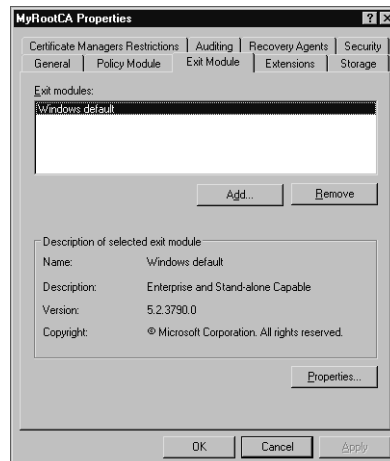
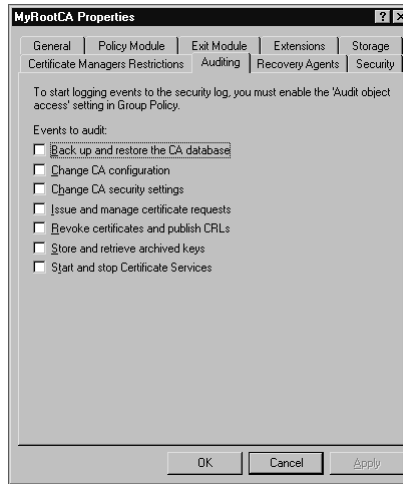


Figure 12.12 Publication Settings Tab of the Default Exit Module

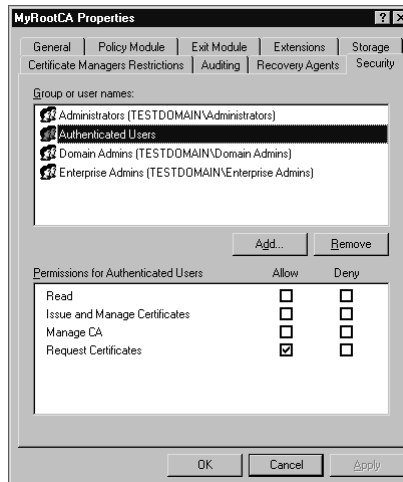
6. Click the **Storage** tab. Note that the default settings cannot be changed because Active Directory is being used. We'll discuss more about the relationship between Active Directory and enterprise CAs later in this chapter.
7. Click the **Certificate Managers Restrictions** tab. The default setting here tells the CA to not restrict certificate managers. As an administrator, you can designate certificate managers by giving them the **Issue and Manage Certificates** permission. By changing the default, you can specifically restrict the users, groups, and computers over which a certificate manager has control.
8. Click the **Auditing** tab. As seen in Figure 12.13, there are many events that you can monitor – each concerned with a different aspect of security. Especially important are the **Change CA configuration**, **Change CA security settings**, and **Issue and manage certificate requests** events. Skip the Recovery Agents tab; we'll cover it during our discussion of key archival and recovery.

Figure 12.13 Auditing Tab of the CA Property Sheet



- Click the **Security** tab. The **Security** tab, shown in Figure 12.14, enables you to grant or deny access to users over several key areas of the CA. Note that the **Issue and Manage Certificates** permission denotes a certificate manager, whereas the **Manage CA** permission gives authoritative access to the entire CA. Click **Cancel** to return to the CA snap-in.

Figure 12.14 Security Tab of the CA Property Sheet



Analyzing Certificate Needs within the Organization

You've just concluded a tour of most of the properties associated with a CA, but knowing what you *can* do does not mean that you know what you *should* do. To find out more about what you should do, you need to analyze the certificate needs of your organization and then move on to create an appropriate CA structure.

According to Microsoft's TechNet, the analysis of certificate needs springs primarily from "the analysis of business requirements and the analysis of applications that benefit from PKI-based security." In other words, when designing a PKI/CA structure, you need to understand the different uses for certificates and whether your organization needs to use certificates for each of these purposes. Examples include SSL for a secure Web server, EFS for encryption of files, and S/MIME for encryption of e-mail messages. The use of S/MIME might dictate that your CA hierarchy has a trust relationship with external CAs, and the use of SSL might lead you to implement a stand-alone CA instead of an enterprise CA. Thus, analyzing these needs *before* you implement your PKI can save you a lot of time and trouble.

Determining Appropriate CA Type(s)

For most administrators, the most significant factor in designing a CA structure is the amount of PKI-related traffic on the network. If you run a small organization without an Internet presence, for example, a single-root CA that issues certificates directly to users will probably fit the bill. However, in a larger organization, a CA hierarchy is likely to be more appropriate.

The first choice when determining appropriate CA types for your PKI is how many subordination levels to use. One level, the root, is required. Two, three, and even four subordination levels are relatively common, but the three-tier model is the one most referenced and most-frequently used. So how does the three-tier model work? We've discussed previously the differences between a root CA and a subordinate CA, and that a root CA issues certificates to the second-tier subordinates. In the standard three-tier model, the root CA has the job of issuing certificates to the second-tier. That's all it really does. Certainly it has the capability of doing more – it could even issue certificates to users. However in a large company, the amount of traffic generated by even a few PKI-aware applications could easily overwhelm a single CA. Also, if you shift the responsibility of issuing certificates to subordinate CAs, you can take the root CA *offline* – meaning that you detach it from the network entirely. This provides a very high level of security, because attackers have no way of getting to the machine. When a subordinate CA requires a certificate from the root, you can either briefly connect the root CA to the network and then remove it again, or you can literally use a floppy disk.

The intermediate level of CAs, the one just below the root, has the responsibility for controlling certificate policy and issuing certificates to the bottom-level CAs. These

bottom-level CAs are the ones that actually issue certificates to users, machines, and applications. The question then becomes: why don't the intermediate CAs just issue the user certificates directly? The answer is that although they can, it just isn't as scalable as the three-tier model. It is easier to add CAs to the hierarchy that are concerned only with issuing certificates and not involved with policies such as key length and CSP choice.

After you have determined the hierarchical structure of your CAs, you will need to determine which CAs are set up as enterprise CAs and which ones are set up as stand-alone CAs before implementing them. You may recall that in Exercise 12.01 you installed certificate services and chose an enterprise root CA. The choice in your network will depend on several different factors, such as your needed level of security. Both enterprise and stand-alone CAs have advantages and disadvantages. We'll explore some of them in the following sections.

Enterprise CAs

An enterprise CA is tied into Active Directory (AD) and is required to use it. In fact, a copy of its own CA certificate itself is stored in Active Directory. Perhaps the biggest difference between an enterprise CA and a stand-alone CA is that enterprise CAs use Kerberos or NTLM authentication to validate users and computers before certificates are issued. This provides additional security to the PKI because the validation process relies on the strength of the Kerberos protocol and not a human administrator. Enterprise CAs also use templates, which are described later in this chapter, and they can issue every type of certificate.

There are also several downsides to an enterprise CA. In comparison to a stand-alone CA, enterprise CAs are more difficult to maintain and require a much more in-depth knowledge about Active Directory and authentication. Also, because an enterprise CA requires Active Directory, it is nearly impossible to remove it from the network. If you were to do so, the Directory itself would quickly become outdated – making it difficult to resynchronize with the rest of the network when brought back online. This forces an enterprise CA to remain attached to the network, leaving it vulnerable to attackers.

Stand-Alone CAs

Stand-alone CAs do not require Active Directory (although they *can* use AD information if it is available), and are usually used as either secure root CAs or as an issuer to such applications as stand-alone Web servers. Stand-alone CAs are generally not suitable for enterprise-type applications. Because certificate templates are not used on a stand-alone CA, a standalone is more basic and easier to maintain than an enterprise CA. A stand-alone CA keeps a copy of its CA certificate in a shared folder and if Active Directory is not used, users that need to request certificates need to know the location of the CA. Finally, stand-alone servers can be secured by removing them from the network.

The disadvantages to a stand-alone CA are that an administrator must manually approve or deny every certificate request individually, a stand-alone CA cannot issue log-on certificates, and templates cannot be used with a stand-alone CA, so a key recovery agent cannot be established (we discuss the key recovery agent template below).



EXAM WARNING

A stand-alone CA does not need Active Directory as an enterprise CA does. For test day, remember that without Active Directory, all certificate requests made to a *stand-alone* CA are automatically tagged as *pending*. This means that automatic fulfillment is not available and an administrator must manually approve or deny each incoming request. Under Windows Server 2003, stand-alone CAs *can* be configured to accept requests automatically, but that is not the default setting.

Planning the CA Hierarchy

There is more than meets the eye when planning a CA hierarchy. We've already discussed choices you will need to make between root and subordinate and between enterprise and standalone. You will also need to consider possible cross-trust hierarchies and the establishment of the key recovery agent.

Cross-Trust Hierarchies

For a PKI entity to use a certificate provided by a CA, the entity must trust that CA. This trust is established when the entity has a copy of the CA's certificate located in its local certificate store. Using the public key contained in the certificate, the entity can verify the CA's digital signature. How, then, does the certificate get from the CA to the entity's local store? Unfortunately, there is not just one answer. Group policies under Active Directory, preloaded certificates in Windows Server 2003, and downloads from the Windows Update Web site are the most common ways.

The chain of trust from an issuing CA all the way up to the root CA must be verified by an entity requesting a certificate for the certificate to be accepted. In a small, local network operation this is easy to accomplish. However, when your organization must exchange data with external parties, there needs to be a way to recognize and trust a third-party CA as if it were a part of your local chain of trust. There are two ways to do this:

- You can use a certificate trust list, or CTL.
- You can create a cross-trust hierarchy, which enables an external CA to be viewed as a subordinate CA in your local trust chain.

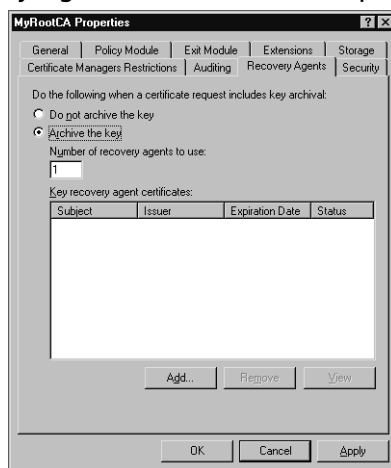
Using a CTL or a cross-trust hierarchy under previous versions of Windows presented a central problem. When an external CA gained trust status, every certificate issued by it and all of its subordinate CAs were automatically trusted. New to Windows Server 2003 is a feature called *qualified subordination*. Qualified subordination enables you to specify how many subordinates can be trusted, and it also enables you to specify the purposes of certificates that can be accepted from the external CAs.

Key Recovery Agent

As when a person has locked his or her keys inside the car, lost encryption keys in a PKI can be troublesome. Luckily, Windows Server 2003 provides a locksmith of sorts (called a Registration Authority, or RA) that earlier versions of Windows did not have. A key recovery solution, however, is not easy to implement and requires several steps. The basic method follows:

1. Create an account to be used for key recovery.
2. Create a new template to issue to that account.
3. Request a key recovery certificate from the CA.
4. Have the CA issue the certificate.
5. Configure the CA to archive certificates by using the **Recovery Agents** tab of the CA property sheet (shown in Figure 12.15).
6. Create an archive template for the CA.

Figure 12.15 Recovery Agents Tab of the CA Property Sheet



Each of these steps requires many substeps, but can be well worth the time and effort. It is worth noting again that key recovery is not possible on a stand-alone CA, because a stand-alone cannot use templates. It is also worth noting that only encryption keys can be recovered—private keys used for digital signatures cannot be.



TEST DAY TIP

Key archival and recovery rely on a version 2 template, which is only available in Windows Server 2003 Enterprise or datacenter Editions. If you're using Windows Server 2003 Standard Edition, the **Recovery Agents** tab won't even be visible because the Standard Edition only supports version 1 templates.

Planning CA Security

The two fundamentals of CA security are to guard the CA hierarchy from attackers and to configure the hierarchy for disaster recovery. The first of these requires a good deal of planning. For starters, you need to know the physical and logical location of the root CA. For extreme security, the CA can be physically located in a locked closet with a lights-out configuration (“lights out” refers to a server that has neither a monitor nor a keyboard attached). In most cases, however, lights out would be appropriate only after the entire PKI is set up. Remember that you will need to use the root CA every time a subordinate CA needs to request a certificate.

As we have already discussed, configuring the root CA as a standalone is probably the most important measure you can take to prevent accidental or intentional tampering. With no network connectivity, attacks become virtually impossible, since a user would have to log on while sitting at the physical location of the server. Other security considerations are really more a function of general server security—things such as requiring complex passwords and implementing file encryption.

In guarding the hierarchy, you cannot solely concentrate on the root CA. After all, if a subordinate CA is tampered with, every entity below it in the PKI hierarchy becomes compromised. Most subordinate CAs are attached to the network. This obviously increases their vulnerability. Beyond securing the network itself (by using IPSec and group policies, for example), there is another part of a standard PKI that helps maintain CA integrity. That part is *certificate revocation*, which we will go into in greater detail shortly. Certificate revocation enables an administrator to warn PKI clients about certificates that might not be authentic or that might have been issued by a rogue CA.

Disaster recovery applies to every CA in the hierarchy, but especially at the root. That being said, the importance of performing proper backups cannot be overstated. A periodic full backup (for example, weekly) with more frequent incremental backups (for example, daily) is recommended. For your organization, configuring the hierarchy for a disaster may also include installing additional CAs that are responsible for more narrow responsibilities. For example, you might want one CA to issue smart card certificates and nothing more. That way, if the CA is lost, it is not as difficult to replace. Finally, remember Windows Server 2003’s capability to archive and recover keys.



TEST DAY TIP


As mentioned above, the first concern of PKI security is keeping the root CA secure. Because Microsoft recommends that you configure the root CA as offline and standalone, the machine should not be a domain controller. Domain controllers need to be available for replication and cannot be offline a majority of the time.

Certificate Revocation

A CA's primary duty is to issue certificates, either to subordinate CAs or to PKI clients. However, each CA also has the capability to revoke those certificates when necessary. The tool that the CA uses for revocation is the *certificate revocation list*, or CRL. The act of revoking a certificate is simple: from the **Certification Authority** console, simply highlight the **Issued Certificates** container, right-click the certificate and choose **All | Revoke Certificate**. The certificate will then be located in the **Revoked Certificates** container.

When a PKI entity verifies a certificate's validity, that entity checks the CRL before giving approval. The question is: how does a client know where to check for the list? The answer is the CDPs, or CRL Distribution Points. CDPs are locations on the network to which a CA publishes the CRL. In the case of an enterprise CA under Windows Server 2003, Active Directory holds the CRL and for a standalone, the CRL is located in the *certsrv\certenroll* directory. Each certificate has a location listed for the CDP, when the client views the certificate, it then understands where to go for the latest CRL. Figure 12.16 shows the Extensions tab of the CA property sheet, where you can modify the location of the CDP.

Figure 12.16 Extensions Tab of the CA Property Sheet



For a CA to publish a CRL, use the **Certification Authority** console to right-click the **Revoked Certificates** container and choose **All Tasks | Publish**. From there, you can choose to publish either a complete CRL or a Delta CRL.



NOTE

Delta CRLs are new to Windows Server 2003. They enable a CA to publish only changes made to the original CRL. Since they are much smaller than the entire CRL, network traffic is minimized.

Whether you select a New CRL or a Delta CRL, you are next prompted to enter a publication interval (the most frequent intervals chosen are one week for full CRLs and one day for Delta CRLs). Clients cache the CRL for this period of time and then check the CDP again when the period expires. If an updated CDP does not exist or cannot be located, the client automatically assumes that all certificates are invalid.

EXAM
70-293
OBJECTIVE
6
6.1
6.2.2

Planning Enrollment and Distribution of Certificates

For a PKI client to use a certificate, two basic things must happen. First, a CA has to make the certificate available and Second, the client has to request the certificate. Only after these first steps can the CA issue the certificate or deny the request. Making the certificate available is done through the use of certificate templates and is a topic that we discuss in detail section. As for the client, there are three methods of requesting certificates – all three of which are essential to a thorough understanding of PKI:


- Auto-enrollment
- The Certificates snap-in
- The Certificates Web page

We will discuss each in more detail in the section titled *Certificate Requests*.

Certificate Templates

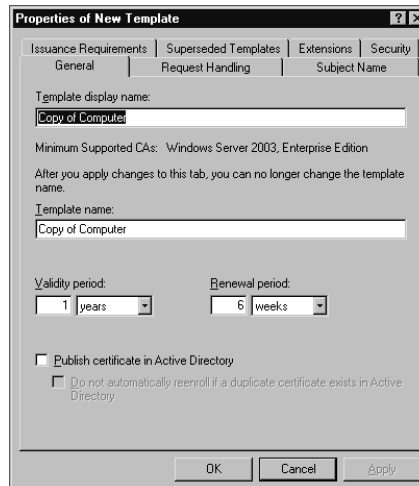
A *certificate template* defines the policies and rules that a CA uses when a request for a certificate is received. Many built-in templates can be viewed using the **Certificate Templates** snap-in (see Figure 12.17). The snap-in can be run by right-clicking the **Certificate Templates** container located in the **Certification Authority** console (described in Exercise 12.02) and clicking **Manage**. You can use one of the built-in templates or create your own.

Figure 12.17 Certificate Templates Snap-In



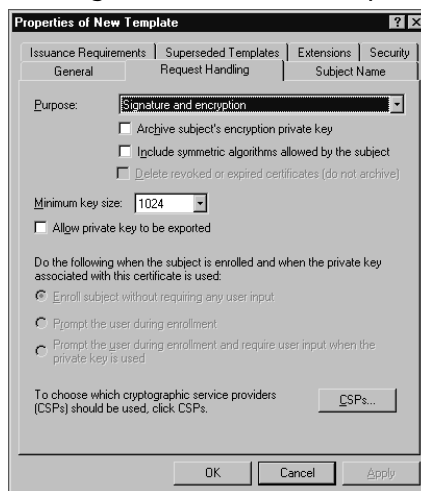
When creating your own template, you have multiple options that will guide the CA in how to handle incoming requests. The first step in the creation process is to duplicate an existing template. You do this by using the **Certificate Templates** snap-in, then right-clicking the template you wish to copy and selecting *Duplicate Template*. On the **General** tab that appears by default (seen in Figure 12.18), there are time-sensitive options such as validity period and renewal period. Note the default validity period of one year and the default renewal period of six weeks. There are also general options such as the template display name and a check box for publishing the certificate in Active Directory.

Figure 12.18 General Tab of the New Template Property Sheet



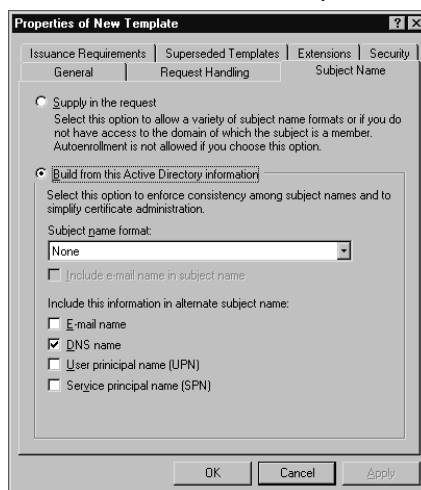
The **Request Handling** tab, shown in Figure 12.19, has options including minimum key size and certificate purpose. The certificate purpose can be encryption, signature, or signature and encryption. There is also an option to allow the export of the private key. Finally, you can instruct the CA how to act when the subject's request is received and which CSPs to use.

Figure 12.19 Request Handling Tab of the New Template Property Sheet



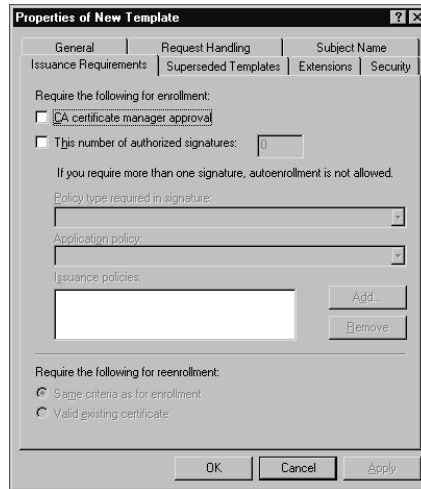
The **Subject Name** tab seen in Figure 12.20 gives you the choice of obtaining subject name information from Active Directory or from the certificate request itself. In the latter case, auto-enrollment (which we'll discuss later in the chapter) is not available.

Figure 12.20 Subject Name Tab of the New Template Property Sheet



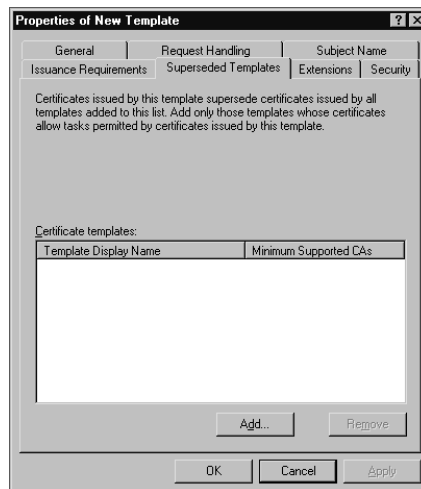
The **Issuance Requirements** tab seen in Figure 12.21 enables you to suspend automatic certificate issuance by selecting the CA certificate manager approval check box.

Figure 12.21 Issuance Requirements Tab of the New Template Property Sheet



The **Superseded Templates** tab, as shown in Figure 12.22, is used to define which certificates the current template supersedes. Usually, this tab is used to configure a template that serves several functions; e.g., IPSec and EFS. In this case, a template used *only* for IPSec or a template used *only* for EFS would be placed on the superseded templates list.


Figure 12.22 Superseded Templates Tab of the New Template Property Sheet



The **Extensions** tab, as seen in Figure 12.23, can be used to add such things as the Application Policies extension, which defines the purposes for which a generated certificate


can be used. The Issuance Policies extension is also worth mentioning, because it defines when a certificate may be issued.

Figure 12.23 Extensions Tab of the New Template Property Sheet



The **Security** tab is similar to the **Security** tab that we saw in Exercise 12.02, except that this tab is used to control who may edit the template and who may request certificates using the template. Figure 12.24 shows the default permission level for the **Authenticated Users** group. For a user to request a certificate, however, the user must have at least the **Enroll** permission assigned to him or her for manual requests and the **Autoenroll** permission for automatic requests.

Figure 12.24 Security Tab of the New Template Property Sheet




After you have configured a particular template, the CA still cannot use it to issue certificates until it is made *available*. To enable a template, use the **Certification Authority** console and right-click the **Certificate Templates** container. Selecting **New | Certificate Template to Issue** completes the process (you will use this procedure in Exercise 12.04).

Certificate Requests

A client has three ways to request a certificate from a CA. The most common is auto-enrollment, and we'll discuss its deployment shortly. A client can also request a certificate by use of the **Certificates** snap-in. Clicking **Start | Run**, typing in **certmgr.msc** and pressing **Enter** can launch the snap-in, shown in Figure 12.25. Note that the **Certificates** snap-in does *not* appear in the **Administrative Tools** folder as the **Certification Authority** snap-in does after installing certificate services.

Figure 12.25 Certificates Snap-In



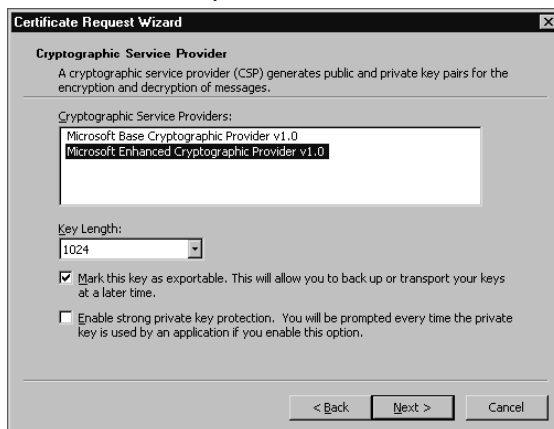
Next, by expanding the **Personal** container and right-clicking the **Certificates** container beneath it, you can start the **Certificate Request Wizard** by choosing **All Tasks | Request New Certificate**. After the welcome screen, the first screen of the wizard enables you to choose the certificate type. Figure 12.26 shows you the available options. You can only choose a type for which the receiving CA has a template.

Figure 12.26 Certificate Type Screen of the Certificate Request Wizard

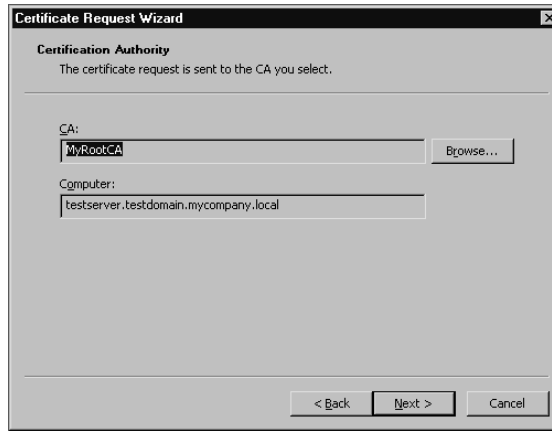


If you select the **Advanced** check box, the next screen (Figure 12.27) enables you to choose the Cryptographic Service Provider (CSP) and key length. You can also mark the key as exportable and/or enable strong private key encryption.

Figure 12.27 Cryptographic Service Provider Screen of the Certificate Request Wizard



Continuing with the advanced options, you can choose **Browse the domain** to choose a CA to which you want to send the request (Figure 12.28).

Figure 12.28 Certification Authority Screen of the Certificate Request Wizard

Finally, the wizard finishes by prompting you for a friendly name and description for the certificate.

The last method for requesting a certificate is to use a Web browser on the client machine. Note that if you use this option, IIS must be installed on the CA. Exercise 12.03 shows the steps for requesting a certificate using a client machine in this manner.



TEST DAY TIP

The order of component installation can be important when dealing with CAs. If you install certificate services *before* you install IIS, a client will *not* be able to connect as in the exercise below until you run the following from the command line: **certutil -vroot**. This establishes the virtual root directories necessary for Web enrollment. Note also that you must have selected the Web enrollment support option during the certificate Services installation procedure that you completed in Exercise 12.01.

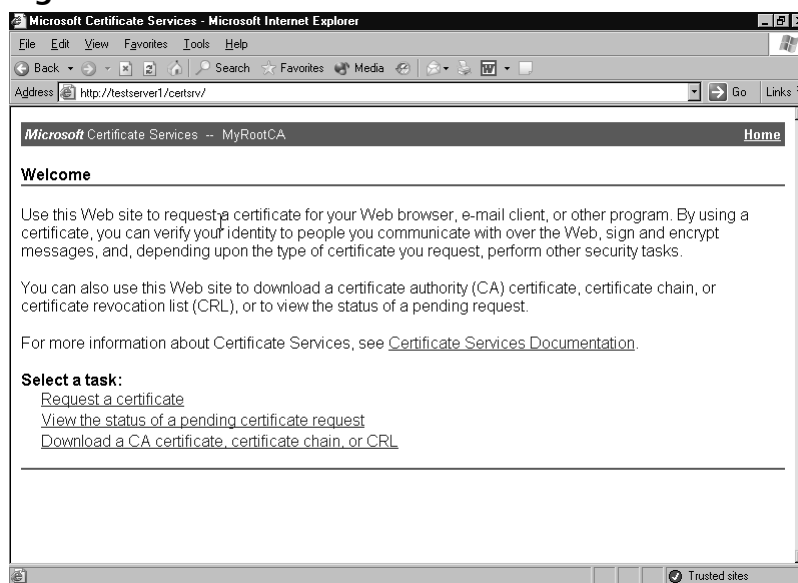
EXERCISE 12.03

REQUESTING A CERTIFICATE FROM A WEB SERVER

1. On any computer for which you want to request a certificate, launch Internet Explorer (version 5.0 or later) by clicking **Start | Programs or All Programs | Internet Explorer**.
2. In the address bar, type **http://servername/certsrv**, where *servername* is the name of the issuing CA.

3. When the screen appears, as shown in Figure 12.29, click **Request a Certificate**.

Figure 12.29 Welcome Screen of the CA's Web Site



4. Click **User Certificate**, then **Submit** when the next screen appears.
5. When the **Certificate Issued** page appears, click **Install This Certificate**. Close the browser.

Auto-Enrollment Deployment

Perhaps the most exciting new feature of the Windows Server 2003 PKI is the ability to use auto-enrollment for user certificates as well as for computer certificates. The request and issuance of these certificates may proceed without user intervention. There are, however, some strict requirements:

- Only Windows Server 2003 clients or Windows XP clients can use auto-enrollment.
- Windows Server 2003 Enterprise Edition or Datacenter Edition is required to configure auto-enrollment for version 2 templates.

Group policies are used in Active Directory to configure auto-enrollment. In **Computer Configuration** | **Windows Settings** | **Security Settings** | **Public Key**

Policies, there is a group policy entitled **Automatic Certificate Request Settings**. The property sheet for this policy enables you to choose to either **Enroll certificates automatically** or not. Also, you will need to ensure that **Enroll subject without requiring any user input** option is selected on the **Request Handling** tab of the certificate template property sheet. Finally, be aware that doing either of the following will cause auto-enrollment to fail:

- Setting the **This number of authorized signatures** option on the **Issuance Requirements** tab to higher than one.
- Selecting the **Supply in the request** option on the **Subject Name** tab.



TEST DAY TIP

Remember that auto-enrollment is available for user certificates only if the client is Windows XP or Windows Server 2003, and you must be logging on to a Windows Server 2003 domain. Machine certificates can be issued via auto-enrollment with Windows 2000.

Role-Based Administration


In a small network of one or two servers and just a handful of clients, administration is generally not a difficult task. When the size of the network increases, however, the complexity of administration seems to increase exponentially. Microsoft's recommendations for a large network include dividing administrative tasks among the different administrative personnel. One administrator may be in charge of backups and restores, whereas another administrator may have complete control over a certain domain, and so on. The role of each administrator is defined by the tasks that he or she is assigned to, and individual permissions are granted based on those tasks. PKI administration, which can be as daunting as general network administration, can be similarly divided. Microsoft defines five different roles that can be used within a PKI to facilitate administration:

- CA Administrator
- Certificate Manager
- Backup Operator
- Auditor
- Enrollee

At the top of the hierarchy is the CA administrator. The role is defined by the *Manage CA* permission and has the authority to assign other CA roles and to renew the CA's certificate. Underneath the CA administrator is the certificate manager. The certificate man-

ager role is defined by the *Issue and Manage Certificates* permission and has the authority to approve enrollment and revocation requests.

The Backup Operator and the Auditor roles are actually operating system roles and are not CA-specific. The Backup Operator has the authority to back up the CA and the Auditor has the authority to configure and view audit logs of the CA. The final role is that of the Enrollees. All authenticated users are placed in this role and are able to request certificates from the CA.



Implementing Smart Card Authentication in the PKI

If security is a primary concern for your organization, you might want to consider the use of smart cards for both local and remote authentication. This adds a second level of security to the authentication process. Whereas traditional authentication via password requires only “something you know” (the password), smart card authentication also requires “something you have” (the card).

Along with biometric devices such as fingerprint readers and retinal scanners, smart cards represent a more secure way for users to gain access to the network. Smart cards are not as secure as most biometric devices, but they are more widely implemented and have a longer history of use (more than 11 years). In fact, there are many companies that issue smart cards and smart card readers along with Windows Server 2003–compliant drivers and software. Primarily because of several competing standards, smart card adoption has been slow, but their popularity continues to grow. They might not replace the standard log-on password anytime soon, smart card technology is full of potential.

What Are Smart Cards?

Most smart cards today look and feel like a credit card. The difference is that smart cards have either integrated circuit technology (with gold-colored metallic contact points on the surface of the card) or magnetic technology (located inside the card). Smart cards use these technologies to house an embedded microprocessor that is capable of storing everything from encryption keys to medical information (at least in theory). At present, smart cards are generally used for authentication, and sometimes for encrypted e-mail.

In the case of authentication, the user inserts a smart card into a smart card reader and enters his or her personal identification number, or PIN, similar to the process of using an ATM bank card. The reader can then forward the secure information contained in the card, eliminating the need for the user to type in a name and password (the actual process of authentication is described in more detail below). It is a more secure method of authentication because theft of a smart card does not compromise security—without the PIN, the card is useless, and without the card, knowledge of the PIN is useless.



TEST DAY TIP

For the exam, be aware that when smart card authentication is used, you cannot promote a server to a domain controller and you cannot join a computer to a domain. To accomplish these administrative tasks, a normal password logon is required from a user with appropriate permissions.

How Smart Card Authentication Works

After setting up an enrollment station (described below), any user with the enrollment agent certificate can issue smart cards to users. Enrollment is the process by which a CA grants a certificate to the card. The card itself generates a public/private key pair, and the certificate is used to protect the public key during transport. After enrollment, the user can insert the card at any workstation on the network, including terminal services clients and remote access clients, as long as a smart card reader is present.

If possible, clients logging on to a Windows Server 2003 network will be authenticated with the Kerberos protocol. In traditional authentication, a username and password typed in via the keyboard are used to encrypt communication between the client and the Key Distribution Center (KDC). With smart cards, however, the private key stored in the card digitally automatically signs the timestamp that is sent to the KDC, eliminating the need for a password. In addition to the encrypted timestamp, the card's certificate (including of course the card's public key) is sent as well. When the KDC receives the package, known as a ticket-granting ticket (TGT) request, it verifies the public key and then uses the public key to verify the digital signature on the request. If everything checks out, the server authenticates the client by returning a ticket that is also encrypted with the card's public key. Finally, the ticket is decrypted at the client's workstation by the private key stored in the smart card.



EXAM WARNING

The CA that issues the smart card certificates must reside in the same network forest as the users. Users from a different forest will use local domain controllers to authenticate, and these domain controllers will not be able to validate the certificates they receive.

Deploying Smart Card Logon

Even though smart cards have been around for some time, many different standards still exist. This can complicate the deployment of a smart card solution, especially if Windows Server 2003 does not natively support the hardware you've chosen. In that case, several extra steps are required. Windows Server 2003, out of the box, contains drivers for two

companies that manufacture smart cards and readers – Schlumberger and Gemplus. For any other vendor's equipment, you'll need to install drivers and the CSP that the vendor uses.

The first step in deployment is to prepare the appropriate certificate templates. These templates include the following:

- Enrollment agent
- Smart card logon
- Smart card user certificates

The templates are not enabled by default and require some configuration. The second step is to issue the enrollment agent certificate. Finally, the smart cards need to be enrolled at the enrollment station. We'll guide you through the step-by-step deployment process later in this chapter.

Smart Card Readers

Most smart card readers in today's market attach to the computer's USB or serial port. USB equipment is strongly recommended if your clients have USB ports. Readers are available in external or internal models, and many cost less than fifty dollars at retail. Readers that are built into a keyboard are also gaining in popularity. Make certain that the readers you choose will read the kind of smart card you plan on issuing.

Smart Card Enrollment Station

The enrollment station you choose should be a secure system and must be running Windows 2000 or higher. Of course a smart card reader must be installed and appropriate drivers and CSPs loaded if necessary. Finally, you should install any vendor-supplied utility software.

Using Smart Cards To Log On to Windows


Smart cards can be used for more than secure authentication. In fact, there are two different templates in Windows Server 2003 that are both used for smart card certificates. The first is the smart card log-on certificate, which, as the name implies, is used only for logons. The second is the smart card user certificate, which, in addition to logons, provides secure e-mail services. For the following exercise, you'll use the more common of the two which is the smart card logon certificate. You will have to have a PKI implemented with at least one CA already running before beginning. You will also need a smart card reader and a smart card to complete this exercise.

EXERCISE 12.04

IMPLEMENTING AND USING SMART CARDS


1. On the system acting as the CA, log on as an administrator and open the Certification Authority console by clicking **Start | Programs | Administrative Tools | Certification Authority**.
2. Expand the appropriate CA container, right-click **Certificate Templates**, and choose **New | Certificate Template to Issue** (see Figure 12.30).

Figure 12.30 Enable Certificate Templates Window



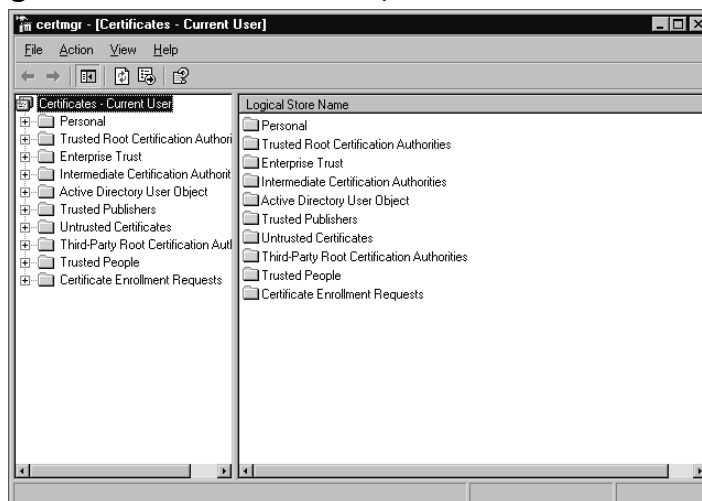
3. Select **Smartcard Logon** and click **OK**. Repeat step 2 and select **Enrollment Agent** and then click **OK**.
4. Right-click **Certificate Templates** and choose **Manage**. As shown in Figure 12.31, this displays the Certificate Template snap-in.

Figure 12.31 Certificate Templates Snap-In



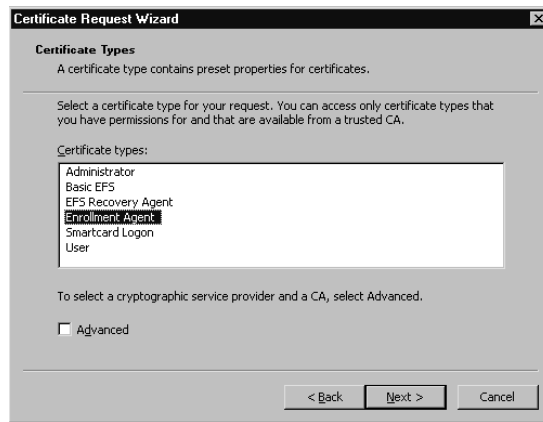
5. Right-click the **Smartcard Logon** template and choose **Properties**. Click the **Security** tab.
6. For this exercise, assign the administrator the role of enrollment agent. Add the **Administrator** by clicking the **Add** button. After selecting the **Administrator**, select the **Read** and **Enroll** check boxes. Click **OK** to finish. Close the console.
7. Log on to the enrollment station system as an administrator. Click **Start** | **Run**, type **certmgr.msc**, and then click **OK**. This launches the **Certificates** snap-in, as seen in Figure 12.32.

Figure 12.32 Certificates Snap-In



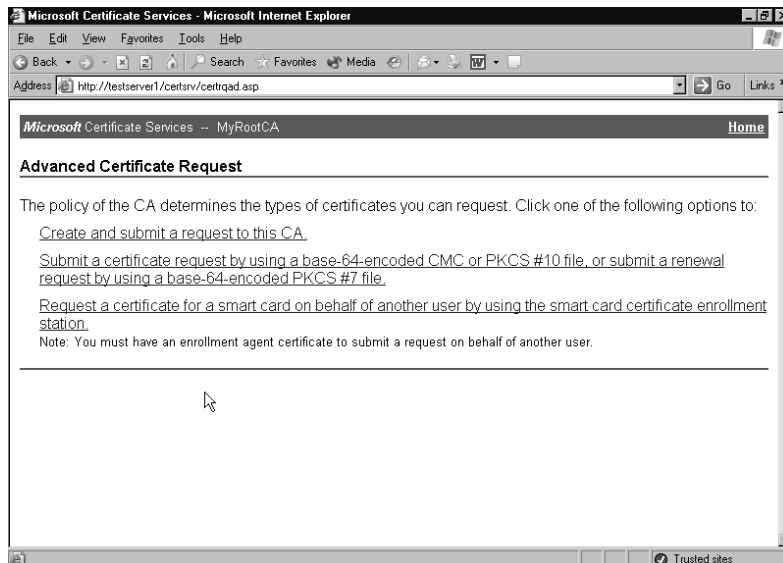
8. Expand the **Personal** container, right-click **Certificates**, and choose **All Tasks** | **Request New Certificate**. Proceed past the **Certificate Request Wizard**'s opening screen by clicking **Next**.
9. Figure 12.33 shows the **Certificate Types** screen. Choose **Enrollment Agent** and click **Next**. On the next screen, do not type anything in for the **Certificate Friendly Name** and **Description** fields. These fields are optional, and you will not use friendly names or their descriptions in this exercise. Click **Next**, and then click **Finish**. A message appears when the certificate has been issued. Close the console.

Figure 12.33 Certificate Request Wizard's Certificate Types Screen



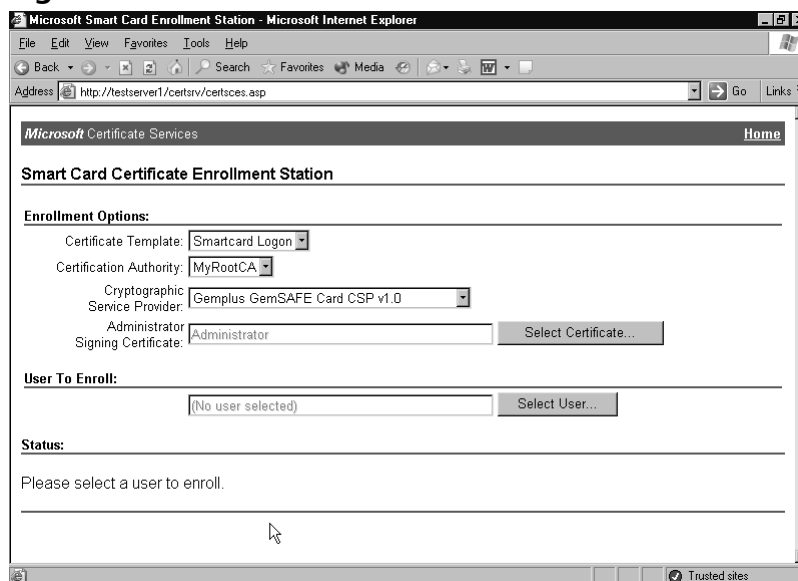
10. You will now use a similar certificate-requesting technique to Exercise 12.03, but with more advanced options. Launch **Internet Explorer** and type **http://servername/certsrv** in the **Address** bar, where **servername** is the server name of the CA you used in step 1.
11. Click **Request a certificate** and then click **Advanced certificate request** on the next screen.
12. Figure 12.34 shows the **Advanced Certificate Request** screen. Click **Request a certificate for a smart card on behalf of another user by using the smart card certificate enrollment station**.

Figure 12.34 Advanced Certificate Request Screen



13. Figure 12.35 shows the Smart Card Certificate Enrollment Station screen. Select **Smartcard Logon** from the **Certificate Template** drop-down box.

Figure 12.35 Smart Card Certificate Enrollment Station Screen



14. Select the CA used in Step 1 from the **Certification Authority** drop-down box.
15. Select the appropriate CSP from the **Cryptographic Service Provider** drop-down box.
16. Click the **Select User** button and choose the user you are enrolling.
17. Place the smart card into the attached reader and click **Enroll**.
18. The CSP will now enable you to enter a PIN for the card. Enter the PIN and click **OK**.
19. Distribute the card to the user for testing.

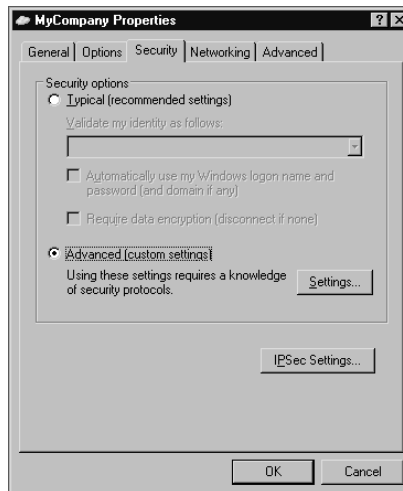
Using Smart Cards for Remote Access VPNs

The use of smart cards for local logons has met with limited, albeit recently growing, success. One reason for the limited acceptance is that local authentication traffic does not usually pass over insecure public networks; therefore, the added cost and administrative effort

required for a smart card implementation is not justified. For remote access users, however, authentication communications are vulnerable, and smart cards can provide needed extra security.


To use smart cards to log on to a remote access VPN server, the server must first be configured to enable it. This includes selecting a protocol, as discussed below. It also includes obtaining a machine certificate for the VPN server. When the server is able to accept smart card certificates, the client must be configured to send them. This means attaching a smart card reader and establishing a VPN connection. If you view the Properties of the client's VPN connection, you will notice a **Networking** and a **Security** tab. For smart card use, the type of VPN selected under the **Network** tab should be the Level 2 Tunneling Protocol, or L2TP. The **Security** tab, shown in Figure 12.36, is a bit more complex. There are two options, **Typical** and **Advanced**.

Figure 12.36 Security Tab of the VPN Client's Properties Sheet



Choose **Advanced (custom settings)** and click the **Settings** button. Choose the **Use Extensible Authentication Protocol (EAP)** option and select **Smart Card or other certificate (encryption enabled)** from the drop-down box. Click the **Properties** button, and the **Smart Card or Other Certificates** dialog box appears as shown in Figure 12.37. Choose the **Use my smart card** option. Your configuration of the VPN client is now complete.

Figure 12.37 Smart Card or Other Certificate Properties Sheet



Enabling the Extensible Authentication Protocol (EAP) on a Remote Access Server

The steps required to prepare a server that is already running the routing and remote access services (RRAS) to use smart card authentication are fairly straightforward. First, from the RRAS console, display the properties sheet for the server and proceed to the **Security** tab. Next, choose **Windows Authentication** and select **Authentication Methods**. Choose the **Extensible Authentication Protocol (EAP)** option and select **EAP Methods**. Finally, choose the **Smart Card or Other Certificate** option.

Configuring **Remote Access Policies** is also relatively simple. You can create a new policy or edit the existing **Allow Access If Dial-In Permission Is Enabled** policy. After going into the policy's property sheet, choose to **Edit the profile** and proceed to the **Authentication** tab. Select **EAP Methods** and click **Add** to choose the **Smart Card or Other Certificates** option. Clicking **Edit** brings up the property sheet for the option. To complete the edit, select the RRAS server's fully qualified domain name (FQDN) in the **Certificate Issued To** field.

Using Smart Cards To Log On to a Terminal Server

Using smart cards to log on to a terminal server is inherently more secure than using passwords, as we've discussed previously. Similar to using a smart card on a local workstation, using a smart card on a terminal client enables the server to verify your identity and give you appropriate access. Also, if you want the information contained in the card to be available for the entire terminal session, perform the following steps:

1. Click **Start | Programs** or **All Programs | Accessories | Communications | Remote Desktop Connection**.
2. Click **Options** and proceed to the **Local Resources** tab.
3. Under **Local Devices**, click the **Smart Card** option and click **Connect**.

Summary of Exam Objectives

The purpose of a PKI is to facilitate the sharing of sensitive information such as authentication traffic across an insecure network. This is done with public and private key cryptography. In public key cryptography, keys are generated in pairs so that every public key is matched to a private key and vice versa. If data is encrypted with a particular public key, then only the corresponding private key can decrypt it. A digital signature means that an already encrypted piece of data is further encrypted by someone's private key. When the recipient wants to decrypt the data, he or she must first "unlock" the digital signature by using the signer's public key, remembering that only the *signer's* public key will work. This might seem secure, but because anyone at all can sign the data, how does the recipient know for certain the identity of the person who actually signed it?

The answer is that digital signatures need to be issued by an authoritative entity, one whom everyone trusts. This entity is known as a certification authority (CA). An administrator can use Windows Server 2003, a third-party company such as VeriSign, or a combination of the two to create a structure of CAs. Certification authorities, as the name implies, issue certificates. In a nutshell, certificates are digitally signed public keys. Certificates work something like this: party A wants to send a private message to party B and wants to use party B's public key to do it. Party A realizes that if B's public key is used to encrypt the message, then only B's private key can be used to decrypt it, and since B and no one else has B's private key, everything works out well. However, A needs to be sure that he's really using B's public key and not an imposter's, so instead of just asking B for B's public key, he asks B for a certificate. B has previously asked the CA for a certificate for just such an occasion (B will present the certificate to anyone who wants to verify B's identity). The CA has independently verified B's identity and has then taken B's public key and signed it with its own private key, creating a certificate. Party A trusts the CA and is comfortable using the CA's well-known public key. When A uses the CA's public key to unlock the digital signature, he can be sure that the public key inside really belongs to B, and he can take that public key and encrypt the message.

The "I" in PKI refers to the infrastructure, which is a system of public key cryptography, certificates, and certification authorities. CAs are usually set up in a hierarchy, with one system acting as a root and all the others as subordinates at one or more levels deep. By analyzing the certificate requirements for your company, you can design your CA structure to fit your needs. Most organizations use a three-tier model, with a root CA at the top, an intermediate level of subordinates who control CA policy, and a bottom level of subordinates who actually issue certificates to users, computers, and applications. In addition to choosing root and subordinate structure for the CA hierarchy, each CA during installation needs to be designated as either an enterprise or a standalone. Each of these choices has distinct advantages and disadvantages. Most CA configuration after installation is done through the Certification Authority snap-in. In addition to issuing certificates, CAs are responsible for revoking them when necessary. Revoked certificates are published to a CRL that clients can download before accepting a certificate as valid.

Enterprise CAs use templates to know what to do when a certificate request is received and how to issue a certificate if approved. Windows Server 2003 includes several built-in templates, or you can configure new ones. After a CA is ready to issue certificates, clients need to request them. Auto-enrollment, Web enrollment, or manual enrollment through the Certificates snap-in are the three ways by which a client can request a certificate. Auto-enrollment is available for computer certificates, and in Windows Server 2003 for user certificates as well.

Finally, using smart cards for authentication requires the use of a PKI. Using a card reader, a local or a remote user can insert his or her card and enter a PIN in place of typing in a username and password. This method of authentication uses EAP and is extremely secure, especially for remote access users using a corporate VPN. An enrollment agent (a user who holds an Enrollment Agent certificate) uses an enrollment station that has been pre-configured to put information such as a certificate on the cards before they're issued to users. Also, smart cards may be used for secure e-mail or for logging on to a terminal server.

Exam Objectives Fast Track

Planning a Windows Server 2003 Certificate-Based PKI

- ☑ A PKI combines public key cryptography with digital certificates to create a secure environment where network traffic such as authentication packets can travel safely.
- ☑ Public keys and private keys always come in pairs. If the public key is used to encrypt data, only the matching private key can decrypt it.
- ☑ When public key-encrypted data is encrypted again by a private key, that private key encryption is called a digital signature.
- ☑ Digital signatures provided by ordinary users aren't very trustworthy, so a trusted authority is needed to provide them. The authority (which can be Windows-based) issues certificates, which are basically digitally signed containers for public keys and other information.
- ☑ Certificates are used to safely exchange public keys and to provide the basis for applications such as IPsec, EFS, and smart card authentication.

Implementing Certification Authorities

- ☑ Certificate needs are based upon which applications and communications an organization uses and how secure they need to be. Based on these needs, CAs are created by installing certificate services and are managed using the Certification Authority snap-in.

- ☑ A CA hierarchy is structured with a root and one or more level of subordinates – three levels is common. The bottom level of subordinates issues certificates. The intermediate level controls policies.
- ☑ Enterprise CAs require and use Active Directory to issue certificates, often automatically. Stand-alone CAs can be more secure and need an administrator to manually issue or deny certificate requests.
- ☑ CAs need to be backed up consistently and protected against attacks. Keys can be archived and later retrieved if they are lost. This is a new feature for Windows Server 2003.
- ☑ CAs can revoke as well as issue certificates. After a certificate is revoked, it needs to be published to a CRL distribution point. Clients check the CRL periodically before they can trust a certificate.

Planning Enrollment and Distribution of Certificates

- ☑ Templates control how a CA acts when handed a request and how to issue certificates. There are quite a few built-in templates, or you can create your own using the Certificate Template snap-in. Templates must be enabled before a CA can use them.
- ☑ Certificates can be requested with the Certificates snap-in or by using Internet Explorer and pointing to *http://servername/certsrv* on the CA.
- ☑ Machine and user certificates can be requested with no user intervention requirement by using auto-enrollment. Auto-enrollment for user certificates is new to Windows Server 2003.
- ☑ Role-based administration is recommended for larger organizations. Different users can be assigned permissions relative to their positions, such as certificate manager.

Implementing Smart Card Authentication in the PKI

- ☑ Smart cards are credit card-like devices that embed a microprocessor. They can securely hold public/private keys, certificates, and other information.
- ☑ Users insert smart cards into readers and enter a PIN to use information contained on the card. Authentication is the most popular application of the technology, followed by secure e-mail services.
- ☑ To deploy smart cards, you need to configure the CA to issue smart card and enrollment agent certificates, set up an enrollment agent, and set up an enrollment

station. Smart cards have to be enrolled, or set up with appropriate information, before someone can use one.

- ☑ Smart cards are increasingly used for remote access authentication, such as over a company VPN. Cards can also be used for securely logging on to a terminal server.

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: In what format do CAs issue certificates?

A: Microsoft certificate services use the standard X.509 specifications for issued certificates and the Public Key Cryptography Standard (PKCS) #10 for certificate requests. The PKCS #7 certificate renewal standard is also supported. Windows Server 2003 also supports other formats, such as PKCS #12, DER-encoded binary X.509, and Base64 Encoded X.509, for exporting certificates to computers running non-Windows operating systems.

Q: If certificates are so important in a PKI, why don't I see more of them?

A: Many portions of a Windows PKI are hidden to the end user. Thanks to features such as auto-enrollment, some PKI transactions can be completely handled by the operating system. Most of the work in implementing a PKI comes in the planning and design phase. Operations such as encrypting data via EFS use certificates, but the user does not "see" or manually handle the certificates.

Q: I've heard that I can't take my laptop overseas because it uses EFS. Is this true?

A: Maybe. The backbone of any PKI-enabled application such as EFS is encryption. Although the U.S. government now permits the exporting of "high encryption" standards, some countries still do not allow their import. The Windows Server 2003 PKI can use high encryption, and so the actual answer depends on the country in question. For information on the cryptographic import and export policies of a number of countries, see <http://www.rsasecurity.com/rsalabs/faq/6-5-1.html>.

Q: Can I create my own personal digital signature and use it instead of a CA?

- A:** Not if you need security. The purposes behind digital signatures are privacy and security, and a digital signature at first glance seems to fit the bill. The problem, however, is not the signature itself, but the lack of trust in a recipient. Impersonations become a looming security risk if you can't guarantee that the digital signatures you receive came from the people with whom they were supposed to have originated. For this reason, a certificate issued by a trusted third party provides the most secure authentication.
- Q:** Can I have a CA hierarchy that is five levels deep?
- A:** Yes, but that's probably overkill for most networks. Microsoft's three-tier model of root, intermediate, and issuing CAs will more than likely meet your requirements. Remember that your hierarchy can be wide instead of deep.
- Q:** Do I have to have more than one CA?
- A:** No. Root CAs have the capability to issue all types of certificates and can assume responsibility for your entire network. In a small organization, a single CA might be sufficient for your purposes. For a larger organization, however, this structure would not be suitable.
- Q:** How can I change the publishing interval of a CRL?
- A:** From the **Certification Authority** console, right-click the **Revoked Certificates** container and choose **Properties**. The **CRL Publishing Parameters** tab enables you to change the default interval for full and Delta CRLs.
- Q:** Why can't I seem to get auto-enrollment for user certificates to work?
- A:** Remember that auto-enrollment for machines is a feature that has been around since Windows 2000, but auto-enrollment for user certificates is new to Windows Server 2003. To use this feature, you need to be running either a Windows Server 2003 or Windows XP client and you must log on to a Windows Server 2003 domain. Finally, auto-enrollment must be enabled through Active Directory's group policy. Also, you won't be able to auto-enroll a user unless the user account has been assigned an e-mail address.
- Q:** What is the default validity period for a new certificate?
- A:** The default, which can be changed on the **General** tab of a new template's property sheet, is one year. Other important settings, such as minimum key size and purpose of the certificate, can be found on the sheet's other tabs.

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 a Windows Server 2003 Certificate-Based PKI

1. You are setting up a procedure to keep documents exchanged between members of the R & D department secret. They will be sending these documents across the Internet to each other. Which PKI process will you need to employ to achieve this?
 - A. Confidentiality
 - B. Non-repudiation
 - C. Authentication
 - D. Data Integrity

Implementing Certification Authorities

2. You are the administrator for a large and very busy network and your bandwidth is nearing its limits. Your users are complaining about the time it takes to access the payroll server to update their hours. All users are required to have certificate authentication to access the server. What can you change in your current setup to help reduce network traffic and speed access to the payroll server?
 - A. Configure the CA to use complete CRLs for replication.
 - B. Assign times for each user to update their payroll.
 - C. Use DES for the encryption method.
 - D. Configure the CA to use Delta CRLs.
3. Your department has completed preliminary testing of a newly established PKI, and before actual deployment begins, you've been assigned the task of revoking the test certificates. So far, there is only a single enterprise CA installed, and Active Directory is of course in use. Which of the following steps should you take?
 - A. In the **Certification Authority** console, expand the **Issued Certificates** container, and revoke all certificates by right-clicking each certificate and choosing **All | Revoke Certificate**.

- B. In the **Certification Authority** console, expand the **Issued Certificates** container, and revoke all certificates by right-clicking each certificate and choosing **All | Revoke Certificate**. Right-click the **Revoked Certificates** container, and choose **All Tasks | Publish**.
 - C. Using the **Certificates** snap-in, expand the **Personal** container, and highlight the **Certificates** container found beneath. In the right pane of the console, right-click each certificate and choose **Add to Certificate Revocation List**.
 - D. Using the **Certificates** snap-in, expand the **Personal** container, and highlight the **Certificates** container found beneath. In the right pane of the console, right-click each certificate and choose **Add to Certificate Revocation List**. Right-click the **Trusted Root Certification Authority**, and choose **Publish to Directory**.
4. You decide to implement a Windows Server 2003 based-PKI for your network, and because you want the most secure method of issuing and maintaining certificates, you decide to use a stand-alone server to issue a certificate to a subordinate, which in turn issues certificates to users. You take the root CA offline. Your users complain that they are unable to access some resources. After investigating the problem you discover that they can log on to the network and access everything except those resources protected by certificates. They also can connect to the servers by both name and IP address. What is preventing the users from gaining access to those resources?
 - A. The root CA server is offline.
 - B. The subordinate CA is offline.
 - C. The certificates have been compromised.
 - D. The certificates are still pending.
 5. You have a two-tier hierarchy for your certificate PKI. OurRoot is an enterprise root CA. OurIssuer1 and OurIssuer2 are OurRoot subordinates. These two CAs issue all the certificates for your company. OurIssuer1 issues to the northern region and OurIssuer2 issues to the southern region. An ex-employee appears to have obtained the issuing certificate for OurIssuer2. What steps would you take to prevent users from using certificates issued by the compromised server?
 - A. Add the compromised certificate to the CRL from *OurRoot*.
 - B. Delete all certificates on *OurIssuer2* and reissue them.
 - C. Reinstall certificate services on *OurIssuer2*.
 - D. Add all certificates issued by *OurRoot* to the CRL.

6. As a member of the PKI design team in your company, you are charged with integrating one of your subsidiaries that already has a PKI with your office's PKI. The current proposal on the table has a second-tier CA located in your local PKI issuing certificates to a second-tier CA located on the subsidiary's PKI, and vice-versa. Both infrastructures are Windows Server 2003 based. Your company's security goals, however, mandate that only certain certificates be used on your PKI if they are issued from the subsidiary's CA, but all your CA's certificates need to be trusted by the subsidiary. What is your assessment?
- A. Both your office and the subsidiary will need to create a CTL that has a limited trust chain length on your side.
 - B. The subsidiary's CA needs to be reconfigured as your CA's subordinate.
 - C. A cross-trust needs to be created, and the type of acceptable certificates for your CA narrowed by using qualified subordination policies.
 - D. This arrangement is not possible under Windows Server 2003. The company needs to implement a third-party PKI.
7. Your company has a partner with whom you need to communicate securely. You have an existing root CA and need to allow usage for partner-issued certificates as well. In which of the following ways can you accomplish this? Choose all that apply.
- A. Create a CTL.
 - B. Install an issuing CA at the partner's site.
 - C. Create a cross-trust hierarchy.
 - D. Install a partner's issuing CA at your site.
8. You are the administrator of an existing three-tier PKI including a stand-alone Root CA, three mid-level CAs, and twelve issuing CAs. You fear that your Root certificate has been compromised. What steps should you take to secure your infrastructure with the least amount of administrative effort?
- A. Add the twelve issuing CAs' certificates to the mid-level CAs' CRL.
 - B. Add the three mid-level CAs' certificates to the Root CA's CRL.
 - C. Add the Root CA's certificate to the three mid-level and twelve issuing CAs' CRL.
 - D. Create a new CA hierarchy and issue new certificates to all clients.

Planning Enrollment and Distribution of Certificates

9. You are attempting to request a certificate by using Internet Explorer, but fail to display the welcome screen of the Web site. You have typed in the address `http://mycertauthority/certsrv` and you've double-checked the name of the CA. Also, you have confirmed with the network administrator that the CA is configured with IIS, and the Web enrollment support option was chosen during the certificate services installation. What is the most likely cause of the problem?
 - A. The CA is configured as a standalone.
 - B. IIS was installed after certificate services.
 - C. The EAP protocol has not been installed.
 - D. You are using a Windows 2000 Professional client.

10. The Ecstatic Llama Company wants your consulting firm to implement a two-tier private CA design made specifically for their PKI. Because the plans for ELC call for high security, the root CA will be designated as standalone and offline. Your job is to install an enterprise subordinate CA while maintaining the security needs of your client. What are the two best methods to accomplish this task? Choose two answers.
 - A. In the Certification Authority console, configure the subordinate to use auto-enrollment and reboot the machine.
 - B. In the Certification Authority console, point the subordinate to use Active Directory and configure the subordinate to trust the root CA.
 - C. Put the root CA briefly online and use Web enrollment to obtain the root CA certificate, then take the root CA back offline.
 - D. Save the subordinate request as a PKCS #10 file, transport the file to the root CA, issue the certificate, and then transport the certificate back to the subordinate.

11. You are the CA administrator for your branch office and want to have greater control over your certificate managers. Your plan is to have each manager manage certificates over a different Active Directory group, but you do not want to give any manager the capability to renew the CA's certificate. What is your best course of action?
 - A. In the **Certification Authority** snap-in, use the **Security** tab of the CA's property sheet to configure manager restrictions.
 - B. Using the **Certificate Templates** snap-in, right-click the **Certificate Templates** container, and choose **Properties**. On the **Security** tab, give the Certificate Managers group the *Issue and Manage Certificates* permission.
 - C. In the **Certification Authority** snap-in, use the **Certificate Managers Restrictions** tab of the CA's property sheet and choose the **Restrict certificate managers** option.
 - D. It cannot be done.

12. As the network administrator for B & H Day Care Centers, you are attempting to configure a third-tier CA to issue a particular type of certificate. From the **Certificate Templates** snap-in, you have duplicated an existing template and modified it to B & H's specifications. However, users are still unable to successfully install the certificate governed by the new template. You have checked the structure of the CA hierarchy and are comfortable that no intentional attacks have taken place. What first step can you take to ensure the proper distribution of the certificate?
 - A. Launch the **Certificate Templates** snap-in, right-click the **Certificate Templates** container, and select **New | Certificate Template to Issue**. Select the new certificate template.
 - B. Launch the **Certificate Templates** snap-in and highlight the **Certificate Templates** container. In the right pane of the console, right-click the new certificate template, and choose **Properties**. From the **Publish** tab, select the **Publish to Directory** option.
 - C. From any PKI client's browser, point to **http://servername/certsrv**, where *servername* is the name of the CA that contains the new certificate template. Select the **Issue a Certificate Template** link.
 - D. Using an account with appropriate permissions, copy the new certificate template to the root CA's certificate store. From the root CA, enable the template by using the **Certificate Templates** snap-in.

Implementing Smart Card Authentication in the PKI

13. You have been designated as the enrollment agent for the entire Pants, Inc. organization during the smart card deployment that has just been completed. Your supervisor

- has now assigned you the project of updating the company's VPN solution by configuring the current RRAS server to accept smart card remote access. However, when you log on to the server and attempt to configure it, you are unsuccessful. What is the most likely reason for the failure?
- A. The Extensible Authentication Protocol (EAP) has not been installed.
 - B. You are not a member of the Administrators group.
 - C. The Routing and Remote Access Service does not have the required application certificate.
 - D. A smart card reader has not been installed on the server.
14. Your company uses smart card authentication for its local network. You are an administrator and have been directed to install a new domain controller in the main office. You install Windows Server 2003 on the new hardware and begin the *dcpromo* process. When the install process asks you for authentication, what will you need to supply to finish the promotion?
- A. Username and password
 - B. Smart card and PIN
 - C. Username and PIN
 - D. Smart card and password
15. You are the administrator of a small network, and you have recently assigned yourself as an enrollment agent for your firm's new smart card system by making sure that you have Read and Enroll permissions on the Smart Card Logon template's **Security** tab. However, when you begin testing the implementation, you discover that you are unable to fully complete a request for a certificate on behalf of another user. You are using Internet Explorer on the enrollment station computer. Which of the following, if true, could be reasons for the failure? Choose all that apply.
- A. The smart card manufacturer's CSP has not been installed on the enrollment station.
 - B. IIS has not been installed on the enrollment station.
 - C. The Write permission has not been assigned to your account.
 - D. Neither the Smart Card Logon nor the Smart Card User templates have been enabled on the CA.
 - E. You logged on to the enrollment station using your administrator account.

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

MCSE 70-293

Self Test Questions, Answers, and Explanations

This appendix provides complete Self Test Questions, Answers, and Explanations for each chapter.

Chapter 1: Using Windows Server 2003 Planning Tools and Documentation

Overview of Network Infrastructure Planning

1. You are proposing the purchase of a new e-mail server for your corporate network. You have specified a new server from a major OEM manufacturer that is configured with a powerful quad-processor configuration, hot-swappable hard drives, and redundant power supplies and network adapters, with a three-year onsite warranty. Due to a budget crunch, the chairperson of the budget committee has suggested that the company can make do with a less powerful workgroup server from a local computer store. This server has only a single processor and no redundancy features, and a one-year onsite warranty. What reasons can you provide the budget committee members that might convince them to authorize the purchase of the server that you specified, even though it has a higher price tag?
 - A. A more powerful server will provide better performance and scalability as the company's needs grow over time.
 - B. Redundant hardware components will increase the server's availability to service the needs of the company's users and customers.
 - C. The extended warranty on the more powerful server will increase support costs over time, since you're paying to cover the machine under warranty for three times as long.
 - D. Windows Server 2003 requires at least a dual-processor configuration.

A, B. When calculating total cost of ownership (TCO), you should take into account the potential that your company will outgrow a less powerful server, and thus you will need to replace it more quickly. This has the potential to cost the company more money over time than if they had made the initial investment in more powerful hardware. TCO should also take into account "soft costs," such as lost productivity and lost customer revenue due to downtime caused by hardware malfunction. Redundant hardware features will help to decrease this potential for downtime and increase the server's overall availability.

C, D. Answer C is incorrect because an extended warranty will likely reduce support costs over time, since repairs and parts will be covered for a longer period. Answer D is incorrect because Windows Server 2003, although it will offer better performance in a multiprocessor configuration, requires only a single processor to meet the installation requirements.
2. You are the network administrator of a Windows NT 4 domain for a shipping warehouse that operates 24 hours a day, 6 days a week. You perform a full nightly backup of all user files at 3:00 A.M. Users on the overnight shift are complaining that they are often locked

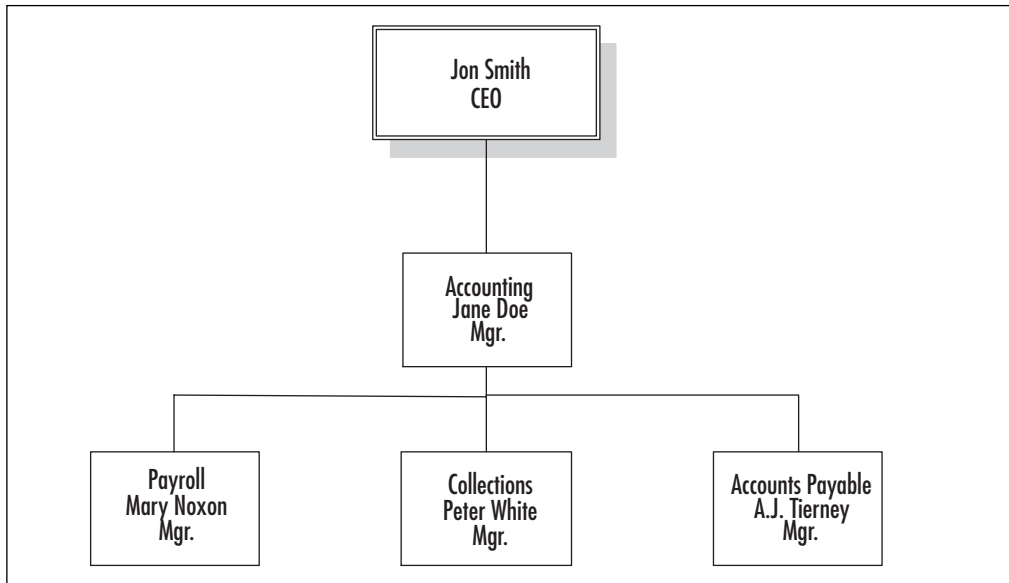
out of files that they need access to while the backup process is running. You are proposing a network upgrade to Windows Server 2003 in the near future. What Windows Server 2003 feature will assist you in addressing this problem?

- A. Disk quotas
- B. NTFS file security
- C. Volume Shadow Copy
- D. Network Load Balancing

C. Volume Shadow Copy is a new feature in Windows Server 2003 that allows users to access files and folders on a server while they are being backed up, thus increasing the availability of centrally stored information on a Windows file server.

A, B, D. Answer A is incorrect because disk quotas will assist you in monitoring and limiting disk usage on your network, but they will not address the issue of users being locked out of files that are being backed up. Answer B is incorrect because NTFS file permissions can dictate only which users or groups can read, modify, and/or delete a file or folder. The users are not being locked out of the files due to permissions issues. Answer D is incorrect because Network Load Balancing is a server clustering technology; it does not address the issue of accessing files that are being backed up.

3. A portion of your company's organizational structure is shown in Figure 1.14. Third-level department managers report to the second-level department managers directly above them in the organizational chart. Second-level managers report to their corresponding vice presidents, who then report to the company CEO. Your company CEO would like a consistent security policy to be implemented across the entire network, but each subdepartment has specific desktop and application installation settings that you would like to be able to control and deploy centrally. What is the most efficient AD structure to design for this company?

Figure 1.14 Organizational Structure

- A. Configure a single domain for the organization, and configure a series of nested OUs for each second-level and third-level department. Configure the domain with a single security policy, and link a GPO to each OU to enable each specific department's desired settings.
 - B. Configure a parent domain for each second-level department, and configure a child domain for each third-level department. Create and link a separate GPO to each domain to control security and application settings.
 - C. Configure a single domain for the organization, and configure a global security group for each department. Configure the domain with a single security policy, and link a GPO to each global group to enable each specific department's desired settings.
 - D. Create a separate forest for each second-level department, and create a child domain for each third-level department. Configure a security policy for each forest, and configure a domain GPO for each third-level department.
- A.** A hierarchical organizational structure such as the one described in this question lends itself to a domain structure with nested OUs. Using OUs will also allow you to apply departmental settings centrally through GPOs.
- B, C, D.** Answer B is incorrect because a single domain is more efficient if a network environment has unified security requirements. Answer C is incorrect because GPOs cannot be linked to group objects; they can be linked to only sites, domains and OUs. Answer D is incorrect because it creates more administrative overhead than is necessary for the environment described in the question.

4. You are the administrator for a network that supports a mixture of Windows NT 4 Workstation, Windows 2000, and Windows XP Professional. You are preparing to upgrade your network servers from Windows NT Server to Windows Server 2003. What is the strongest level of network authentication that you can configure your Windows domain to use in its current configuration (without installing third-party software)?
- A. Kerberos
 - B. LM
 - C. NTLM
 - D. NTLM version 2
- D.** The strongest authentication type that Windows NT 4 Workstation or Server can use without installing a third-party Kerberos client is NTLM version 2.
- A, B, C.** Answer A is incorrect because Kerberos authentication can be used only by the Windows 2000 and Windows XP Professional workstations on your network, not the NT 4 workstations. Answer B is incorrect because LM authentication is a weaker authentication protocol than NTLM or NTLM version 2. Answer C is incorrect because NTLM, while stronger than LM authentication, is still not as effective as NTLM version 2.

Analyzing Organizational Needs

5. You are the administrator of a Windows 2000 network and are planning an upgrade to Windows Server 2003. As part of the upgrade process, you are attempting to determine whether you need to upgrade your network cabling from Token Ring cabling to 100MB Ethernet. What is the best way to go about making this determination?
- A. Use Performance Monitor to capture a baseline of network utilization at several points during the day over the course of several weeks.
 - B. Use Network Monitor to capture network frames being sent to and from your domain controller's network adapter.
 - C. Use the IPsec Monitoring utility to view network traffic being sent between your domain controllers and your Windows 2000 Professional clients.
 - D. Use Performance Monitor to capture a single snapshot of network utilization when most users are in the office, such as mid-morning.
- A.** To obtain a network baseline, you need to view network utilization over an extended period of time, to get an accurate picture of your network's overall bandwidth utilization. You can use the Windows Performance Monitor utility to collect this information quickly and easily.

B, C, D. Answer B is incorrect because you would use the Performance Monitor utility to view network bandwidth utilization most efficiently. Network Monitor is used to capture individual packets. Although it does provide some bandwidth usage information, it does not directly monitor network performance. Answer C is incorrect because the IPsec Monitoring utility will examine only the IPsec protocol, not overall network utilization. Answer D is incorrect because taking only a single snapshot of network traffic will not provide an accurate picture of network utilization, which can fluctuate throughout the course of a day and week.

6. After returning from a two-day technology management seminar, your CEO tells you that he would like to create a fault-tolerant configuration for the company's heavily trafficked Web and database servers. Your network is currently running the Standard Edition of Windows NT 4. You have recently proposed an upgrade to Windows Server 2003. What features offered by this proposed upgrade would provide an attractive option to meet your CEO's request?

- A. SMP processing
- B. Volume Shadow Copy
- C. Network Load Balancing
- D. Server clustering

C, D. Windows Server 2003 provides the ability to configure Network Load Balancing and server clustering to increase the fault tolerance of your network services and resources.

A, B. Answer A is incorrect because SMP processing allows a Windows server to effectively utilize hardware with multiple processors, but it does not provide fault tolerance. Answer B is incorrect because Volume Shadow Copy is used to provide uninterrupted file and folder access during network backups. Although regular backups are part of your disaster recovery solution, a backup solution should be implemented in addition to fault-tolerance schemes, not in place of them.

7. You are the network administrator for a medium-sized company that consists of Sales, Customer Service, Accounting, Human Resources, and Data Entry departments. You have been receiving complaints that your company's e-mail server has been performing more slowly than usual over the past several weeks. Several users have mentioned that their e-mail clients have "frozen" in the middle of sending an e-mail message, forcing them to reboot their machines. Upon investigating, you find that one user's mailbox is roughly ten times the size of the second largest mailbox on the server, and this user is receiving approximately 1000 messages per day, compared to a company average of 46. The user in question is a data-entry clerk who does not use e-mail for sales inquiries or other business-related contacts. When you ask the user about her e-mail usage, she reports that she has been surfing the Web signing up for Internet coupons and contests, and she has been

deluged with spam as a result. Since the user does not require e-mail access to perform her job function, you disable her e-mail account, and server performance slowly returns to normal. What measures can you implement to prevent this sort of incident from recurring? (Select all that apply.)

- A. Implement disk quotas on the e-mail server so that users' inboxes cannot exceed a certain size.
 - B. Increase the level of authentication security so that only Kerberos-authenticated users can access the e-mail server.
 - C. Distribute an Acceptable Use Policy to your user base so that they understand what they can and cannot do while using their office PCs.
 - D. Use NTFS file permissions to restrict network access to personnel in your Sales and Customer Service department only.
- A, C.** Implementing disk quotas will prevent a single user's disk usage from interfering with the operation and performance of the entire server. An Acceptable Use Policy will provide all of your users with a list of appropriate uses for their office computers so that they can make the correct decision regarding their e-mail, Web, and other computer use.
- B, D.** Answer B is incorrect because the level of user authentication was not the source of the problem; inappropriate use of the company's computer resources is what caused the performance slowdown on the e-mail server. Answer D is incorrect because this solution will adversely affect the network access of the other users on your network who require network access.

Developing a Test Network Environment

8. You are the network administrator for a law firm that has multiple locations throughout the United States. Your firm has purchased a customer relationship management (CRM) application that will be hosted in the firm's main office in Key Biscayne, Florida, and accessed by other offices using dedicated WAN links. You would like to test the performance of this software over a WAN link before deploying it to the other offices in the firm. Unfortunately, you only have access to test equipment in the Key Biscayne office location. What is the best way to test the performance of this application?
- A. Use the average network bandwidth utilization in each office to estimate the performance of the application over the WAN.
 - B. Install routers within the test lab to simulate the latency of the dedicated WAN links between offices.
 - C. Access the CRM application from your home computer using your high-speed Internet connection.
 - D. Test the application using production systems in each of the remote offices.

- B.** You can simulate WAN traffic in a single location by installing routing equipment identical to the equipment used to connect the remote offices to the main headquarters where the application is being housed.
- A, C, D.** Answer A is incorrect because this will provide only the bandwidth usage on each office's LAN; it does not analyze the traffic that would be passed over the WAN link. Answer C is incorrect because you would be using a different access method than the one your offices would be using. This would not give you an accurate indication of the performance of the application over the dedicated WAN links. Answer D is incorrect because you should not introduce untested technologies into a production environment if you have the ability to assess them in a test environment first.
9. You are in the process of building a lab environment to test a new network application. You would like to isolate the test environment from your production equipment as much as possible to prevent any test changes from affecting your users' daily tasks. What can you do to protect your production environment from changes performed in your test lab? (Select all that apply.)
- A. Place a router or firewall between the network infrastructures connecting the test lab to your production machines.
- B. Keep the network cabling for the test lab physically separated from the network hardware that provides connectivity to your production environment.
- C. Contain the test lab in a separate OU.
- D. Use 100MB Ethernet for your production machines, but only 10MB Ethernet for the test lab.
- A, B.** To prevent unauthorized network traffic from traversing between test environment and production equipment, you can place a router or firewall between the two locations to control network traffic between the two. You can also completely segregate the two environments by using two completely disconnected cabling systems.
- C, D.** Answer C is incorrect because separating the test lab into an OU in an existing domain will not do much, if anything, to protect the production environment from changes made to the test lab. Answer D is incorrect because using two different speeds of Ethernet cabling will not prevent traffic from passing between the two environments.
10. You are designing a lab environment to test a proposed upgrade to Windows Server 2003. You are in the process of creating a domain structure in the test lab to assess various features and functions of the upgrade process, including switching the domain from mixed mode to native mode and moving from a standard DNS zone to AD-integrated DNS. At the same time that the Windows Server 2003 testing is taking place, you would also like to use the test lab to evaluate a new accounting package that will be implemented on the production network before the Windows Server 2003 upgrade takes place. You do not

want the two batteries of tests to interfere with each other. Which of the following would be good design choices for the domain structure of the test lab? (Select all that apply.)

- A. Create two separate domains: one to test the accounting software and one to test the domain mode and DNS functionality of Windows Server 2003.
 - B. Create a single domain in the test lab to encompass the entire test environment.
 - C. Create a separate OU to test the accounting software so that it will not be affected by the switch in domain mode.
 - D. Create two separate forests: one to test the DNS configuration and the switch from mixed mode to native mode and one to perform the tests on the accounting software package.
- A, D.** Changes to DNS settings, as well as switching from mixed mode to native mode functionality, will affect all users and computers in an entire domain. To isolate these changes, you should create a separate domain or forest to test them.
- B, C.** Answer B is incorrect because, in a single domain, the Windows Server 2003 testing might adversely affect the accounting software testing. Answer C is incorrect because a switch from mixed mode to native mode will still affect an OU within a single domain.
11. You have received a critical software update from the vendor of your accounting software suite. The software vendor has indicated that you should apply this patch as quickly as possible to correct a potential security breach. As the administrator for your network, what should you do when you receive this notice?
- A. Install the patch on all production systems as quickly as possible.
 - B. Install the patch in your network's test lab to ensure that it functions properly and without any adverse side effects, and then apply it to all of your production systems as soon as possible.
 - C. Install the patch on a single workstation on your production environment to see if there are any bugs or malfunctions. When you are satisfied, apply the patch to the remainder of your workstations.
 - D. Send the software patch to Microsoft Product Support Services for testing before applying it to your network computers.
- B.** Even when a patch is designed to correct a potential security breach, you should still test it for proper functionality before applying it to your production network.
- A, C, D.** Answer A is incorrect because you should test any patches before applying them to any production computers. Answer C is incorrect because installing a patch on even a single production workstation without testing it first can adversely affect your entire network. Answer D is incorrect because you should test the software patch in-house before applying it to your production network.

12. You are the network administrator for a small company that is considering purchasing a Windows 2003 Server machine to replace an aging Windows NT 4 Server machine. The client workstations run a mix of Windows 98, Windows NT Workstation, and Windows XP Professional. Each network client needs to be able to access the network server after it is upgraded, since the client workstations will be upgraded on a one-by-one basis over the course of several months. You have been informed that you will need to use the production server itself for testing, and that there is only sufficient budget to allot one representative workstation PC for test purposes. What is the best way for you to test client connectivity to Windows Server 2003?
- A. Configure the test workstation with Windows Server 2003. Connect a production Windows 98, Windows NT 4, and Windows XP Professional workstation to the test server.
 - B. Use a utility like VMware to simulate how each operating system on your network will function with the new Windows 2003 server.
 - C. Check each client operating system one at a time, reformatting the test PC after you've finished testing each operating system.
 - D. Connect a production Windows 98 and Windows NT 4 Workstation to the Windows 2003 Server. Configure the test workstation to use Windows XP Professional.
- B.** When you are working with limited hardware resources, a third-party utility like VMware will allow you to mimic several different operating systems on a single machine, allowing you to test the interoperability of multiple client operating systems with Windows Server 2003.
- A, C, D.** Answer A is incorrect because configuring the test workstation with Windows Server 2003 will not give an accurate representation of how the actual server hardware will behave. Since you have the new server that will later be placed into production, it will be the most accurate server hardware to test with. Answer C is incorrect because testing each client operating system one at a time is an inefficient method. Using a utility like VMware will allow you to test far more conveniently and quickly. Answer D is incorrect because you should not use current production equipment for testing, since you don't know how it will affect the behavior of the systems and hence the productivity of your users.

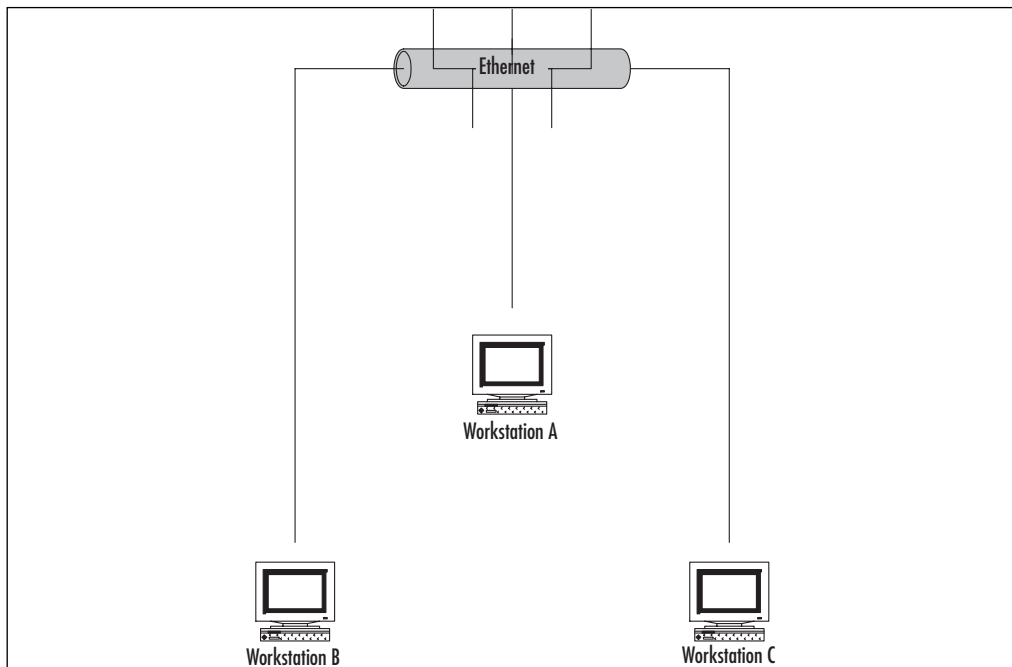
Documenting the Planning and Network Design Process

13. You have recently started working as a network administrator for a company whose network consists of multiple Windows Server 2003 domains. The previous network administrator left you with little documentation detailing how the network is configured, and you've discovered that many client workstations are behaving inconsistently—sometimes the Run line is unavailable, sometimes a user cannot access the Control Panel, and so on.

You suspect that this is the result of Group Policy settings, and want to put together a list of all GPOs that are present within each domain on your network. What is the most efficient way of accomplishing this task?

- A. View each domain's settings within the Group Policy Management Console (GPMC) and take note of the values listed under the Group Policy node in each domain.
 - B. Use a GPMC script to list all GPO objects within each domain.
 - C. Load the Resultant Set of Policies (RSOP) snap-in to view the various GPOs that are causing client settings to be applied.
 - D. Examine the Group Policy tab of each domain's Properties sheet in Active Directory Users & Computers.
- B.** A significant advantage to GPMC is that it includes several preconfigured scripts to ease and automate many administrative tasks, one of which will allow you to list all GPOs that are present in a given domain.
- A, C, D.** Answer A is not the most efficient way to determine the total number of GPOs present on your network. Answer C is incorrect because this will indicate only the "winning GPO" that applied each setting; there might be other GPOs that exist elsewhere in the domain structure. Answer D is also inefficient and will require more manual intervention than using an easily automated script like the one suggested in Answer B.
14. A portion of your network is shown in the Figure 1.15. You are using Network Monitor from WorkstationB to capture network traffic for analysis. You suspect that there is an Internet Relay Chat (IRC) connection between WorkstationA and WorkstationC, but the Network Monitor trace does not show any sign of that connection. What is the most likely reason for this?

Figure 1.15 Network Portion



- A. Network Monitor captures broadcast traffic only on a Windows network.
 - B. Windows workstations do not support IRC connections.
 - C. The version of Network Monitor that ships with Windows Server 2003 products does not operate in promiscuous mode.
 - D. You need to use Performance Monitor to capture and analyze network traffic between machines on a Windows network.
- C. In order to capture and analyze all traffic on a network segment, you need to use a network traffic analyzer that can operate in promiscuous mode, such as the one that ships with Microsoft Systems Management Server (SMS). The version included in Windows Server 2003 captures only packets going directly to or from the machine on which it is running.
- A, B, D. Answer A is incorrect because Network Monitor will capture any network packets that are directed towards the network adapter on the machine running the monitor. Answer B is incorrect because Windows workstations support IRC chat and relay applications. Answer D is incorrect because Performance Monitor is used to analyze performance metrics on a given machine, including (but not limited to) processor speed, RAM, paging file, and hard drive usage.

15. Your company, airplanes.com, has recently undergone a merger with southern-airplanes.com, and you have taken over the network management of both halves of the newly formed company. Airplanes.com has a strict policy of desktop and software installation restrictions, while southern-airplanes.com has historically been more lenient with allowing users to customize their computers and install personal software. Several of the users from southern-airplanes.com have complained about the policy restrictions that have been placed on their desktops. You have been asked to present a report to the management group detailing which restrictions are in place on various OUs. What is the most efficient way to present this information to the management group in an easily readable format?
- A. Capture a screen shot of the Properties sheet of the various OUs' Group Policy settings and save the screen shot using a desktop publishing software package.
 - B. Export the GPO settings to a text file, then import the text file into an Excel spreadsheet.
 - C. Demonstrate the use of the Group Policy Editor to apply GPO settings during the meeting with the management group.
 - D. Use the Group Policy Management Console (GPMC) to present the various GPO settings in an organized HTML-formatted report.
- D.** The GPMC presents GPO settings in a unified and easy-to-read format that can be saved as an HTML file and printed or published to a Web page.
- A, B, C.** Answer A is incorrect because an OU's Properties sheet will show only which GPOs have been applied against the OU; it will not list the settings that are in effect. Answer B is not as efficient as using the GPMC to generate a preformatted HTML report of GPO settings. Answer C is incorrect because this will not provide the management group members with the most complete information available to them; they will see only the GPO settings that you demonstrate to them.

Chapter 2: Planning Server Roles and Server Security

Understanding Server Roles

1. Your network consists of two machines running Windows Server 2003 Standard Edition, one machine running Windows Server 2003 Datacenter Edition, one machine running Windows Server 2003 Web Edition, and two machines running Windows Server 2003 Enterprise Edition. You want two of these machines to be domain controllers on the network. Which machines will you promote to domain controllers and how will you configure them in this role?
- A. Configure the two machines running Windows Server 2003 Enterprise Edition to be domain controllers using the `secedit /configure` tool.

- B. Promote the Windows Server 2003 Datacenter Edition and Windows Server 2003 Web Edition using the DCPROMO tool.
 - C. Configure a machine running Windows Server 2003 Standard Edition and a machine running Windows Server 2003 Enterprise Edition to be domain controllers using the Configure Your Server Wizard.
 - D. Configure machines running Windows Server 2003 Standard Edition and Windows Server 2003 Web Edition using the Manage Your Server tool.
 - C. Configure a machine running Windows Server 2003 Standard Edition and a machine running Windows Server 2003 Enterprise Edition to be domain controllers using the Configure Your Server Wizard. The Configure Your Server Wizard allows you to add and remove roles, including the domain controller role. This tool can be used to make servers into domain controllers, as long as the servers are running the Standard Edition, Enterprise Edition, or Datacenter Edition of Windows Server 2003.
 - A, B, D. Answer A is incorrect because `seccedit /configure` is a command-line tool that is used to configure the security settings of a computer. It isn't used to promote member servers to domain controllers. Answer B is incorrect because servers running the Web Edition of Windows Server 2003 cannot be domain controllers. Answer D is incorrect for this same reason. It is also incorrect because the Manage Your Server tool can be used to invoke the Configure Your Server Wizard (which can configure servers be domain controllers), but doesn't actually create the domain controller itself.
2. Your network is upgrading from Windows NT 4 to Windows Server 2003 and will consist of two domains in a single forest. One domain is a child of the other domain and dedicated to the Sales departments in the organization. During the upgrade, all workstations will be upgraded to Windows XP and Windows 2000 Professional. When the last BDC is removed from the network, what role will the PDC emulator play on the network?
- A. The PDC emulator will be used to modify object classes and attributes.
 - B. The PDC emulator will receive preferred replication of password changes performed by other domain controllers in the domain.
 - C. The PDC emulator in the child domain will be used to synchronize the time on all domain controllers in the forest.
 - D. The PDC emulator will be used to add new domains and remove unneeded ones from the forest.
 - B. The PDC emulator will receive preferred replication of password changes performed by other domain controllers in the domain. When a password is changed on a domain controller, 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 domain controllers in a domain.

- A, C, D.** Answer A is incorrect because the schema master is used for making changes to the schema, including modifying classes and their attributes. Answer C is incorrect because, although the PDC emulator synchronizes the time on domain controllers, it only does so within the domain (not the entire forest). The PDC emulator is a domain-wide operations master role and affects only the domain. Answer C is also incorrect because the PDC emulator in a child domain will look to the PDC emulator in the forest root for time synchronization. Answer D is incorrect because the domain naming master is in charge of adding new domains and removing unneeded ones from the forest.
3. The only protocol used by your network is TCP/IP, despite the fact that workstations in the organization do not have access to the Internet. A user has been accessing files on server on your network and now wants to connect to a Web server that is used as part of the company's intranet. The user enters the URL of the Web site into Internet Explorer. Which of the following servers will be used to provide information needed to connect to the Web server?
- A. DHCP server
 - B. DNS server
 - C. WINS server
 - D. File server
- B.** DNS servers map fully qualified domain names (like `www.syngress.com`) to IP addresses. When a user enters a DNS name into a Web browser or other application, it is sent to a DNS server, which looks up the IP address for the requested name. This IP address is sent back to the client, which uses it to locate and communicate with the server.
- A, C, D.** Answer A is incorrect because DHCP servers are used to issue IP addresses to clients. Because TCP/IP is the only protocol used on the network, and the user already has been accessing resources on a file server, this means that the user already has an IP address. Answer C is incorrect because a URL has been entered and WINS servers are used to resolve NetBIOS names to IP addresses (and vice versa). Answer D is incorrect because servers configured in the role of a file server would not need to provide any information to clients accessing an intranet Web site.
4. You want to set up a discussion group that can be accessed over the corporate intranet, so that users can view and post messages in a forum that can be viewed by other employees. Which of the following services would you use to implement this functionality?
- A. HTTP
 - B. FTP
 - C. NNTP
 - D. SMTP

- C.** NNTP is the Network News Transfer Protocol. The NNTP Service in IIS allows users to distribute news messages, which can be viewed using a newsreader program. Users can browse through messages stored on the server, respond to existing messages, and post new messages.
- A, B, D.** Answer A is incorrect because HTTP is the Hypertext Transfer Protocol, which is used by the World Wide Web Publishing Service in IIS. It allows users to access Web pages. Answer B is incorrect because FTP is the File Transfer Protocol. It is used for transferring files between clients and servers. Answer D is incorrect because SMTP is the Simple Mail Transfer Protocol, which is used for transferring e-mail.

Planning a Server Security Strategy

5. You are planning to use a server on your network as a Windows Server 2003 domain controller. The server has 128MB of RAM, 2GB of hard disk space, and four processors. Which of the following editions of Windows Server 2003 can you install on this server? (Select all that apply.)
- A. Windows Server 2003 Standard Edition
 - B. Windows Server 2003 Enterprise Edition
 - C. Windows Server 2003 Datacenter Edition
 - D. Windows Server 2003 Web Edition
- A, B.** Windows Server 2003 Standard Edition and Windows Server 2003 Enterprise Edition both support a computer running a minimum of 128MB of RAM and 1.5GB of hard disk space. The Standard Edition supports up to four processors. The Enterprise Edition supports up to eight processors.
- C, D.** Answer C is incorrect because the Datacenter Edition of Windows Server 2003 requires a minimum of 512MB of RAM. Answer D is incorrect because the Web Edition of Windows Server 2003 supports a maximum of two processors.
6. You are concerned about insecure methods of authentication being used on a network. You are currently upgrading your network to Windows Server 2003, but some servers are still running Windows NT 4 and Windows 2000 Server. Even after the upgrade, some Windows 2000 Server computers will exist in the domain. You want to implement Kerberos authentication within the domain. Which of the following operating systems will be able to use it? (Select all that apply.)
- A. Windows NT 4
 - B. Windows 2000 Server
 - C. Windows Server 2003
 - D. None of the above

- B, C.** Windows 2000 Server and Windows Server 2003 both support Kerberos authentication. Kerberos was first implemented in Windows 2000 and continues to be used in Windows Server 2003 as the default authentication service.
- A, D.** Answer A is incorrect because Kerberos was never supported in Windows NT 4. Answer D is incorrect because Windows 2000 and Windows Server 2003 both support Kerberos authentication.
7. Your network consists of two Windows Server 2003 domain controllers, a Windows 2000 server that is used as a Web server, and a Windows NT 4 server that runs an older version of SQL Server. Your company does not have the budget to immediately replace these servers, but you want to raise the domain functional level of your domain to the highest possible level. What functional level will you raise this domain to?
- A. Windows 2000 mixed
B. Windows 2000 native
C. Windows Server 2003 interim
D. Windows Server 2003
- D.** Because the only servers being used as domain controllers are running Windows Server 2003, the domain can be raised to the Windows Server 2003 domain functional level. The Windows Server 2003 level is used when there are only Windows Server 2003 domain controllers in the domain.
- A, B, C.** Answer A is incorrect because the Windows 2000 mixed level is used when there are Windows NT, Windows 2000, and Windows Server 2003 domain controllers. Because there are no Windows NT BDCs or Windows 2000 domain controllers, this isn't the highest level that can be used. Answer B is incorrect because this level is used when there are only Windows 2000 and Windows Server 2003 domain controllers. Because there are not any Windows 2000 domain controllers, a higher level can be used. Answer C is incorrect because Windows Server 2003 interim is used when your domain consists of Windows NT and Windows Server 2003 domain controllers, and you are upgrading Windows NT domains directly to Windows Server 2003. Because there are not any Windows NT BDCs, this isn't the highest level that can be used.

Planning Baseline Security

8. You have just promoted a Windows Server 2003 computer to be a domain controller. After the promotion, you accidentally apply the wrong security template to it. It now has security settings that are too high. You can automatically change the security settings back to their previous configuration using which of the following security templates?

- A. Setup security
- B. Rootsec
- C. IesacIs
- D. DC security

D. The DC security template is created when a server is first promoted to a domain controller, and it contains default settings for the file system, Registry, and system services. Applying this template will restore the settings to the state they were in after the server was first promoted.

A, B, C. Answer A is incorrect because the setup security template allows you to reapply default security settings on either clients or servers, but should not be used on servers that have been configured as domain controllers. Answer B is incorrect because rootsec is a template that is used to define settings for root of the system volume. Answer C is incorrect because iesacIs is a template that is used to define settings to lock down Internet Explorer.

9. You want to apply an existing security template to the local computer policy of a Windows Server 2003 computer. Which of the following tools would allow you to do this from the command line?

- A. Security Configuration and Analysis
- B. `secedit /configure`
- C. `secedit /import`
- D. `gpupdate`

B. The Secedit tool is a command-line utility that can be used to apply configuration settings stored in a security template to a local computer policy. To apply a policy, use the `secedit` command with the `/configure` switch.

A, C, D. Answer A is incorrect because the Security Configuration and Analysis tool is a graphical utility. Although it can be used to apply security templates to local computer policy, the question states that a command-line tool is required. Answer C is incorrect because `secedit /import` is used to import a template into the database so that it can be used to either analyze security on the machine or configure its security settings. Answer D is incorrect because `gpupdate` is a command that is used to trigger an update of GPO settings.

10. You have performed an analysis of a Windows Server 2003 domain controller using Security Configuration and Analysis. Once the analysis is complete, a red X appears beside the Enforce Password History policy. What does this mean?

- A. The policy does not match a corresponding setting for the associated entry in the database.

- B. The entry in the database and the policy's setting match.
 - C. An entry exists in the database that does not correspond to any setting on the computer.
 - D. A setting exists on the computer that does not correspond to any entry in the database.
 - A.** The policy does not match a corresponding setting for the associated entry in the database. Although the Enforce Password History entry exists in the database, the value of the entry is different from what is currently configured in the policy.
 - B, C, D.** Answer B is incorrect because a green check mark indicates that the entry in the database and the computer's setting match. Answer C is incorrect because an exclamation mark indicates that an entry in the database does not correspond to any setting on the computer. Answer D is incorrect because a question mark indicates that the setting is on the computer, but there is no corresponding entry in the database.
11. You have created a security template and now want to apply its settings to a GPO that can be linked to containers in Active Directory. Which containers can you link a GPO to in Active Directory? (Select all that apply.)
- A. Domains
 - B. Trusts
 - C. Sites
 - D. Local computer policy
 - A, C.** Security templates can be imported into GPOs in Active Directory. These GPOs can be linked to domains, sites, or OUs in the Active Directory structure.
 - B, D.** Answer B is incorrect because trusts cannot have group policies applied to them. Answer D is incorrect because the local computer policy is stored on the machine and cannot have a GPO linked to it.

Customizing Server Security

12. You have installed a new file server on the network and formatted it to use NTFS. After formatting is complete, you use EFS to encrypt a folder containing files belonging to users. If a user accesses a file belonging to him in this folder, and then copies it across the network for another user to access, which of the following will occur?
- A. The file on the hard disk and the data sent over the network will remain encrypted.
 - B. The file on the hard disk and the data sent over the network will be decrypted and remain that way.

- C. The file on the hard disk will be decrypted, so EFS can send it encrypted over the network.
- D. The file on the hard disk will remain encrypted, but data sent over the network will be unencrypted.
- D.** The file on the hard disk will remain encrypted, but data sent over the network will be unencrypted. EFS only encrypts data on NTFS volumes. When data that is encrypted with EFS is sent over the network, it isn't encrypted. For data to be encrypted during transmission, other methods like IPSec are needed.
- A, B, C.** Answer A is incorrect because EFS only encrypts data on hard disks. It does not encrypt data transmitted over the network. Answer B is incorrect because when a file is transmitted over the network, the original file on the hard disk isn't decrypted and left that way. EFS will keep the file on the hard disk encrypted, so others cannot access it. Answer C is incorrect, because EFS isn't used for transmitting encrypted data over the network.
13. You have created a custom security template that you now want to import into a GPO that is linked to the domain level. Which of the following tools will you use to invoke the Group Policy Object Editor to view and modify the GPO at this level?
- A. Active Directory Users and Computers
- B. Active Directory Sites and Services
- C. gpupdate
- D. Securedc
- A.** Active Directory Users and Computers is used to view GPOs linked at this level. Active Directory Users and Computers can then be used to invoke the Group Policy Object Editor, where you can import security templates into group policies at the domain and OU levels.
- B, C, D.** Answer B is incorrect because Active Directory Sites and Services is used to access GPOs at the site level and can be used to invoke the Group Policy Object Editor to edit these objects. Answer C is incorrect because gpupdate is used to refresh group policies on Windows Server 2003. Answer D is incorrect because securedc is a security template that can be applied to domain controllers.
14. Your network consists of servers running Windows 2003 Server and workstations running Windows 2000 Professional. You have applied several custom security templates to GPOs linked to the OU, domain, and site levels in Active Directory. In addition to this, there are security settings that have also been applied at the local computer level of all machines that are on the network. Because several policies now affect the computer accounts within the domain, site, and OU, which of the following will occur when the user logs on to the domain?

- A. The policy setting at the local computer level will be overwritten by the OU-level GPO, which will be overwritten by the domain-level GPO, which will finally be overwritten by the site-level GPO. For this reason, major security settings must be made at the site-level GPO; all others will be overwritten.
 - B. Security settings in the GPOs will not be applied to machines running Windows 2000 that have joined the domain.
 - C. The security settings at the local computer level will override those of the GPOs.
 - D. The policy settings will be cumulative and applied in the order of policies at the site level, domain level, and finally OU level.
 - D.** The policy settings will be cumulative and applied in the order of policies at the site level, domain level, and finally OU level.
 - A, B, C.** Answer A is incorrect because policy settings are cumulative and applied in the following order to computer accounts: site-level GPOs, domain-level GPOs, OU- and sub-OU level GPOs. Answer B is incorrect because GPOs can be applied to any Windows 2000 or later computer that has joined a domain. Answer C is incorrect because security settings configured in GPOs override those made at the local computer level.
15. You apply custom security templates to the local computer policy on a member server and to a GPO linked to an OU in Active Directory. All servers on the network are running Windows Server 2003. After performing these actions, you find that the local computer policy has taken effect, but the group policy has not taken effect on member servers within the domain. Which of the following is the reason for this, and how can you fix it?
- A. Group policy settings take effect immediately. The problem must be that the security policy was not applied properly.
 - B. Group policy settings are refreshed on member servers every 90 minutes. To force the server to refresh the group policy, use the `secedit /refresh` command.
 - C. Group policy settings are refreshed on servers every 5 minutes. To force the server to refresh the group policy, use the `gpupdate` command.
 - D. Group policy settings are refreshed on servers every 90 minutes. To force the server to refresh the group policy, use the `gpupdate` command.
 - D.** Group policy settings are refreshed on servers every 5 minutes. To force the server to refresh the group policy, use the `gpupdate` command. Local computer policies are stored on the computer, and they take effect immediately. Group policy settings are stored in Active Directory and need to be downloaded to the machine. Because of this, the group policy settings are refreshed at regular intervals. To force a refresh, the `gpupdate` command can be used.

- A, B, C.** Answer A is incorrect because group policy settings do not take effect immediately. The group policy settings are refreshed on computers at regular intervals. Workstations have group policy settings refreshed every 90 minutes, member servers are refreshed every 90 minutes, and domain controllers are refreshed every 5 minutes. Answer B is incorrect because the `secdit /refresh` command isn't used in Windows Server 2003. It has been replaced by the `gpupdate` command. Answer C is incorrect because member servers are refreshed every 90 minutes. Domain controllers are refreshed every 5 minutes.

Chapter 3: Planning, Implementing, and Maintaining the TCP/IP Infrastructure

Understanding Windows 2003 Server Network Protocols

1. You are implementing a network that will include UNIX workstations that will share files and information with the Windows users. What protocols will you need to implement to provide integration with UNIX machines?
 - A. IPX/SPX
 - B. NetBEUI
 - C. TCP/IP
 - D. NetBIOS over TCP/IP

C. UNIX uses TCP/IP by default, and since the TCP/IP suite of protocols is standard across platforms, HTTP and FTP can be used to share information and files.

A, B, D. IPX/SPX is used with NWLink and NetWare servers. NetBEUI is no longer supported on Windows Server 2003. NetBIOS is used for NetBIOS name resolution, which is not required to access UNIX machines. It is possible to implement SAMBA, a service that provides Server Message Block (SMB) and NetBIOS encapsulation over TCP/IP on the UNIX clients, which will enable them to use NetBIOS name resolution and function as NetBIOS clients.
2. You purchased a new desktop computer running Windows XP for your small office and a server running Windows Server 2003. Your old desktop is running Windows 95. It has a network adapter and can access files on another Windows 95 machine. The Windows XP machine has not arrived, but you want to back up the data from the Windows 95 computer to the Windows Server 2003 machine. However, from the Windows Server 2003 computer, you are unable to see the shares on the Windows 95 computer. What should you do to allow the Windows Server 2003 machine to access the Windows 95 machine?

- A. Install NetBEUI on Windows Server 2003 computer.
 - B. Install NWLink on the Windows 95 client.
 - C. Install TCP/IP on the Windows 95 client.
 - D. Ensure the server has a valid IP address and implement a DHCP server on the Windows Server 2003 machine with a valid scope.
- C, D.** This solution will provide a means for the TCP/IP client on Windows 95 to obtain a valid IP address and communicate with the server. Windows 95 does not support APIPA features, and it uses NetBEUI to access workgroup computers.
- A, B.** NetBEUI is not supported on Windows Server 2003. NWLink will not communicate with the Windows Server 2003 machine, unless it also has NWLink installed, which is not the case by default.

Planning an IP Addressing Strategy

3. You are implementing a test lab that contains three Windows Server 2003 machines, twenty Windows XP Professional machines, and two IP-based printers. You have been given the network address of 155.1.50.0 and a subnet mask of 255.255.255.224. What is the CIDR notation for your subnet?
- A. 155.1.50.0/27
 - B. 155.1.50.0/5
 - C. 155.1.50.0/24
 - D. 155.1.50.0/3
- A.** The subnet mask is 255.255.255.224, which equals 11111111.11111111.11111111.11100000, for a total of 27 mask bits.
- B, C, D.** Answer B is 11111000.00000000.00000000.00000000, or 248.0.0.0, which is a supernet containing 134,217,726 hosts per network. C is 11111111.11111111.11111111.00000000, or 255.255.255.0. D is 11100000.00000000.00000000.00000000, or 192.0.0.0, which is a supernet containing 536,870,912 hosts per network.
4. You are given a task to create eight subnets on your LAN, and you have been assigned the address space 172.16.128.0/23. How many hosts will you have and what is the CIDR notation for the new subnet's address space?
- A. 2032 hosts on 172.16.128.0/24
 - B. 240 hosts on 172.16.128.0/27
 - C. 496 hosts on 172.16.128.0/26
 - D. 48 hosts on 172.16.128.0/29

- C.** You need eight subnets, so you need 3 bits (111 binary = 8). Since the current subnet is using 23 bits, you would add 3 to that for your subnet, which is 26 bits, or /26. The remaining number of bits for hosts is 6. $2^6 - 2 = 62$, and you have eight subnets, so $62 * 8 = 496$.
- A, B, D.** Answer A is incorrect because it provides only one more subnet. Answers B and D are incorrect because you were instructed to create eight subnets, and these options create more than eight subnets.
5. Which of the following addresses is suitable for dividing into at least nine subnets, each with the ability to support 200 hosts per network?
- A. 10.1.1.0/24
 B. 10.1.1.0/20
 C. 10.1.1.0/19
 D. 10.1.1.0/22
- B, C.** You will need 4 bits to expand your subnet mask to divide the network into at least nine subnets, since $9 = 1001$. You need at least 200 hosts per network, so you need at least 8 bits for the host IDs, since $200 = 11001000$. If you add $8 + 4$ bits, you get a total of 12 bits needed as a minimum to support the required architecture. $32 - 12 = 20$, so any CIDR notation with /20 or less can be used to meet these requirements.
- A, D.** These addresses don't leave enough bits available to meet the requirement for 200 hosts. Answer A leaves 4 bits (4 are needed for the subnet), for $2^4 - 2$, or 14 hosts per network. Answer D leaves 6 bits, which is $2^6 - 2$, or 60 hosts per network.
6. You are having trouble accessing Microsoft's Web site. When you ping www.microsoft.com, the request times out. How should you proceed in troubleshooting this problem?
- A. Ping the loopback adapter, the IP address of this machine, then the default gateway and determine if your connectivity is valid. If there are no issues, run tracert and identify where the communications stop.
 B. Ping the default gateway, the IP address of a remote host other than Microsoft, such as Yahoo, then ping the IP address of this machine and then the loopback adapter.
 C. Use Network Monitor to analyze the traffic to www.microsoft.com.
 D. Use System Monitor to look at counters on the local machine to determine the error.
- A.** The order in which you should ping to troubleshoot your TCP/IP configuration is always from the loopback adapter outward.
- B, C, D.** Answer B is incorrect because the ping tasks were not in the right order to best isolate the problem. Answers C and D are wrong because they don't take into account the local LAN traffic and its routes.

7. You implement a Windows Server 2003 machine that is functioning as a file server on your LAN. The server name is FileServer01. Users attempting to browse the shares on \\FileServer01\ are unable to see any of the shares you created. What is likely the problem?
- A. You do not have DNS installed on the LAN.
 - B. DHCP is unavailable.
 - C. NetBIOS encapsulation is not enabled on the Windows Server 2003 machine.
 - D. FileServer01 FTP service is stopped.
- C. NetBIOS encapsulation is required to browse file shares on Windows Server 2003 machines.
- A, B, D. DNS is used for host name-to-IP address resolution and is not related to browsing resources. DHCP provides IP addresses and configuration information. Although it could provide a WINS server IP address, it would not necessarily solve this problem. FTP services do not provide browser services.
8. A client computer configured as a DHCP client was unable to obtain an address from the DHCP server. Upon investigation, you discovered that the DHCP scope was not activated, so you activated it. The client computer has an APIPA address of 169.254.0.1. What actions are required for the client to obtain an IP address from the DHCP server?
- A. Run ipconfig /all from a command prompt.
 - B. Use Netsh to assign an address to the network adapter.
 - C. Log off Windows XP and log on again.
 - D. Take no action.
- D. When a DHCP client fails to obtain an address, it will continue to request an address every five minutes, until one is obtained.
- A, B, C. Ipconfig /all will only display the current configuration. If ipconfig /renew is run, it would initiate the request immediately, although it is not required. Using Netsh to assign an address would defeat the purpose of having a DHCP client. Logging off Windows XP and logging on has no effect on obtaining a DHCP lease.

Planning the Network Topology

9. Your company is merging with another organization, and you have been tasked with merging the corporate networks. You have determined that the other company has between 50 and 125 hosts on 7 networks. Your company has 25 to 50 hosts on 12 networks. You want the integration to provide room for five percent growth over the next two years. Your routers do not support variable-length subnet masks. You decide to use the private address 192.168.0.0. What is the best subnet mask for your new corporate LAN?

- A. 255.255.0.0
- B. 255.255.255.0
- C. 255.255.255.192
- D. 255.255.224.0

B. You have a maximum total of 1475 hosts ($875 + 600$) with an approximate growth of 74 each year for two years, for a total of $148 + 1475 = 1623$ hosts on 19 networks. Using the private class B address with the subnet mask of 255.255.255.0 allows you to create 255 networks, each with up to 254 hosts.

A, C, D. Answer A is incorrect because it would provide only one network. Answer C is incorrect because the result would give you too few networks. Answer D is incorrect because it would give you too many hosts on too many networks.

10. You want to simplify the configuration and management of TCP/IP clients on your network, which consists of 300 Windows XP Professional machines, 12 Windows Server 2003 machines, and 23 printers on four subnets. Which of the following solutions best suits your needs?

- A. Implement WINS using APIPA. Provide at least one DNS server for each WINS server.
- B. Implement DHCP to provide assigned IP address leases and scope properties that contain the necessary host resolution methods, the IP address of the default gateway, and the DNS servers.
- C. Implement AD integrated DNS and WINS and configure WINS to do reverse lookups.
- D. Provide thorough documentation for each client to manually configure its IP address with a valid subnet mask and DNS server.

B. DHCP is the default option for Windows XP and Windows Server 2003; therefore, if you implemented DHCP, you would need to maintain only the DHCP server. All the clients would automatically obtain the configuration from DHCP.

A, C, D. Answer A is incorrect because APIPA is suitable for only small networks. It doesn't provide a default gateway or support DNS, which is used for host name resolution. WINS is used for NetBIOS name resolution, but requires that the client be configured with the WINS server IP address, which is not part of APIPA. Answer C is incorrect because AD does not assign IP addresses. Answer D is incorrect because it requires manual configuration of every client machine on the network, which is prone to mistakes and not centrally managed.

11. All of the clients on your network are configured to use DHCP for their TCP/IP configuration. You upgrade Internet access to use a T1 line that is connected to a different router

than the current router that is being used by the Digital Subscriber Line (DSL) connection. What actions are required to allow the executive staff to access the Internet using the new default gateway, by configuring each executive's machine only one time, while not allowing the other company employees to use the T1?

- A. Create a logon script for the Executives Group that uses the route add -d command to add the new router information. Set the script to run every time members of the Executive Group log on.
 - B. Create a logon script for the Executives Group that uses the route add -p command to add the new router information. Set the script to run once the next time members of the Executive Group log on.
 - C. Create a new property for the router in the DHCP scope options. Set up reservations for each of the executive's machines.
 - D. Run the command route add with the information for the new router on each executive's machine.
- B.** Using the -p switch with the route add command will allow you to persist the route in the routing table. This will provide a manual configuration for the executives that will override the gateway provided by DHCP and thus allow normal DHCP operations.
- A, C, D.** Answer A is incorrect because there is no route add -d command. The command to delete routes is **route delete *Destination***, where *Destination* is the network destination you want to remove from the routing table. Answer C is incorrect because you are updating the scope, so any machine that obtains a DHCP lease from that scope will use that gateway. Defining reservations for the executive machines will simply ensure that they get the same IP address from the scope. Answer D is incorrect because the route would not be persisted, so it would require multiple repetitive configurations to address.
12. You have integrated a smaller LAN into your network that contains a Novell NetWare server using IPX/SPX. You want to be able to access it from a Windows Server 2003 machine, so you install NWLink. You notice that after you installed NWLink, the Windows XP client machines that connect to Windows Server 2003 are taking longer to connect and read information. What can you do to ensure the best performance for the Windows XP clients?
- A. Install NWLink on the Windows XP machines.
 - B. Install the Novell NetWare Client on the Windows XP machines.
 - C. Move TCP/IP up in the binding order on the Windows Server 2003 machine.
 - D. Install the Novel NetWare Client on the Windows Server 2003 machine.

- C.** By moving the most used protocols up to the top of the stack, you will force that protocol to respond to client requests first and then attempt additional protocols. If the wrong protocol is at the top of the stack, the clients will first attempt to use the protocol that they don't have, fail, and then try the next protocol in the stack.
- A, B, D.** Answers A and C are incorrect because the Windows XP machines do not need to connect to the Novell NetWare server. Answer D is incorrect since it changes only the security settings and protocols on the Windows Server 2003 machine and does not address the issue with the Windows XP connectivity.
13. You are network administrator for a new company. Your LAN is connected to the Internet by a single T1 line. You obtain a single public IP address from your ISP. Your firewall services are outsourced to the ISP. The LAN includes five Windows XP Professional computers and one Windows Server 2003 computer named Server01. All Windows XP client computers are configured to use DHCP to obtain their IP configurations. Server01 is configured as a DHCP server and contains two network adapters. You connect one network adapter to the hardware for the ISP connection and connect the other network adapter to the LAN. You want client computers to access the Internet, including browsing the Web and file transfers via FTP. Which of the following configuration tasks must you complete?
- A. Install the DNS Server service.
- B. Install WINS Services.
- C. Install Routing and Remote Access Services (RRAS).
- D. Assign the public IP address to the external adapter.
- A, C, D.** Since you have only one external IP address, you must assign it to the interface that connects you with the ISP. Only one external address is required. It is also necessary to install RRAS to provide a means for NAT for the private IP addresses on the LAN to map to the external IP address, and to route Internet traffic to the ISP. DNS is required for resolving uniform resource locators (URLs). You may have been provided IP addresses for DNS servers by your ISP, which could be used instead of installing your own DNS server, but you would not be able to use DNS for internal host name resolution.
- B.** WINS is used for NetBIOS name resolution and is not required to access the Internet.

Planning Network Traffic Management

14. Users are complaining about slow network performance. Using Network Monitor, you have identified the source of the excessive traffic is inbound and outbound traffic from your DNS server. How would you identify the source of the excessive DNS traffic?

- A. Using the host IP addresses from Network Monitor, perform a tracert command to each host and determine the time it takes to get to each requested destination.
 - B. Use System Monitor to watch performance counters on the DNS server and identify the cause of the slow performance.
 - C. Use System Monitor to watch performance counters on the client machines to identify the machine that is using the DNS server heavily.
 - D. Ping the DNS server using the `-t` option from different host machines to identify the subnet that is causing the increase in network traffic.
- B.** Using System Monitor, you can identify the problem areas and zero in on the exact operations in DNS that are causing the traffic.
- A, C, D.** Using Tracert is incorrect because it provides information only on routing latency and does not address the issue you have identified with the DNS server. System Monitor running on the client machine may provide some insight related to each client, but the source of the network utilization problem has been identified as the DNS server. Ping `-t` would only serve to increase the network load and would do so continuously.
15. You are using Network Monitor to analyze traffic on your Windows Server 2003 machine. You have a lot of data that has been captured, but you are looking for specific information. How do you accomplish this?
- A. Define a filter for the captured data.
 - B. Open the trace in Notepad and do a global search for the information you are seeking.
 - C. Export the data to a `.cap` file and view the reports in Excel.
 - D. Set up the counters for the appropriate data.
- A.** Network Monitor allows you to apply filters to the captured data.
- B, C, D.** Answers B and C are incorrect because the `.cap` files are binary and can be viewed only in Network Monitor. Answer D is incorrect because counters are part of System Monitor.

Chapter 4: Planning, Implementing, and Maintaining a Routing Strategy

Understanding IP Routing

1. Your IT Director has decided the new internal network needs to use private addressing. Which of the following IP addresses are private addresses?

- A. 193.168.0.1
- B. 171.17.0.1
- C. 10.0.0.1
- D. 172.16.0.15

C, D. The three blocks of private address space are 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. Answer C, 10.0.0.1, falls in this range, as does Answer D, 172.16.0.15.

A, B. Neither of these IP addresses fall within any of the three blocks of address spaces assigned to private addressing. Remember, the three address blocks set aside and defined as private address space are as follows:

- **10.0.0.0 with a subnet mask of 255.0.0.0, or 10.0.0.0/8** This network is a private address space that has 24 host bits that can be used.
- **172.16.0.0 with a subnet mask of 255.240.0.0, or 172.16.0.0/12** This network is a private address space that has 20 host bits that can be used. This provides a range of 16 class B network IDs from 172.0.0.0/16 through 172.31.0.0/16.
- **192.168.0.0 with a subnet mask of 255.255.0.0, 192.168.0.0/16** This network is a private address space that has 16 host bits that can be used. This provides a range of 256 class C network IDs from 192.168.0.0/24 through 192.168.255.0/24.

2. Your IT Director has determined that your network should use dynamic routing. You've determined that a route is now being considered unreachable. What has happened to that route in the routing table?

- A. It has been marked as unreachable in the routing table.
- B. Nothing has happened to that route in the routing table.
- C. It has been removed from the routing table.
- D. You must manually go into the routing table and remove the entry.

C. In dynamic routing, when a route is unreachable, the route is removed from the routing table.

A, B, D. Because it is removed from the routing table, it is not marked as unreachable, so both answers **A** and **B** are incorrect. Answer **D** is incorrect because this is dynamic routing, not static, so you don't need to manually enter or change the router table entry.

3. Your newest hire has been assigned the task of configuring a Windows Server 2003 computer as a router and has asked you how to determine if a machine address or an IP address is being used at the router. You explain that routers use IP addresses, while bridges and hubs use machine addresses. You continue to explain that the OSI reference model has seven layers and that IP, or the Internet Protocol, operates at what layer?

- A. The Physical layer
 - B. The Data Link layer
 - C. The Network layer
 - D. The Transport layer
- C.** The Network layer implements protocols that can transport data across a LAN segment. These protocols are known as routable protocols because their data can be forwarded by routers past the local network. IP is the dominant routable protocol. IP, and other Layer 3 protocols, are considered asynchronous because they send the data with no attempt to verify that the data arrived at its destination.
- A, B, C.** Answer **A** is incorrect because the Physical layer of the OSI reference model is responsible for the transmission of data. This layer operates only with ones and zeros. Answer **B** is incorrect because the Data Link layer is responsible for providing end-to-end validity of the data being transmitted. Answer **D** is incorrect because the Transport layer is also responsible for the end-to-end integrity of data transmissions. An example of a protocol used at the Transport layer is TCP. TCP and other Layer 4 protocols are considered synchronous because they verify the successful arrival of the data at its destination.
4. Your IT Director has opened a command prompt window on your Windows Server 2003 computer and is trying to figure out what routes are available to this computer. Which of the following commands should you tell him to use to list the active routes from the command prompt?
- A. route list
 - B. route print
 - C. show route
 - D. dump
- B.** To produce a list of the active routes from the command prompt, type **route print** and press the **Enter** key.
- A, C, D.** None of these are route commands. They are Netsh commands.
5. Your IT Director is determined to use static routing on your large corporate network. You need to convince him that static routing probably is not the best choice, and you want him to think that decision was his idea. You decide to do this by asking him which of the following is an advantage of using static routing?
- A. Fault tolerance
 - B. Scalability
 - C. Manual configuration
 - D. Classless routing

- D.** Static routing works well with classless routing because each route must be added with a network mask.
- A, B, C.** Static routing is not fault tolerant. Although it works well for small networks, it does not scale well. It requires manual configuration and makes no attempt at discovery of other networks or other systems on the network.
6. RRAS is enabled on your Windows Server 2003 computer, and you have three network adapter cards in the computer configured for subnet IDs of 192.168.32.0/20, 192.168.64.0/20, and 192.168.96.0/20. Which subnet ID can you use if you need to support another subnet with this RRAS server?
- A. 192.168.20.0/20
- B. 192.168.40.0/20
- C. 192.168.48.0/20
- D. 192.168.60.0/20
- C.** The subnet mask of 20 bits means that the first two octets and the first four bits of the third octet are used to define the subnet ID. The rightmost bit of the four bits used in the third octet represents a value of 16. Each subnet ID must therefore have a value in the third octet that is divisible by 16. The only one in the list that meets this criterion is Answer A.
- A, B, D** These answers are incorrect because none of them meets the criterion. In Answer A, 192.168.20.0/20, the 20 in the third octet is not divisible by 16. The same situation exists for Answer B, 192.168.40.0/20. The 40 in the third octet is not divisible by 16. Finally, in Answer D, 192.168.60.0/20, the 60 in the third octet is not divisible by 16.
7. You want to configure a multiple gateway on a Windows Server 2003 machine, but you have only one NIC installed. How do you accomplish this goal?
- A. Assign the IP addresses 192.168.0.10 and 192.168.1.10 to the interface.
- B. Assign the IP addresses 10.0.0.1 and 172.16.0.1 to the interface.
- C. Assign the IP addresses 172.16.0.1 and 192.168.0.1 to the interface.
- D. You cannot configure multiple gateways on a machine with one NIC.
- A.** When using a single NIC, the IP addresses must be assigned to either the same network segment or to segments that are part of the same single logical network. Answer A is the only answer that meets this criterion.
- B, C, D.** In Answer B, the two network addresses, 10.0.0.1 and 172.16.0.1, to be assigned to the interface are located on two different logical networks. The same holds true for Answer C. Here, the addresses 172.16.0.1 and 192.168.0.1 are also on separate and different logical networks. A single NIC can be used when the IP addresses

assigned are from the same network segment or to segments that are part of the same single logical network. This statement rules out Answer D.

8. Your IT Director has been reading again. He has decided that he wants to convert the network to OSPF, but he is having some difficulty with terminology. He knows that an OSPF router can serve one of four roles. His problem is that he can't remember which role exists when one of the router's interfaces is on the backbone area. Help him out. Which of the following is it?
- A. Internal router
 - B. Area border router
 - C. Backbone router
 - D. Autonomous system boundary router
- C.** If one of a router's interfaces is on the backbone area, that router is considered a backbone router.
- A, B, D.** When all interfaces are connected to the same area, the router is considered an internal router. This rules out Answer **A**. When a router's interfaces are connected to different areas, that router is an area border router (ABR). This rules out Answer **B**. When the router exchanges routes with sources outside the network area, it is known as an autonomous system boundary router (ASBR). This rules out Answer **D**.

Security Considerations for Routing

9. As the network administrator, you are asked to set up network access so that a group of contract developers can work via a VPN connection connecting to your network's Windows Server 2003 VPN server. The contract developers are all using either Windows 2000 Professional or Windows XP Professional workstations. You must meet the following requirements:
- The contract developers must be allowed to connect to the network via the Internet.
 - You must use PPP encryption.
 - You must use a protocol that provides tunnel authentication.
 - You must use a protocol that secures the data between the endpoints of the tunnel.

You configure a VPN using PPTP. Which of requirements are met? (Select all that apply.)

- A. The contract developers are able to connect to the network via the Internet.
- B. PPP encryption is used.
- C. Tunnel authentication is used.

- D. Data between the endpoints of the tunnel is secure.
- A, B, D.** Using a VPN allows these developers to connect to the local network in a secure manner. The VPN allows users to exchange data between computers in the network as if there were a point-to-point private link between them. In order for this to succeed, a protocol such as PPTP must be used to encapsulate the PPP frames. PPTP creates the tunnel and uses a modified version of GRE to encapsulate the PPP frames as tunneled data.
- C.** Unlike L2TP, PPTP does not support tunnel authentication.
10. You have enabled RRAS on your Windows Server 2003 computer. You want to set up IP packet filtering to help you manage access from remote clients. Where in the Routing and Remote Access console will you enable IP packet filters?
- A. The properties of the remote-access ports
- B. The properties of the remote-access server
- C. The profile of a remote-access policy
- D. The conditions of a remote-access policy
- C.** IP packet filters are managed in the remote-access profile of a remote-access policy. The ability to request an IP address, IP packet filters, idle time allowed before being disconnected, and length of session are settings defined in the remote-access profile.
- A, B, D.** IP address assignment management cannot be performed via the properties sheets of either the remote-access server or the remote-access ports or via the conditions of a remote-access policy.
11. You have set up an isolated, secure subnet with only an RRAS server running on Windows Server 2003 connecting the two parts of your internal network. You are protecting your internal network against unauthorized access with your firewall, and authorized users on the intranet establish VPN tunnels to your secure subnet through the RRAS server. You do have a problem, however. It seems that remote VPN clients cannot access the secure subnet through your configuration. How should you reconfigure the system to allow remote VPN clients access to the secure subnet?
- A. Ask your ISP to create the necessary filters to allow IPSec traffic to pass.
- B. Create filters on the RRAS server to allow only VPN traffic to pass.
- C. Define filters on the firewall to allow the VPN traffic to pass.
- D. Configure the router in front of the firewall to allow IPSec traffic to pass.
- C.** The most likely reason that VPN traffic is unable to access the secure subnet through the RRAS server is that the firewall isn't configured to allow VPN traffic to pass from the Internet. Correct the problem by configuring filters on the firewall to allow this traffic to pass.

- A, B, D.** Data packet transmission is transparent to all hosts between the source and the recipient. This includes all routers on the ISP's network and any router you might have in front of the firewall. As a result, Answer A, asking your ISP to create the necessary filters to allow IPsec traffic to pass, isn't correct because your ISP's routers are between the source and the destination and therefore transparent. Because internal VPN traffic is occurring, you know that filters are already created on the RRAS server allowing the VPN traffic to pass. This means Answer B, create filters on the RRAS server to allow only VPN traffic to pass, is incorrect also. Finally, Answer D, configure the router in front of the firewall to allow the VPN traffic to pass, is incorrect because the traffic is transparent between the source and the recipient.
12. You've been asked to provide Internet access for clients on your network. You decide to use NAT. You try to establish a secure VPN session from a remote site unsuccessfully. You try again using L2TP. Again the connection fails. You are able to successfully connect when in the same office. Why are you unable to make a connection from the remote location?
- A. You haven't configured the NAT server to translate the IP Security packets.
 - B. You cannot establish an L2TP connection behind a computer running NAT. The L2TP session fails because the IP Security packets become corrupted.
 - C. L2PT does not work with Windows Server 2003 VPNs.
 - D. NAT does not allow for remote networking.
- B.** You cannot use NAT with L2TP.
- A, C, D.** Answer A, regarding not having configured the NAT server to translate the IP Security packets, is incorrect because there is nothing to configure. L2PT does work with Windows Server 2003 VPNs, so Answer C is incorrect. Finally, Answer D stipulates that NAT doesn't allow for remote networking. This is incorrect, because it does. It just doesn't allow for using L2PT security.
13. You've just been asked to set up things so that a group of developers can work from home and still connect to your office network. The developers are using either Windows 2000 Professional or Windows XP Professional. You must meet the following requirements:
- Allow the developers to connect to the network through the Internet.
 - Use PPTP encryption.
 - Use a protocol that provides tunnel authentication.
 - Use a protocol that secures data between the endpoints of the tunnel.

You plan to configure a VPN that uses L2TP. Which requirement or requirements are met?

- A. The developers can connect to the network through the Internet.
 - B. PPTP encryption is used.
 - C. Tunnel authentication is provided.
 - D. Data between the endpoints of the tunnel is secured.
- A, C.** The VPN allows the developers to work from home over the Internet. L2TP is a combination of PPTP and Layer 2 forwarding that can be used as a tunneling protocol.
- B, D.** L2TP doesn't use PPTP, but rather uses IPSec to encrypt data. This means that the data between the two endpoints of the tunnel will not be secure in this situation.

Troubleshooting IP Routing

14. You've installed RRAS on a Windows Server 2003 computer in your network. The network is not connected directly to the Internet, and the private IP address range you are using is 192.168.0.0. When you dial in, you connect successfully, but you're unable to access any resources. Pinging other servers using their IP addresses results in the message "Request timed out." Running the ipconfig command shows you that your dial-up connection is being given the IP address 169.254.75.182. What should you do to resolve the problem?
- A. Configure the remote-access server to act as a DHCP Relay Agent.
 - B. Ensure that the remote-access server is able to connect to a DHCP server that has a scope for its subnet.
 - C. Configure the remote-access server with the address of a DHCP server.
 - D. Authorize the remote-access server to receive multiple addresses from a DHCP server.
- B.** Your dial-up connection is being assigned a default IP address because it is unable to obtain an assigned IP address from a DHCP server. This is because the remote-access server is unable to connect to a DHCP server that has the proper scope.
- Answers A, C, D.** The IP address being assigned to the dial-up connection, 169.254.75.182, is an automatically assigned IP address that computers that have either Windows 2000 Professional or Windows XP Professional installed will assign themselves when no DHCP server is available or can be contacted. Answer A is incorrect because, unless the remote-access server can connect to a DHCP server in the first place, being able to relay DHCP information will be of no assistance in this situation. Answer C is incorrect because you don't assign a computer the address of a DHCP server. DHCP works through broadcast. Finally, Answer D is incorrect because a single network interface cannot receive multiple addresses from a DHCP server. Even if it could, if the machine is not receiving DHCP broadcasts, it still won't be assigned an address, and the automatic assignment will still take place.

15. You think you may have a problem on your network. You need to open a command line window and troubleshoot your network. Which of the following lists of commands represent the command-line utilities most often used in maintaining and testing routing functionality?
- A. show helpers, Trace, PING, Route
 - B. pathping, Tracert, show helpers, show routing
 - C. pathping, PING, Route, Tracert
 - D. pathping, PING, Route, Trace
- C. The four commonly used command-line utilities most often used in maintaining and testing routing functionality are pathping, PING, Route, and Tracert.
- A, B, D. Answer A is wrong because show helpers is a Netsh command. Answer B is wrong because both show helpers and show routing are Netsh commands. Answer D is wrong because Trace is a C++ debugger command.

Chapter 5: Planning, Implementing, and Maintaining an Internet Connectivity Strategy

Connecting the LAN to the Internet

1. You have five Windows XP clients on a network with a Windows Server 2003 server. The server has an always-on Internet connection with an ISP. What service can you install on the server to allow the clients to access the Internet, without requiring you to obtain additional IP addresses from your ISP?
- A. PPTP
 - B. NAT
 - C. DHCP
 - D. DNS
- B. The Network Address Translation (NAT) service allows multiple clients on a LAN to share an existing Internet connection through a single IP address.
- A, C, D. Answer A is incorrect because the Point-to-Point Tunneling Protocol (PPTP) is a VPN protocol. Answer C is incorrect because, although the Dynamic Host Configuration Protocol (DHCP) is one of the functions provided by the NAT service, it is not a complete solution for sharing an Internet connection. Answer D is incorrect because although a Domain Name System (DNS) proxy service is provided by the NAT service, DNS is not a service for sharing Internet connections.

2. You are configuring a simple network with two computers, both running Windows Server 2003. Both will be used as Web servers and must be accessible over the Internet. You have chosen to assign an Internet IP address to each machine, and you want to configure a single Internet connection for use by both machines. Which of the following is the best strategy?
- A. Use a routed connection.
 - B. Use NAT.
 - C. Use ICS.
 - D. Two separate connections are required.
- A.** A router provides a simple way to connect both machines to the Internet. Each will require an IP address.
- B, C, D.** Answer B is incorrect because NAT is unnecessary for a simple network where all machines will have an Internet IP address. Answer C is incorrect because Internet Connection Sharing (ICS) provides the same service as NAT, and connection sharing is not required in this case. Answer D is incorrect because a single connection can be used, although two separate public IP addresses will be required.
3. Your network includes a Windows Server 2003 computer and several workstations running Windows 2000 and Windows XP. You need to configure the server to provide shared Internet access to all machines on the network. The server will also act as a Web server. In addition, one of the workstations is providing an FTP service and requires its own Internet IP address. Which solution will address all of these requirements?
- A. ICS
 - B. A hardware router
 - C. NAT
 - D. IAS
- C.** The NAT service can provide a shared Internet connection for all workstations and allow more than one computer to have an IP address accessible to the Internet.
- A, B, D.** Answer A is incorrect because ICS does provide shared Internet connections, but does not allow more than one Internet IP address. Answer B is incorrect because a hardware router does not provide address translation or connection sharing. Answer D is incorrect because Internet Authentication Service (IAS) provides authentication for users, not shared Internet access.
4. You have a DHCP server on the network that automatically assigns IP addresses to clients. You are configuring a NAT server to provide shared Internet access. You want clients to use internal addresses from the same pool, whether or not they are using the Internet. What is the most efficient way to do this?

- A. Divide the address pool between the NAT server and the DHCP server.
 - B. Define identical address pools on the NAT server and the DHCP server.
 - C. Configure NAT to forward IP addressing requests to the DHCP server.
 - D. Remove the DHCP server from the network and use NAT exclusively.
- C. NAT can forward addressing requests to the existing DHCP server. This way, the NAT server does not need its own address pool.
- A, B, D. Answer A is incorrect because dividing the address pool would cause either the DHCP or NAT server to run out of IP addresses sooner than necessary. Answer B is incorrect because using the same address pool on both servers could create conflicts. Answer D is incorrect because there is no need to remove the already functioning DHCP server in favor of the limited addressing abilities of the NAT service.

Implementing Virtual Private Networks (VPNs)

5. You are planning a VPN to allow traveling employees to access the network from remote locations. Employees will be using a variety of ISPs to connect to the Internet. You want to ensure that the VPN offers end-to-end encryption between the VPN client and server for maximum security. Which VPN protocol should you use?
- A. PPTP
 - B. L2TP only
 - C. L2TP and IPSec
 - D. PPP
- C. L2TP and IPSec provide VPN connectivity with end-to-end encryption.
- A, B, D. Answer A is incorrect because PPTP does not support end-to-end encryption. Answer B is incorrect because L2TP does not provide encryption, but requires the use of IPSec. Answer D is incorrect because PPP is a dial-up protocol, not a VPN protocol.
6. You have configured a VPN server running RRAS under Windows Server 2003. A number of remote workstations are able to access the network by connecting to the Internet using local access methods and establishing a VPN connection. Which of the following terms describes this type of VPN?
- A. Router-to-router
 - B. Point-to-point
 - C. Internet-based
 - D. One-way

- C.** This type of VPN is called an Internet-based (or client-server) VPN.
- A, B, D.** Answer A is incorrect because a router-to-router VPN connects two networks, rather than offering remote access to clients. Answer B is incorrect because point-to-point is not a type of VPN. Answer D is incorrect because one-way is a type of initiation for router-to-router VPNs.
7. You have configured a router-to-router VPN using two Windows Server 2003 computers as VPN servers, each with a local Internet connection. You have configured the VPN servers at each end of the VPN to use the PPTP protocol. Which of the following types of encryption will the VPN use in this configuration?
- A. L2TP
B. MPPE
C. IPSec
D. EAP
- B.** MPPE is used to encrypt VPN traffic when PPTP is used.
- A, C, D.** Answer A is incorrect because L2TP is a tunneling protocol and does not provide encryption. Answer C is incorrect because IPSec encryption is used with L2TP and not with PPTP. Answer D is incorrect because EAP is an authentication protocol and does not provide encryption.
8. You need to configure a VPN connection between the local network and a remote branch. The remote branch has access to a dial-up ISP and will be billed by the hour by the ISP for the time spent online. Which of the following is the best strategy to configure the VPN?
- A. Use a demand-dial connection.
B. Use a persistent connection.
C. Use dial-up access via RRAS.
D. Create a dedicated WAN link.
- A.** A demand-dial VPN can provide connectivity to the remote branch while minimizing the expense of the dial-up ISP.
- B, C, D.** Answer B is incorrect because a persistent connection cannot be used with a dial-up connection. Answer C is incorrect because using dial-up access via RRAS would require a long-distance call and would not take advantage of VPN features. Answer D is incorrect because a dedicated WAN link is not part of a VPN solution.

Using Internet Authentication Service (IAS)

9. You have three RRAS servers configured for VPN access for remote clients. The servers are currently using Windows authentication, and you wish to use IAS for centralized authentication. You have installed the IAS component on a Windows Server 2003 computer. What additional task is necessary to enable IAS authentication?
- A. Install IAS on all RRAS server computers.
 - B. Configure each RRAS server to use RADIUS authentication.
 - C. Install a RADIUS client.
 - D. Choose authentication protocols.
- B.** You need to configure each RRAS server to use the RADIUS (IAS) server for authentication.
- A, C, D.** Answer A is incorrect because IAS needs to be installed on only one computer. Answer C is incorrect because the existing RRAS servers will act as RADIUS clients. Answer D is incorrect because the default authentication protocols will be used if you do not choose protocols.
10. You have installed the IAS component on a Windows Server 2003 server. You are planning the authentication strategy for the IAS server and have configured the IAS server to use EAP for authentication. Which of the following protocols are supported by EAP? (Select all that apply.)
- A. MD5 CHAP
 - B. PAP
 - C. SPAP
 - D. EAP-TLS
- A, D.** EAP supports the MD5 CHAP and EAP-TLS authentication types.
- B, C.** Answer B is incorrect because PAP is a basic authentication method and is not part of EAP. Answer C is incorrect because SPAP is not supported by EAP.
11. You have an IAS server running Windows Server 2003. It supports a group of RRAS servers used to manage VPN connections for clients. You are configuring the authentication methods for the IAS server and want to allow the clients to use smart cards for secure and convenient authentication. Which of the following authentication protocols should you select?
- A. MS-CHAP
 - B. EAP-TLS
 - C. MD5 CHAP
 - D. MS-CHAP v2

- B.** The EAP-TLS protocol supports smart card authentication.
- A, C, D.** Answer A is incorrect because MS-CHAP is a password authentication method and does not support smart cards. Answer C is incorrect because MD5 CHAP is an implementation of the same CHAP protocol under EAP. Answer D is incorrect because MS-CHAP v2 is also a password authentication protocol.
12. You have configured an RRAS server on one Windows Server 2003 computer and an IAS server on another, and configured the RRAS server to use the IAS server for authentication. In RADIUS terminology, which computer(s) are referred to as network access servers?
- A. The IAS server
- B. The RRAS servers
- C. The clients of the RRAS server
- D. Both the IAS and RRAS servers
- B.** The RRAS server is the network access server (NAS).
- A, C, D.** Answer A is incorrect because the IAS server is the RADIUS server, not the access server. Answer C is incorrect because the clients do not communicate with the IAS server. Answer D is incorrect because only the RRAS server is a network access server.
13. During a security audit, you are monitoring network traffic and notice that plaintext versions of passwords are passing through the network. You are using an IAS server to handle authentication. Which protocol do you need to disable at the IAS server to prevent this security risk?
- A. MS-CHAP
- B. PAP
- C. EAP-TLS
- D. CHAP
- B.** PAP uses plaintext passwords and should be disabled unless required for legacy clients.
- A, C, D.** Answer A is incorrect because MS-CHAP uses a challenge-response system and does not transmit passwords across the network. Answer C is incorrect because EAP-TLS is an encrypted protocol. Answer D is incorrect because CHAP, like MS-CHAP, does not transmit plaintext passwords.
14. You have an IAS server running Windows Server 2003. You need to enable and configure EAP to support clients that use EAP authentication. In the IAS MMC snap-in, where do you find the options for configuring EAP?

- A. Properties
 - B. Remote Access Policies
 - C. Protocols
 - D. Connection Request Processing
- B.** The options for EAP are configured under Remote Access Policies.
- A, C, D.** Answer A is incorrect because the Properties dialog box does not include authentication options. Answer C is incorrect because there is no Protocols section or dialog box. Answer D is incorrect because the Connection Request Processing options relate to forwarding requests to external RADIUS servers.
15. You wish to create client software for VPN clients to connect to the network so that clients do not need to manually specify the VPN server, tunneling protocol, and other settings. Which program allows you to customize the client software?
- A. Connection Manager
 - B. Connection Manager Administration Kit
 - C. RRAS MMC snap-in
 - D. IAS MMC snap-in
- B.** The Connection Manager Administration Kit (CMAK) allows you to create custom client software.
- A, C, D.** Answer A is incorrect because Connection Manager is the actual client software, not the customization program. Answer C is incorrect because the RRAS MMC snap-in configures the RRAS server, not clients. Answer D is incorrect because the IAS MMC snap-in configures an IAS server.

Chapter 6: Planning, Implementing, and Maintaining a Name Resolution Strategy

Planning for Host Name Resolution

1. You are the administrator of a Windows Server 2003 network. Recently, your company made a sudden and unexpected announcement that it would be merging with another company called Syngress Industries, a large company that has more than 20,000 employees. You learn that, in the short term, communications between the two companies will need to take place over persistent VPNs using each company's respective connections to the Internet, both of which are operating at about 75 percent capacity. You will need to set up trust relationships between two AD forests. Furthermore, you plan to move significant amounts of data between the two networks. You learn the Syngress Industries uses a

child domain of its Internet domain namespace for its AD forest root. The name of the internal domain is ad.syngress.com. You want to ensure that your DNS infrastructure can resolve names for internal hosts of Syngress Industries. You also want to ensure that your solution is the most effective in terms of resource usage. What should you do to enable name resolution for internal hosts of Syngress Industries?

- A. Create a secondary zone for ad.syngress.com on you DNS servers.
 - B. Create a stub zone for syngress.com on your DNS servers.
 - C. Create an Active Directory-integrated zone for ad.syngress.com
 - D. Create a conditional forwarding configuration on your DNS servers for ad.syngress.com
- D.** Configuring conditional forwarding is the correct answer because it best satisfies the condition to be the most effective in terms of resource usage, which primarily is bandwidth in this case. After a time, the forwarding servers would acquire a cache of frequently accessed resources in the ad.syngress.com domain.
- A, B, C.** Answer A is incorrect because creating a secondary zone would enable name resolution, but would cause a significant amount of zone replication traffic over the VPN. Answer B's solution might work if the syngress.com zone contained NS records to delegate authority to the ad.syngress.com domain. However, this would be a bad security practice, since syngress.com is used for Internet clients to resolve names of the publicly available syngress.com servers. Furthermore, the presence of a firewall between the syngress.com DNS servers and the ad.syngress.com servers would mean that the NS and A glue address records would resolve to external IP addresses of the firewall and not IP addresses on the internal network. Answer C is incorrect because your organizations are in two separate AD forests.

2. You are the administrator of a Windows Server 2003 network. Your boss has just read an article on how DNS servers can be compromised so that they will redirect recursive queries to bogus Web sites that can cause potential harm. Your boss has asked you to ensure that the DNS servers in the DMZ have the highest level of protection possible against this and other types of common attacks on DNS servers. You have two DNS servers. DNS-A is used to resolve name mappings for your public Web and mail server. The other DNS server, DNS-B, is used by the internal proxy server to resolve Web site addresses to IP addresses. What actions should you take to carry out your boss's order to provide the highest possible security against common multiple DNS attacks? (Select the best answer.)
- A. Enable protection against cache pollution on DNS-B and disable recursion on DNS-A
 - B. Enable protection against cache pollution on DNS-A and disable recursion on DNS-B
 - C. Disable recursion on DNS-A and configure the firewall to not allow any inbound traffic with destination ports of TCP or UDP port 53 to reach DNS-B
 - D. Disable recursion on DNS-B and configure the firewall to not allow any inbound traffic with destination ports of TCP or UDP port 25 to reach DNS-A

- C.** The problem your boss is describing is cache pollution. Although you can enable protection against cache pollution to mitigate this risk, you should try to stop the potential risk at the firewall, if possible. By configuring the firewall to not allow any inbound traffic that uses the DNS ports from reaching DNS-B, you are preventing any potentially malicious traffic in the form of bogus DNS queries from reaching DNS-B in the first place. You can't use the same restriction for DNS-A, because it provides name resolution for Internet hosts that wish to connect to your Web and mail servers. However, if recursion is disabled on DNS-A, it will still answer queries for zones that it is authoritative for, but it will send a negative response to recursive queries. Disabling recursion also has the added benefit of providing a degree of protection against DoS attacks.
- A, B, D.** Answer A is workable and provides additional security. However, the boss wants the highest level of protection against *multiple* common attacks on DNS servers, so this choice is not as good as Answer C. Answers B and D are wrong because they compromise the ability of DNS-A to resolve the names of your Web and mail servers.
3. You are the administrator of a Windows network that consists of a mixture of Windows NT 4, Windows 2000, and Windows Server 2003 servers, providing a mix of file, print, messaging, and other services critical to your network. You are currently running WINS, DNS, and DHCP services on your network. You have already enabled dynamic DNS on your forward and reverse lookup zones, but you want to ensure that all of your client computers can find the name-to-address mapping of all your servers using DNS. You want to minimize the administrative effort for this project. What action should you take? (Select the best answer.)
- A. Place the DHCP servers in the DnsUpdateProxy group.
- B. Enable DHCP to update forward and reverse lookup zones on behalf of all DHCP clients.
- C. Manually enter the records for servers that have static addresses.
- D. Create a WINS resource record in the forward and reverse lookup zones.
- D.** Windows NT 4 operating systems are not able to update static addresses in a dynamic zone. You must either manually enter resource records for these servers or configure the DNS to query the WINS server when it cannot resolve a name mapping. Since the latter involves the least administrative effort, Answer D is the correct choice.
- A, B, C.** Answer A is incorrect because it will not have an effect on whether resource records for clients are created in the DNS zones. Answer B is incorrect because it is unlikely a server is going to be configured as a DHCP client. Answer C would work, but it involves more administrative effort than the correct response and has a greater risk of introducing error.

4. You are using ISA Server 2000 as a firewall and Web proxy server to protect your internal AD network and provide Web proxy and caching services for HTTP requests. You currently are using three DNS servers to support the DNS queries. DNS-A is used for your internal AD root. DNS-B is used to provide name resolution for Internet clients that want to connect to your public Web and mail servers. DNS-C is used to provide Internet name resolution. How should you configure the DNS and ISA Server access rules to provide the maximum security and functionality for your DNS infrastructure?
- A. On DNS-A, remove the root hints file and enable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, remove the root hints file and disable recursion. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, enable recursion and update the root hints file. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 with a source port of ANY.
 - B. On DNS-A, remove the root hints file and disable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, remove the root hints file and disable recursion. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, enable recursion and update the root hints file. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 with a source port of ANY.
 - C. On DNS-A, remove the root hints file and enable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, remove the root hints file and disable recursion. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, enable recursion and update the root hints file. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 with a source port of ANY.
 - D. On DNS-A, remove the root hints file and disable recursion. Configure ISA Server to allow no traffic to or from this server. On DNS-B, update the root hints file and enable recursion. Configure ISA Server to allow inbound traffic on TCP and UDP port 53 to the DNS server with a source port of ANY. On DNS-C, disable recursion and update the root hints file. Configure ISA Server to allow outbound traffic on TCP and UDP port 53 with a source port of ANY.
- A.** DNS-A is used for internal DNS resolution. You do not want it to perform recursion to the Internet or be accessible through the firewall. You need to remove the root hints file and prevent ISA Server from forwarding Internet traffic to it. However, it should still be able to perform recursion on your internal network. DNS-B is used to provide authoritative responses to requests from Internet clients who wish to connect to Web and mail servers, but it should not be able to perform recursion. You should disable recursion and remove the root hints file on this server. ISA Server needs to be configured to allow inbound traffic to this server on TCP and UDP port 53 with a source port of ANY. DNS-C is used by ISA Server itself to provide name resolution for Web proxy requests. It

needs to be able to perform recursion. ISA Server should be configured to allow it to communicate with external DNS servers using TCP and UDP port 53 with a source port of ANY.

- B, C, D.** The remaining responses are incorrect because they do not meet the requirements, as explained above.
5. You are the administrator of a Windows Server 2003 network. Your company has recently merged with another company and you have set up trusts between the AD forests and have set up conditional forwarding on your DNS servers to resolve names in the AD forest of the newly merged company. You would like your users to be able to resolve names in the newly merged company with the least possible effort and typing on their part. You would like to implement a solution with the least possible effort on your part. What should you do?
- A. Using ADSI, create an msDS-AllowedDNSSuffixes attribute in the domain object container and include the domain suffix of the newly merged AD forest in the list of allowable suffixes.
 - B. Create a group policy that configures the DNS clients with a custom DNS suffix search list.
 - C. Configure the DHCP server option 81 to supply the name of the domain suffix of the newly merged AD forest to DHCP clients.
 - D. Configure a stub zone for a root domain of the newly merged company on your DNS servers.
- B.** To enable DNS clients to resolve unqualified names (single computer names that require the least typing on the part of the client) in a disjointed namespace, you must create a custom DNS suffix search list. You can manually configure this on the DNS clients. However, Group Policy is the most efficient means of implementing this configuration on the client computers.
- A, C, D.** Answer A would allow the primary computer name to be different from the AD domain name the computer is a member of and is not a relevant solution. Answer C is incorrect because DHCP option 81 allows you to specify only one domain name, which should be the domain name used for your own AD domain. Answer D is incorrect because a stub zone would only accomplish what your conditional forwarding is already doing.
6. You are a DNS administrator of a large, distributed Windows Server 2003 network. The AD domain tree consists of a number of child domains that reflect the geographic locations of the different offices of the company. You are responsible for the DNS root domain of the AD forest and the child domain of the office where you work. All administrative responsibility for the remaining child domains is performed by locally based administrators in their respective offices. The capacity of the WAN links connecting the various offices is

showing signs of being insufficient. You want to ensure that DNS resolution for the child domains outside your administrative control will work company-wide in a fault-tolerant manner without adding additional strain to available resources. What should you do? (Select the best answer.)

- A. On the root DNS servers, configure conditional forwarding for the child domains.
 - B. On the DNS servers in the child domain under your control, configure secondary zones for the other child domains.
 - C. On the root DNS servers, configure stub zones for the child domains.
 - D. On the DNS servers in the child domain under your control, configure secondary zones for the other child domains.
- C.** When you configure stub zones on the DNS servers responsible for the root, the SOA, NS, and A records that indicate the authoritative servers for the child domains are automatically updated whenever a local administrator makes changes to these records in the primary zone. These DNS servers for these subdomains are not under your control, so, if you were to configure conditional forwarding on the root DNS servers, the local administrators would need to inform you so that you could manually make the required configuration changes. Stub zones provide the most fault-tolerant solution. Configuring secondary zones on the root DNS servers would also allow fault-tolerant name resolution, but would increase replication traffic across the WAN.
- A, B, D.** Answers B and D are incorrect because the solution must ensure DNS resolution for the entire company. If you were to implement these solutions in your child domain, the scope of the solution would be limited to your domain and not the other child domains. Of course, you and the other administrators may want to implement such solutions to minimize the amount of DNS referral traffic that would occur if DNS servers had to walk the tree to perform iterative queries in an attempt to resolve names in the various child domains.

7. You are the enterprise administrator of a Windows network that comprises a number of Windows 2000 and Windows 2003 domain controllers. You want to use Active Directory-integrated zones for your zone data to enhance security and optimize replication of zone data. What should you choose as the replication scope? (Select the best answer.)

- A. To all DNS servers in the forest
 - B. To all domain controllers in the AD domain
 - C. To all DNS servers in the AD domain
 - D. To all domain controllers specified in the scope of an application partition
- B.** Because you still have Windows 2000 domain controllers in your environment, your only choice is store the zone data in the domain partition.
- A, C, D.** These answers are incorrect because they require the presence of an application directory partition, which is not available on Windows 2000 domain controllers.

Planning for NetBIOS Name Resolution

8. You are an administrator of a Windows Server 2003 network. You want to automate the backups of the WINS database. You want this backup to occur at least once every 24 hours. What should you do? (Select the best answer.)
- A. Configure the Windows Backup utility to back up the contents of the *%systemroot%\System32\Wins* folder once every 24 hours.
 - B. Using the AT command scheduler, create a batch file that temporarily stops the WINS service, copies the WINS database to another location, and then restarts the service.
 - C. Use a third-party backup solution that is capable of backing up open files and configure it to back up the contents of the *%systemroot%\System32\Wins* folder once every 24 hours.
 - D. In the WINS server console, configure a path to store backups of the database and initiate a manual backup.
- D.** The WINS service includes the ability to back up the WINS database automatically once every 24 hours and on the WINS service shutdown, or to back it up manually. To configure WINS to perform automatic backups of the database, you must specify a path for the backup and perform at least one manual backup of the database. You can subsequently use Windows Backup or a third-party backup solution to back up the contents of the WINS backup folder without needing to be concerned about the consequences of backing up an open file.
- A, B, C.** The remaining answers are partially viable to varying degrees, but do not represent the best solution.
9. You are the administrator of a Windows Server 2003 network. You are responsible for a number of WINS servers that are set up as push/pull replication partners to each other. You have a number of static mappings in your WINS database and want to remove one of these mappings from the WINS database. You want to ensure that the record is deleted on all servers with the least administrative effort. How should you delete the WINS static mapping? (Select the best answer.)
- A. On the owner server of the mapping, find the record and perform a simple deletion.
 - B. On the owner server of the mapping, find the record and perform a tombstone deletion.
 - C. On all of the WINS servers, find the record and perform a simple deletion.
 - D. On all of the WINS servers, find the record and perform a tombstone deletion.
- B.** When you perform a tombstone deletion, the record is marked with an attribute that is replicated with the record to other WINS server. The attribute instructs other WINS servers to remove the record through the scavenging process.

A, C, D. Answer A is incorrect because the record will still remain on the replication partners and will eventually be replicated back to the owner WINS server. Answers C and D are incorrect because they require unnecessary administrative effort to accomplish something that can be performed in one simple operation.

10. You are the administrator of a Windows Server 2003 network. You have five WINS servers and need to reconfigure the replication topology as a result of some recent upgrades to your WAN links. All of your WAN links connecting the head office and your four branch offices now have ample bandwidth to handle additional traffic. You want to ensure the shortest convergence time of replicated records, while at the same time keep the number of replication partnership agreements to an absolute minimum. What replication topology should you choose? (Select the best answer.)

A. Ring topology

B. Mesh topology

C. Hub-and-spoke topology

D. Hybrid of ring and hub-and-spoke topology

C. A hub-and-spoke topology ensures the shortest convergence time with the fewest replication partnerships to manage. The longest path from one server to any other is two hops. The number of partnership agreements is eight. (You need to define a push/pull partnership agreement on each side of the replication path between the hub server and the spoke servers.)

A, B, D. Answer A is incorrect because it would require 10 push/pull partnership agreements to establish and would result in replication paths that were three hops in distance. Answer B is incorrect because it is an overly complex replication topology and would require 20 replication partnership agreements to manage. Answer D is an overly complex topology for the number of WINS servers and not required by the design.

Troubleshooting Name Resolution Issues

11. You are an administrator of a Windows Server 2003 network. Your company, Syngress Industries, manages its own DNS for its public Web and mail servers. The primary DNS server for the syngress.com domain is located in a DMZ protected by ISA Server. Your ISP is hosting secondary servers for the syngress.com domain on its BIND 9 servers. While going through your performance logs, you notice a brief but sudden increase in the number of AXFR requests received and AXFR success sent events. Previously, these counters had values of zero in your logs. You suspect your ISP has changed the configuration of its BIND servers, but the ISP denies it and insists that the secondary zones are behaving optimally. You are concerned by these values and decide to investigate the issue and cor-

rect it, if necessary. What is the likely cause of the problem and what should you do? (Select the best answer.)

- A. A rogue DNS server is attempting to pollute the cache on your DNS server by sending bogus queries over TCP, rather than UDP. You should turn on debug logging to determine the source IP address and block all traffic from this address on ISA Server. You should also enable protection against cache pollution and inform the ISP.
- B. A malicious user is issuing an **nslookup -ls** or equivalent command against your DNS server. You should configure the DNS server to allow zone transfers only to the IP addresses of the secondary servers at the ISP. You should also block all external requests destined for the primary DNS server on TCP port 53 with a source port of ANY, except for the external addresses of the secondary servers. You should inform the ISP managers and ask them to confirm an equivalent level of security on their servers.
- C. A malicious user is attempting to launch a DoS attack on your DNS. You should disable recursion on the DNS server. You should also turn on debug logging to determine the source IP address of the attack and block the IP address at ISA Server. You should inform the ISP to be on the lookout for similar attacks against its DNS servers.
- D. A malicious user is issuing an **nslookup -ds** or equivalent command against your DNS server to get detailed information. You should turn on debug logging to determine the source IP address. Once you determine the IP address, you should block it from all communication with your DNS servers at ISA Server. You should inform the ISP managers and ask them to confirm an equivalent level of security on their servers.

B. AXFR is the DNS protocol used for full zone transfers. Counters in your performance logs indicate requests to do a full zone transfer have been received by and successfully responded to by your DNS server. That means that someone has issued an **nslookup -ls** or equivalent command against your DNS server. By default, BIND 9 servers will attempt to use IXFR to perform incremental zones transfers, unless this option is explicitly disabled. Since you experienced only a brief event, it is likely the user got what he or she wanted. You should, however, protect your server against future occurrences of zone transfers to unauthorized IP addresses, which is also known as footprinting or name dumping. TCP port 53 is used for zone transfers, and blocking this port should not affect the DNS server's ability to respond to name queries, which should be taking place on UDP port 53.

A, C, D. Answer A is incorrect because an attempt to pollute the cache would normally occur as a result of rogue DNS server replying with information that is superfluous to a query issued against it by the DNS. Answer C is incorrect because a DoS attack is most effective if it ties up a DNS server with recursive query requests. It is no doubt possible to tie up a DNS server with excessive zone transfer requests, but you would expect this activity to be sustained over a period of time. Answer D is incorrect because the **nslookup -ds** command requests detailed information on a particular record and is used for debugging. It does not display the contents of the entire zone the way

an nslookup -ls command would. (Note that the -ls switch is available only in NSLookup interactive mode; the nslookup -ds switch is available only in NSLookup noninteractive mode.)

12. You are the administrator of a Windows Server 2003 network. Recently, a junior administrator has, on your instructions, rebuilt one of your WINS servers (WINS-A). You don't have a backup of the WINS database and need to restore the database through reregistrations of WINS clients and replication with another WINS server, WINS-B. Both servers are configured as push/pull replication partners of each other. As soon as WINS-A is brought back online, users configured to use WINS-A as their WINS server immediately start to complain that they can't access file server shares on this server. By the time you hear about the complaints and try to reproduce the results, you find that the problem has disappeared. However, you take the complaints seriously and investigate further. You examine the WINS database on WINS-B and see some data that strikes you as odd. Based on the data shown in the table here, what problem is indicated? (Select the best answer.)

Record Name	Type	IP Address	Owner	Version
WINS-A	[00h] Workstation	192.168.100.20	192.168.179.5	20D
WINS-A	[20h] File Server	192.168.100.20	192.168.179.5	20C

- A. There is a problem with the order of service registration. The workstation service needs to be registered before the file server service.
 - B. There is a problem with WINS replication that has caused the wrong owner to be associated with WINS-A.
 - C. The TCP/IP stack on WINS-A is configured with the IP address of WINS-B as its secondary WINS server.
 - D. The TCP/IP stack on WINS-B is not configured to register itself with a WINS server.
- C. WINS-A is registering its NetBIOS names with WINS-B, rather than itself. A comparison of the IP Address and Owner fields show two different addresses. These should match or problems with name resolution on the network can occur. In the scenario described here, users who pointed to the WINS-B server would have no problem connecting to file server shares on WINS-A because the WINS-B server has a mapping for the file server service on WINS-A. However, users pointing to WINS-A would not be able to resolve this mapping until replication had merged the record from WINS-B, hence the transient nature of the problem. A WINS server should always be configured to register NetBIOS names only with itself.
- A, B, D. Answer A is incorrect because the order in which services register has nothing to do with NetBIOS name resolution. There might be problems with replication, but the evidence presented doesn't point to this, so Answer B is incorrect. While you also-

lutely should configure a WINS server to register its NetBIOS records with itself, it will eventually do so even if the configuration is left blank (this could take some time), so Answer D is incorrect.

13. You are the administrator of a Windows Server 2003 network using DNS and WINS to provide name resolution services. You have two WINS servers that are set up with the default push/pull configurations. Users have been complaining for days about problems connecting to a server called File_Server2. You ping File_Server2 and get a response from the computer. However, when you issue a **net view \\File_Server2** command, you get an error message stating that a duplicate name exists on the network. What is the likely cause of the problem? (Select the best answer.)
- A. The underscore character cannot be used in a NetBIOS name. Rename the computer and reboot it.
 - B. There is a problem with the replication of the records for File_Server2. Manually initiate replication with the WINS server that is the owner of the record of File_Server2.
 - C. The WINS database is corrupt. Manually initiate consistency checking to restore database integrity.
 - D. The WINS server contains an incorrect name mapping for File_Server2.
- D.** You can ping File_Server2, so the issue is related to the NetBIOS name resolution. When you invoke the net view command, you force the use of the NetBIOS interface, which will subsequently enforce the rules for NetBIOS names. Computer names are exclusive and must be unique. Because host name resolution resolves the name to a different IP address than the IP address resolved by the NetBIOS name mapping, you will get a duplicate name error message. We know the IP address returned by the ping is correct and that host name resolution is working for this computer.
- A, B, C.** Answer A is incorrect because underscores are valid characters for NetBIOS names. Underscores are problematic in some implementations of DNS, but are not a problem for Windows DNS. Answer B is a possibly correct answer because, if the WINS record has not replicated throughout the environment, you might see a similar problem. However, users have been complaining for some time—much longer than the default replication interval. Answer C is incorrect because if there were problems with database consistency, the problems would be more widespread.
14. You are the administrator of a WINS server. The WINS server has suffered a hardware failure, and you have subsequently been forced to reinstall Windows Server 2003 and the WINS service. Fortunately, you have a recent backup of the WINS database. You restore the database, but notice that none of the former WINS configuration settings are present. What should you do? (Select the best answer.)
- A. You need to use the `%systemroot%\system32\jetpack.exe` file to restore the WINS configuration after you restore the database.

- B. You need to restore the original system state from the backup to the Windows Server 2003 server.
- C. You need to invoke database consistency checking on the database.
- D. You need to set up replication with a WINS server that was a replication of the former WINS server.
- B.** WINS configuration settings are stored in the Registry. The WINS database contains only NetBIOS registration data, not configuration information. You therefore need to restore the Registry in order to restore the WINS configuration settings. You can do this by restoring the system state backup or a backup of the Registry itself.
- A, C, D.** Answer A is incorrect because the Jetpack utility does not have this functionality. Answers C and D are incorrect because the database does not contain any WINS configuration information.
15. You are the administrator of a Windows Server 2003 network. After restoring the Windows Server 2003 domain controller that you had taken off the network for a few hours for maintenance, your Windows 95 and 98 users have begun complaining that they are unable to access resources on this computer. You remember seeing a message about a duplicate name on the network when you turned on the domain controller, but didn't think much of it at the time because you had changed the IP address of the domain controller before you took it offline. What action should you take?
- A. Create static mappings in the WINS database for the domain controller and disable the migrate on setting.
- B. Create static mappings in the WINS database for the domain controller and enable the migrate on setting.
- C. Have the users of Windows 95 and 98 computers issue an **nbtstat -RR** command.
- D. Have the users of the Windows 95 and 98 computers issue an **ipconfig /flushdns** command.
- A.** It is likely that someone on your network has configured a computer with the same name as the domain controller and hijacked the NetBIOS registration of the domain controller, resulting in a redirection attack. Windows 95 and 98 clients will use NetBIOS for logon services and to connect to file sharing resources. Given the circumstances, the duplicate name message is clear evidence of this kind of attack. If another computer is registered with the same name and is online, the WINS server will report a duplicate name error message back to the computer that is trying to initialize with the same name. For mission-critical servers, it is good idea to create static mappings that cannot be overwritten by dynamic registrations. This situation represents one of the few circumstances that can justify the use of static mappings.
- B, C, D.** Answer B is incorrect because enabling the migrate on setting would allow a dynamic registration to overwrite a static registration. Answers C and D are incorrect

because flushing either of the resolver caches on the client would have no effect on the ultimate results of having an incorrect record in the WINS server.

Chapter 7: Planning, Implementing, and Maintaining a Remote Access Strategy

Planning the Remote Access Strategy

1. You are planning a remote access server and need to enable access for several employees. All the employees are in the same city. The company LAN is not currently connected to the Internet, and your security policy specifies that Internet connections should be avoided. Which of the following is the best choice for the remote access solution?
 - A. Dial-in access
 - B. VPN access
 - C. Wireless access
 - D. Dedicated WAN links

A. Dial-in access is a convenient way to offer access to employees within a city; therefore, Answer **A** is correct.

Answer **B** is incorrect because VPN access requires Internet connections. Answer **C** is incorrect because wireless access is typically not feasible over long distances. Answer **D** is incorrect because dedicated links for each employee would add unnecessary expense.
2. You are configuring a remote access server on a Windows Server 2003 computer. The same server is acting as a domain controller and DHCP server, assigning IP addresses to clients. Which of the following is the simplest method of assigning IP addresses for remote clients?
 - A. Manually configure each client with an IP address.
 - B. Configure the RRAS server to use DHCP.
 - C. Configure a static address pool.
 - D. Use APIPA.

B. Because a DHCP server is already available, you can configure the RRAS server to request addresses from DHCP and avoid the need for separate addressing for dial-up clients; therefore, Answer **B** is correct.

Answer **A** is incorrect because manual configuration is not the simplest method. Answer **C** is incorrect because a static address pool would require additional configuration and consideration of potential conflicts with the DHCP server's address range. Answer **D** is incorrect because APIPA is intended for small networks that do not have a DHCP server available.

Addressing Dial-In Access Design Considerations

3. You are configuring a dial-in remote access server on a Windows Server 2003 computer. Employees will use remote access while traveling. You have ten employees with laptops who will require access to the server, but typically only one is traveling at a time. A telecommuting employee will also require access for eight hours a day. How many modems would be the minimum to reliably serve these users?
- A. 1
 - B. 11
 - C. 2
 - D. 3
- C.** Two modems should be sufficient: one for the telecommuting employee and one for any traveling employee who requires access; therefore, Answer **C** is correct.
- Answer **A** is incorrect because one modem would be busy for eight hours a day and traveling employees would not be able to dial in. Answers **B** and **D** are incorrect because two modems should be sufficient.
4. You have several users who dial in to a remote access server using multilink connections, combining two modems into a single link. Although this provides a higher bandwidth to the users, you find the server runs out of modem lines frequently, and most users are not using their connections to their full potential. Which of the following is a solution to this issue?
- A. Disable multilink connections.
 - B. Set the maximum number of multilink ports to one.
 - C. Use VPN instead of dial-in access.
 - D. Enable Bandwidth Allocation Protocol (BAP).
- D.** Bandwidth Allocation Protocol (BAP) can reduce a multilink connection by one line when it is not used to its full capacity, freeing the modem for other users; therefore, Answer **D** is correct.
- Answer **A** is incorrect because disabling multilink entirely would unnecessarily reduce bandwidth for users that required it. Answer **B** is incorrect because setting the maximum number of ports to one would effectively disable multilink. Answer **C** is incorrect because using VPN access is not an immediate solution to this issue.

Addressing VPN Design Considerations

5. You are configuring a Windows XP client machine to access a VPN server that supports L2TP over IPSec. You need to obtain a computer certificate for the client and wish to do

so from the client machine. A CA is present on the local network. Which application can you use to request a certificate?

- A. A Web browser
 - B. The Certificates MMC snap-in
 - C. The Certification Authority MMC snap-in
 - D. Connection Manager
- A.** You can request a certificate by connecting to the CA using a Web browser; therefore, Answer **A** is correct.
- Answer **B** is incorrect because you can use the Certificates MMC snap-in to request a certificate, but MMC is not usually installed on Windows XP. Answer **C** is incorrect because the Certification Authority snap-in is available only for the CA. Answer **D** is incorrect because Connection Manager can be used to make a VPN connection, but not to request a certificate.
6. You have configured a VPN server running Windows Server 2003 and RRAS. Most clients are able to access the server, but clients running Windows 98 are reporting that they are unable to connect. Which of the following is most likely the cause of this problem?
- A. Computer certificates are not installed.
 - B. L2TP is not enabled on the server.
 - C. PPTP is not enabled on the server.
 - D. Windows 98 does not support VPN client access.
- C.** The likely problem is that PPTP is not enabled on the server, since Windows 98 clients do not support L2TP; therefore, Answer **C** is correct.
- Answer **A** is incorrect because computer certificates are used with L2TP, which is not supported by Windows 98. Answer **B** is incorrect because L2TP support would not work with Windows 98 clients. Answer **D** is incorrect because Windows 98 does support VPN access, but requires the PPTP protocol.

Addressing Wireless Remote Access Design Considerations

7. You are setting up wireless access to the network with two WAPs. You want to use a centralized authentication source for both access points. You have an existing IAS server on the network. Which of the following tasks are necessary to support wireless access? (Choose all that apply.)
- A. Create a remote access policy.
 - B. Configure the WAPs to use RADIUS authentication.

- C. Install a RADIUS server.
- D. Add the WAPs as clients in the IAS server's configuration.
- A, B, and D.** You will need to create a remote access policy, configure the WAPs to use RADIUS authentication, and add them as clients of the IAS server; therefore, Answers **A, B, and D** are correct.
- Answer **C** is incorrect because the existing IAS server will act as the RADIUS server.
8. You have configured a WAP using the EAP-TLS protocol. The WAP is connected to a LAN with a Windows Server 2003 server. Which of the following additional tasks may be necessary to ensure that wireless clients can connect? (Choose all that apply.)
- A. Enable PPP authentication.
- B. Issue computer certificates to clients.
- C. Issue user certificates or smart cards to users.
- D. Install and configure IAS.
- B and C.** For wireless access to work, each client needs a computer certificate and either a user certificate or smart card; therefore, Answers **B** and **C** are correct.
- Answer **A** is incorrect because PPP authentication is not used with wireless access. Answer **D** is incorrect because IAS is not needed for wireless access, although it can be used to improve security and to centralize authentication.

Planning Remote Access Security

9. You are planning security for your network and have determined that the domain functional level is Windows 2000 Mixed mode. You have a combination of Windows Server 2003 and Windows 2000 Server domain controllers. Which of the following actions may be necessary to enable all of Windows Server 2003's security features? (Choose all that apply.)
- A. Eliminate or upgrade the Windows 2000 Server domain controllers.
- B. Eliminate all Windows 2000 clients.
- C. Raise the functional level to Windows Server 2003.
- D. Raise the functional level to Windows Server 2003 Interim.
- A and C.** To enable all security features, you can raise the functional level to Windows Server 2003. This will no longer enable Windows 2000 machines to act as domain controllers; therefore, Answers **A** and **C** are correct.

- Answer **B** is incorrect because only the domain controllers must be running Windows Server 2003. Answer **D** is incorrect because the Windows Server 2003 Interim function level does not enable all security features.
10. You have a network with two Windows Server 2003 servers. You have raised the domain function level to Windows Server 2003. You need to install an additional domain controller and are considering an existing Windows 2000 Server. Which of the following tasks is necessary before using this machine as a domain controller?
- A. Lower the function level to Windows 2000 Mixed mode.
 - B. Lower the function level to Windows Server 2003 Interim.
 - C. Upgrade the Windows 2000 Server to Windows Server 2003.
 - D. Demote the existing domain controller to a member server.
- C**. Once the domain function level is raised, it cannot be lowered, so the only solution is to upgrade the server to Windows Server 2003; therefore, Answer **C** is correct.
- Answers **A** and **B** are incorrect because the domain function level cannot be lowered. Answer **D** is incorrect because the existing domain controller does not need to be changed.

Creating Remote Access Policies

11. You have an RRAS server and have configured two remote access policies. The first policy on the list allows access for all members of the Power Users group. The second policy on the list denies access to clients that connect during evening hours. After testing your configuration, you determine that clients in the Power Users group are able to connect at any time. Which of the following actions would correct this problem?
- A. Delete the first policy in the list.
 - B. Change user account properties to deny remote access.
 - C. Change the order of the policies.
 - D. Install an IAS server.
- C**. Because the first policy that matches a client is used, the policy to deny access for evening hours should be first on the list; therefore, Answer **C** is correct.
- Answer **A** is incorrect because the first policy is necessary to grant access to the group. Answer **B** is incorrect because user accounts set to deny access will be denied remote access regardless of the policy. Answer **D** is incorrect because installing IAS is unnecessary to solve this problem.

12. You are operating a remote access server and currently allow VPN access and dial-in access. You have decided to disallow dial-in access after configuring all the clients for VPN access. Which of the following attributes can you check in a remote access policy to deny access to modem users?
- A. Authentication-Type
 - B. NAS-Port-Type
 - C. Framed-Protocol
 - D. NAS-Identifier
- B.** The NAS-Port-Type attribute can be used to check whether dial-in access is in use; therefore, Answer **B** is correct.
- Answer **A** is incorrect because the Authentication-Type option is used to check the authentication method in use. Answer **C** is incorrect because the Framed-Protocol attribute specifies the protocol used to connect. Answer **D** is incorrect because the NAS-Identifier attribute is a string that identifies an RRAS server.

Creating a Plan to offer Remote Assistance to Client Computers

13. One of your users is having problems getting a productivity application to work correctly. You suspect that he is performing the steps involved in using the application incorrectly, but the application interface is complex and it's difficult for you to explain, over the phone, what he needs to do. The user is running Windows XP, and you want to connect to his PC and show him how to perform the task in question so that he can actually see you go through the steps. How would you arrange to do this?
- A. Send the user a Remote Assistance Request.
 - B. Get the user to send a Remote Assistance Invitation.
 - C. Connect to the user's PC using Remote Desktop.
 - D. Connect to the user's PC using the Web Interface for Remote Administration.
- B.** By getting the user to send you a Remote Assistance Invitation, you can connect to the user's desktop and the user can follow what you are doing.
- Answer **A** is incorrect, because sending the user a Remote Assistance Request is the wrong way and it is also not called a Request. Answer **C** is incorrect, because connecting to a user's PC using Remote Desktop logs off anyone at the PC and he will not be able to see what you are doing. Answer **D** is incorrect, because Remote Administration is not available on Windows XP computers.

14. You are attempting to describe the remote assistance process to a co-worker. The co-worker asks what the correct terms are for the person requesting assistance and the person providing assistance so that he can look them up in Windows Help. Which of the following do you reply with? (Select two.)
- A. Administrator
 - B. Novice
 - C. Expert
 - D. End user
- B, C.** In relation to a remote assistance session, Microsoft refers to the person requesting help as the Novice and the person providing it as the Expert.
- A, D.** Although valid terms in computer networking circles, Administrator and End user are not the terms Microsoft uses to officially refer to roles involved in using Remote Assistance.

Planning for Remote Administration by using Terminal Services

15. You are attempting to describe the remote assistance process to a co-worker. The co-worker asks what the correct terms are for the person requesting assistance and the person providing assistance so that he can look them up in Windows Help. Which of the following do you reply with? (Select two.)
- A. Administrator
 - B. Novice
 - C. Expert
 - D. End user
- B, C.** In relation to a remote assistance session, Microsoft refers to the person requesting help as the Novice and the person providing it as the Expert.
- A, D.** Although valid terms in computer networking circles, Administrator and End user are not the terms Microsoft uses to officially refer to roles involved in using Remote Assistance.

Chapter 8: Planning, Implementing, and Maintaining a High-Availability Strategy

Understanding Performance Bottlenecks

1. You have been tasked with the implementation of enhancing the security of your network and have been allocated a modest budget to accomplish the task. You decide to implement IP Security (IPSec) between your three Windows Server 2003 servers and your Windows 2000 Professional and Windows XP Professional workstations. As the implementation proceeds, you begin hearing reports that the network does not seem as responsive. You confirm that performance has decreased. What can you do to return performance to the previous level and still accomplish your objectives?
 - A. Remove IPSec from the workstations, leaving the servers configured with IPSec.
 - B. Remove IPSec from the servers, leaving the workstations configured with IPSec.
 - C. Add NICs to your servers and configure the cards for load balancing.
 - D. Purchase new NICs that support IPSec on the NIC.

D. IPSec is computer-intensive, and NICs that remove this load from the system's main CPU can significantly boost communication performance.

A, B, C. Neither Answer A nor Answer B addresses the real issue. Removing IPSec from either the servers or the workstations may actually stop communications altogether or allow your network to run unsecured. Answer C may actually decrease performance because now multiple network interfaces will require IPSec calculations to be performed by the system CPU.
2. You have inherited the responsibility of supporting a server from a previous administrator. The system has dual 1 GHz CPUs, 2048MB of RAM, and a dual-channel caching hardware RAID controller with sixteen 18GB hard drives configured as a RAID 5 array. The system has been running an important SQL database for some time, but over the last few weeks, responsiveness has decreased as more people have been accessing the SQL databases. Your part-time SQL administrator has told you that recent database growth is not the case. The databases have been consistently using between 40 and 45 percent of the available disk space. You have been asked to resolve this problem. What can you do to increase the responsiveness of the SQL database?
 - A. Install more RAM in the server.
 - B. Change the RAID array to a RAID 0+1 configuration.
 - C. Change the RAID array to a RAID 0 configuration.
 - D. Increase the cache size on the array controller.

- B.** The SQL databases are on a RAID 5 array, which incurs heavy performance hits on disk writes. Since less than half of the disk space is in use, you can reconfigure the array into a RAID 0+1 configuration, which will boost performance and keep the data protected from drive failures.
- A, C, D.** Answer A is wrong because, while adding RAM to a server is a frequent fix of performance issues, there is likely enough RAM already to adequately run the SQL server. Implementing Answer C would solve the performance issues but would make the system susceptible to a drive failure. Answer D may provide some performance boost, but the underlying problem of the RAID array structure would remain, making this answer a stopgap measure at best.
3. You have recently purchased a new single-CPU, Intel Xeon-based server. This hardware will be used to run a multithreaded CPU-intensive application. How can you ensure that the application performs at its best on the hardware provided?
- A. Turn on hyperthreading.
- B. Add a second CPU.
- C. Boost the processing priority of the applications threads.
- D. Disable hyperthreading.
- A.** A recently purchased Xeon server will support hyperthreading. Turning on hyperthreading should yield a performance increase for the multithreaded application.
- B, C, D.** Answer B is incorrect because you would need to purchase additional hardware. Answer C is incorrect because, although performance may improve, hyperthreading will yield a higher performance boost. Answer D is incorrect because disabling hyperthreading will actually have a negative impact on performance.
4. Your server seems slow to respond to file requests from drive D: at times. You have examined the system with Performance Monitor, and the counter LogicalDisk:Current Disk Queue Length for the D: instance consistently varies between 8 and 20 during these periods of slow response. Drive D: resides on an external, 14-slot disk array with 4 slots populated with hard drives. How should you resolve this problem?
- A. Defragment drive D:.
- B. Add more memory to the system to increase file-caching efficiency.
- C. Add more physical drives to the external array; either expand drive D: across the new drives or create another drive and move some heavily accessed files from drive D: to the new logical drive.
- D. Add processors or turn on hyperthreading.
- C.** The problem is that the disk array is not responding to disk requests quick enough and requests are being queued. Adding drives and expanding drive D: to encompass the

new drives will add more “spindles” to service disk requests. Creating a new logical drive on top of the new physical drives and relocating files to the new logical drive can produce the same effect.

- A, B, D.** Answer A is incorrect because the Current Disk Queue Length counter is consistently over two. File system fragmentation does not produce this effect. Answer B is incorrect because adding more memory to increase the cache will likely not increase disk responsiveness. Answer D is incorrect because adding more processors may actually produce more disk requests and increase the queue length even more.
5. You have recently purchased and installed two new name-brand servers. The servers are identical in all respects, except that one server has a single CPU and the other has two. The single-CPU system will be used for basic file and print services, and the dual-CPU system will be used for running Microsoft Exchange Server. Both systems respond adequately. While developing a performance baseline, you notice that the dual-CPU system seems to be experiencing more interrupts per second than the other server. What should you do to resolve this increased level of interrupts?
- A. Do nothing. This is a peculiarity of Microsoft Exchange Server.
- B. Increase the communication buffers on the multiple-CPU server’s NIC.
- C. Remove the second CPU from the dual-CPU system.
- D. Do nothing. This is normal for a multi-CPU system.
- D.** An increased level of interrupts on a multi-CPU system is normal.
- A, B, C.** Answer A is incorrect because any multiprocessor-capable application will generate interrupt activity, not just Microsoft Exchange Server. Answer B is incorrect because, although this may reduce the number of interrupts, the majority of the interrupts are a result of having multiple CPUs. Answer C is incorrect because, even though the high number of interrupts will cease, performance will be greatly reduced.

Planning a Backup and Recovery Strategy

6. You have been asked to develop a backup strategy for your company’s three Windows Server 2003 servers. You have been told that the primary objective is to have the systems up and running again as quickly as possible should a disaster occur. To accomplish this goal, initial funds have been allocated and, if necessary, ongoing funds will be made available. What backup strategy should you adopt?
- A. Full backups nightly to a tape drive installed in each server
- B. Full backups nightly to a single, centralized tape drive

- C. Full backups weekly, with daily differential backups to a tape drive installed in each server
- D. Full ASR backups nightly
- A.** Answer A is correct because only a single tape set will be needed to restore a server and, with a tape drive in each server, a restore can be performed on each server simultaneously and quickly.
- B, C, D.** Answer B is incorrect because, although only a single tape set would be needed to do a restore, multiple restores could not occur simultaneously. Answer C is incorrect because multiple tape sets would be required to perform a restore, increasing restore time. Answer D is incorrect because an ASR backup does not back up data on partitions and volumes that do not contain Windows components.
7. You have been asked to develop a backup strategy for your company's three Windows Server 2003 servers. You have been told that the primary objective is to minimize the ongoing cost of performing backups. To accomplish this goal, you have been given a modest budget. What backup strategy should you adopt?
- A. Full backups monthly, differential backups on the weekends, and incremental backups daily to a tape drive installed in each server
- B. Full backups monthly, differential backups on the weekends, and incremental backups daily to a single, centralized tape drive
- C. Incremental backups daily to a single, centralized tape drive
- D. Periodic full backups and daily incremental backups to a single, centralized tape drive
- D.** Answer D provides for the lowest number of tapes to be used for backups, while still maintaining good restore capability and lowest cost for purchase of hardware.
- A, B, C.** Answer A is incorrect because this scenario would require more tapes. This is a good scenario for balancing the cost of tapes and backup/restore performance. Answer B is incorrect because, although this would reduce the cost of equipment (fewer tape drives), more tapes would be required. This is a good scenario for reducing the impact of tape media failures and lowering the cost of hardware. Answer C is incorrect because, although this is the lowest cost option, it does not provide a point from which restores can be started.
8. You have been asked to develop a backup strategy for your company's three Windows Server 2003 servers. You have been told that the primary objective is to minimize the time required for performing backups on regular business days. You do not have the use of any advanced storage technology, and an older application on the server requires you to shut down the application and disable Volume Shadow Copy to get a successful backup. To accomplish this goal, you have been given a sufficient budget. What backup strategy should you adopt?

- A. Full backups on the weekends and incremental backups daily to a tape drive installed in each server
 - B. Full backups monthly and differential backups daily to a single, centralized tape drive
 - C. Incremental backups daily to a single, centralized tape drive
 - D. Periodic full backups and daily incremental backups to a single, centralized tape drive
- A.** Answer A is correct because an incremental backup to an internal, dedicated tape drive minimizes the time required to perform the backup.
- B, C, D.** Answer B is incorrect because the downtime required for the backups during the working week would likely increase on a daily basis. Answer C is incorrect because it could take longer than Answer A to back up to a central location. It also has the error of not providing a starting point for restores. Answer D is incorrect because the centralized tape drive could again be a performance bottleneck.
9. Your company uses a well-known and respected third-party backup utility for all of its servers. You are adopting Windows Server 2003 early after its release and have upgraded a number of servers to the operating system. You have high hopes about improving backup performance on some of your higher volume file servers (including the ability to back up open files) and have installed the third-party client agent software on your servers. After a few days, you notice that the speed of backups has not increased. What is the most likely reason that backup performance has not increased?
- A. Volume Shadow Copy has not been turned on for the appropriate volumes.
 - B. The third-party backup software does not use the new features present in Windows Server 2003.
 - C. An ASR backup needs to be performed before the third-party utility will show increased performance.
 - D. The drives hosting the files need to be defragmented for performance to improve.
- B.** As with any new operating system, it takes a while for the rest of the market to catch up. It is likely that the third-party backup software does not recognize the new features of Windows Server 2003 and therefore is backing it up as it would an older operating system client. The solution would be to obtain the updated Windows Server 2003 compatible agent and use it on your servers.
- A, C, D.** Answer A is incorrect because, even if Volume Shadow Copy were turned on, it is unlikely that the third-party backup utility will be able to use it. Answer C is incorrect because performing an ASR backup, while always a good idea, cannot affect ongoing backup performance. Answer D is incorrect because, although defragmenting the drives may yield a performance improvement, it will not solve the problem of the third-party software limitations.

Planning System Recovery with ASR

10. You have inherited the responsibility for supporting an important server recently upgraded from Windows NT 4 to Windows Server 2003. When the server was upgraded, it met the hardware requirements, but not by much. Increasing demand on the system has led to lower than desirable performance. Company management has authorized the purchase of new server hardware and would like you to upgrade the server as quickly as possible with the least amount of risk and additional expense. What is the best way to accomplish the upgrade in the fastest possible time, with the lowest risk, and no additional cost?
- A. Use a third-party product to duplicate the server onto the new hardware.
 - B. Create an ASR backup of the existing server. Use the ASR backup on the new hardware. Back up the existing server. Restore the backup to the new hardware.
 - C. Install Windows Server 2003 onto the new hardware. Back up the existing server. Restore the backup to the new hardware.
 - D. Shut down the existing server and move the existing hard drives to the new server. Boot the new server with the old hard drives.
- B.** This method accomplishes the upgrade with the least risk (the existing server is preserved), least expense (all the tools needed are present in Windows Server 2003), and as quickly as possible (as fast as the backup and restore can be done).
- A, C, D.** Answer A is incorrect because you would need to purchase the third-party utility, incurring additional expense. Answer C is incorrect because it would take longer to get the new operating system installed and configured to operate in the same way as the original system. Answer D is incorrect because the risk factor is too high. Moving hard drives might be technically possible, but the drives (containing the only copy of the original server) could be dropped or corrupted in the process.
11. A few weeks ago, you installed a new server. You have been performing regular full and incremental backups for all files on the system. You did not perform an initial ASR backup. When you arrived this morning, you discovered that the hard drive failed sometime last night after the backup completed, and the server will no longer boot. You replaced the failed hard drive with an identical one you had on hand. What is the quickest way to get the server back to its previous operational state?
- A. Start an ASR restore. Since the hard drive is new and identical to the failed drive, ASR will automatically re-create the previous configuration.
 - B. You cannot restore the server. It is permanently lost.
 - C. Reinstall Windows Server 2003 in a minimal configuration, restore the most recent full backup, and then restore all of the incremental backups in sequence.

- D. Reinstall Windows Server 2003 in a minimal configuration, perform an ASR backup, perform an ASR restore, restore the most recent full backup, and then restore all of the incremental backups in sequence.
- C. Answer C is correct because you cannot directly re-create the operational state without an ASR backup. Assuming the system state was backed up with the regular backups, you can re-create the previous operational state from the backups after you have reinstalled the operating system.
- A, B, D. Answer A is incorrect because you must have the ASR media set in order to perform an ASR restore. ASR does not automatically re-create a previous configuration. Answer B is incorrect because you can re-create the server from the backups; it will simply take longer and be more difficult than if you had an ASR backup available. Answer D is incorrect because you do not have an ASR set.
12. You are working on an existing server. The NIC manufacturer has notified you of an updated driver for your card that will greatly improve performance. You download and install the new driver. Before you reboot the system, you perform an ASR backup. When you reboot the system, it reaches the graphical portion of the boot process and presents a STOP message. What is the proper process for recovering from this problem?
- A. Perform an ASR restore from the ASR backup set you created before the reboot.
- B. Reboot the system, press F8 when prompted during the boot process, select Last Known Good Configuration, and press Enter.
- C. Reinstall the operating system and do a restore of the system from tape backup.
- D. Reboot the system, press F8 when prompted during the boot process, select Safe Mode, and press Enter.
- B. The Last Known Good Configuration option will load the drivers that were used during the boot process prior to the last successful logon, provided that they are not missing or corrupt. The new driver that you installed during your last logon session will not be loaded, and the previous one assigned to the NIC will be loaded.
- A, C, D. Answer A is incorrect because the ASR backup set you created before the reboot would contain the newer driver, and it would be configured for use on next boot. Answer C is incorrect because this would destroy your existing operating system and take much longer to fix. Answer D is incorrect because, although Safe Mode may (or may not) allow you to boot successfully, the new driver is still present.

Planning for Fault Tolerance

13. You are responsible for administering a Windows Server 2003 system. The system has a Pentium III 800 MHz CPU, 1024MB of RAM, and four hard drives configured in a

RAID 5 array that reside in an external seven-slot chassis. The array is controlled by a modern, high-performance hardware RAID controller and presents the array to the operating system as a single drive. You arrive on a Monday morning to find your server has crashed. On investigation, you find that two of the hard drives failed. The server has a built-in display that tells you that one drive failed late Friday night and the second drive failed Sunday afternoon. What should you have done to prevent the second drive failure from causing the server to crash?

- A. Ensure that backups complete during business hours.
 - B. Use Volume Shadow Copy to automatically create a backup on the remaining good drive.
 - C. Install a second hardware RAID controller and distribute the drives evenly on the controllers.
 - D. Purchase another hard drive and configure it as a hot spare drive.
- D.** A hot spare drive would most likely have prevented the outage, and there is an available slot for the spare drive. When the first drive failed, the controller would have brought the spare drive online and re-created the missing data on it. When the second drive failed, the array would be in a nonredundant state but would have continued to function.
- A, B, C.** Answer A is incorrect because the time that the backups would complete would not have affected the failure or operation of the drives. Answer B is incorrect because the operating system saw the array as a single drive and would not be able to have performed any operation on a specific disk. Answer C is incorrect because this would have protected against controller or cable failure but not a drive failure.
14. You are replacing a single-port NIC in your server with a new four-port NIC. Your switches support 100 Mbps full-duplex operation. Your switches also support either load-balancing or failover configurations. Which configuration choice is best for increased performance and availability?
- A. Configure the card for two-way load balancing with failover to the remaining two ports.
 - B. Configure the card for four separate links to the switch. Windows Server 2003 automatically determines that the ports connect to the same switch and enables failover.
 - C. Configure the card for four-way load balancing.
 - D. Leave the old NIC in the server and add the new four-port card into an empty slot on the server. Configure the new card as a failover backup for the existing card.
- C.** Answer C is correct because a multiport load-balancing configuration automatically includes failover redundancy.

A, B, D. Answer A is incorrect because the two ports reserved for failover could be used for added communication bandwidth. Answer B is incorrect because Windows Server 2003 does not support this sort of automatic configuration capability. Answer D is incorrect because the new card would be idle instead of providing additional bandwidth.

15. Your data center recently experienced a utility power failure that took down all of the computer systems. Some systems experienced major problems (hard drive and fan failures) when the power was restored. Because of the failures, company management decides to install an Uninterruptible Power Supply (UPS) for the data center to protect the systems from another power failure. A few months later, another power failure hits the data center and the systems run for a time, then go down when the UPS runs out of power. This time, hard drive failures occur and data is lost. What was missed during the implementation of the new UPS that would have prevented the second power failure from impacting the servers?

- A. Neither the proper procedures nor the automated software controls were implemented to enable a controlled shutdown.
- B. The UPS that was purchased did not have a high enough power runtime rating to handle the load of the equipment in the data center.
- C. Windows Server 2003 does not support the use of a UPS.
- D. Windows Server 2003 does not support the use of a nondedicated UPS. Each server must have a dedicated UPS.

A. Answer A is correct because the job of a UPS is to survive short power interruptions and to allow a controlled shutdown in the event of longer power outages. Unless a communication link is established between the UPS and the servers, the servers do not know that a power failure has occurred. They will continue running until the UPS runs out of power. Alternatively, someone could have logged into the servers and manually shut them down.

B, C, D. Answer B is incorrect because the system did continue to operate for a period of time after the power failure, indicating it could handle the required power load. No UPS will run indefinitely. Answer C is incorrect because Windows Server 2003 does support the use of a UPS. Answer D is incorrect because Windows Server 2003 supports dedicated and nondedicated UPSs.

Chapter 9: Implementing Windows Cluster Services and Network Load Balancing

Making Server Clustering Part of Your High-Availability Plan

1. You have purchased a prepackaged solution that uses an eight-node majority node set (MNS) server cluster. Because you have so many nodes, you have decided to install three nodes in your Atlanta data center, three in your Denver data center, and the last two in your Seattle sales office. You notice fairly soon that the server cluster is experiencing some uptime issues. The nodes in your Atlanta data center seem to fail frequently during times of high WAN utilization. What is likely the problem?
 - A. All nodes in an MNS server cluster must be in the same data center.
 - B. The high WAN traffic is making the heartbeats take longer than 500 ms to get to all nodes and back.
 - C. The nodes in Atlanta are failing, and an MNS server cluster can have two nodes fail before losing quorum and shutting down.
 - D. The cluster cannot be in three geographic areas. An MNS server cluster can exist in a maximum of two geographic regions, and high-speed networks must connect the nodes in each region.

B. No more than 500 ms total round-trip time is allowed between nodes in an MNS server cluster. The high WAN utilization is probably making the transit time exceed 500 ms.

A, C, D. Answer A is incorrect because the MNS model is specifically designed for geographic distribution. Answer C is incorrect because an eight-node MNS server cluster can tolerate three failed nodes before the cluster shuts down. Answer D is incorrect because there is no limitation on the number of geographic areas used, nor is there any communications speed requirement beyond the 500 ms round-trip issue.
2. Your data center experiences a power failure, bringing all of your systems down. When power is returned, a single quorum device server cluster you have in use will not start. You examine the event logs and find error messages stating that the quorum drive cannot be found, yet you are able to view the contents of the quorum drive in Windows Explorer. Research reveals that either the disk signature on the quorum drive or the Registry key containing the disk signature for the quorum drive has been corrupted. What steps should you take to recover from this problem?

- A. Evict all other nodes from the server cluster, repartition and reformat the quorum drive, and rejoin the other nodes to the server cluster.
 - B. Do a restore of the quorum drive from tape.
 - C. Change the location of the quorum resource to another drive, repartition and reformat the quorum drive, and move the quorum resource back to the original quorum drive.
 - D. Shut down all nodes except one, perform an ASR restore on that node, and restart all the nodes.
- D.** Answer D is correct because properly performed ASR backups record drive signatures and layouts, including clustered drives. An ASR restore will rewrite the signatures and Registry settings.
- A, B, C.** Answer A is incorrect because this process would destroy the server cluster. Answer B is incorrect because a normal restore would not repair corrupted Registry keys or disk signatures. The Registry does not reside on the quorum drive. Answer C is incorrect because the quorum resource is not able to come online in its corrupted state. Therefore, you would not be able to move the quorum resource to another drive.
3. As a consultant, you have been called in to attempt to fix a high-availability configuration that is not performing as designed. Your client wanted to provide high availability for a high-traffic Web site. The client purchased a preconfigured, mid-range, two-node server cluster and implemented IIS on the nodes. Response time for serving Web pages is unacceptable, although there have been no incidents of the application failing over. What is the correct fix for this situation?
- A. More nodes need to be added to the server cluster. Increase the number of nodes until performance reaches an acceptable level.
 - B. Add NLB to the server cluster to handle more requests from clients simultaneously.
 - C. Convert the server cluster to an NLB cluster.
 - D. Move the server cluster to high-end hardware to provide quicker response times.
- C.** Answer C is correct because the client's stated purpose for the cluster is appropriate for an NLB cluster, not a server cluster. NLB is designed for handling large volumes of traffic. Server clusters are designed to provide increased availability for specific applications.
- A, B, D.** Answer A is incorrect. Adding more nodes will not improve performance because the application is not meant for server cluster use. Answer B is incorrect because NLB cannot coexist with a server cluster on the same hardware. Answer D is incorrect because the application will not take advantage of the higher-end hardware and will not yield an increase in performance.

4. You have been asked to design a server cluster. The server cluster will start small, but it may expand as more applications are added and predicted growth is experienced. Your proposal is for two nodes, a shared storage device, Fibre Channel host bus adapters, and switches for connectivity. When you present your proposal to management, you are asked to justify the high cost of the Fibre Channel solution. What justification can you provide for implementing Fibre Channel?
- A. Fibre Channel supports more than two nodes, allowing for the predicted growth.
 - B. Fibre Channel is the fastest connectivity solution and will therefore yield the highest performance.
 - C. Fibre Channel easily expands to allow more storage to be added to support the future applications.
 - D. All of the above.
- D.** Fibre Channel is fast, supports the full eight-node maximum cluster size, and is easily expandable by adding Fibre Channel storage devices.
- A, B, C.** Each of these answers addresses only part of the benefits of a Fibre Channel implementation.
5. You are configuring a large, single quorum device server cluster consisting of eight nodes and a dozen shared storage cabinets with 30 logical drives among them. Because of the large number of logical drives, you are using mount points instead of drive letters on most of the drives. After running the Wizard to create your first node, you can see only the drives that have been assigned drive letters. How is this resolved?
- A. Install the second node, which will automatically create mount point resources.
 - B. Manually create the disk resources after the first node is created.
 - C. Reconfigure the shared storage to reduce the number of logical drives to less than 16.
 - D. Temporarily assign drive letters to the mount point drives, and then remove the drive letters after the Wizard finishes installing the first node.
- B.** Answer B is correct because mounted drives are not automatically detected during node installation.
- A, C, D.** Answer A is incorrect because this is not the behavior of the installation of the second node. Answer C is incorrect because reducing the number of drives will not affect how mounted drives are configured. Answer D is incorrect because you cannot change drive letters on a node and a sufficient number of drive letters are not available to attempt this process.
6. You are configuring a two-node, single quorum device server cluster with a single public network interface and a single interconnect interface. The network interfaces and storage devices have been configured, and the interconnects on both nodes have been connected

with a direct crossover Ethernet cable. The installation of the first node proceeds without incident, but when attempting to create the second node, the installation fails. The Wizard reports problems communicating with the first node over the interconnect. You have verified that the cables are functional and have been properly inserted into the connectors. What is the most likely problem?

- A. The interconnect adapters are configured for auto-negotiation or for different speed and duplex settings.
- B. The direct crossover cable method cannot be used with this cluster configuration.
- C. A second interconnect is required with this cluster configuration.
- D. A switch must be used to handle heartbeat traffic.

A. The Ethernet link is not established. Most likely, this is due to conflicting speed and duplex settings. Auto-negotiation can have this result as well.

B, C, D. Answer B is incorrect because the direct crossover cable method does work with this configuration, and is quite common. Answer C is incorrect because a second interconnect is not required. The presence of a second interconnect may allow the second node to join the cluster but will not resolve the issues present on the first interconnect. Answer D is incorrect because an additional network device such as a switch is required in only a three-or-more node configuration.

7. You have installed a third-party backup agent on your nodes. The agent is supposed to listen for requests from its control server and send data to it during a backup. Despite this, your backups are failing. The application on the control server reports that it cannot communicate with the agent. You check the node and see that the agent is running properly. What is the most likely problem?

- A. The agent is not server cluster-compatible and cannot be used on a node.
- B. The control server is attempting to communicate with the agent over the interconnect network.
- C. There is a firewall between the control server and the node running the agent.
- D. The agent has configured itself to listen on the interconnect instead of the public network.

D. Answer D is correct because the installation routine for the agent most likely picked the interconnect network for listening for requests. The agent must be reconfigured to listen for requests on the public network interface.

A, B, C. Answer A is incorrect because the agent may or may not be cluster-compatible. There is not enough information provided to make that determination. Answer B is incorrect because the control server will not even be aware of the interconnect network if the interconnect is properly configured. Answer C is incorrect because, while this is a technical possibility, it does not fit into the scenario described.

8. You have created a small, two-node, single quorum device server cluster to act as a print server for several hundred printers. The shared storage is a 4GB drive. Because of the small size of the shared storage and the transient nature of the data, the print spool resource is on the quorum drive. The server cluster operates acceptably for a period of time, and then both nodes are taken down by a sudden power failure. When power is restored, the nodes boot, but the cluster service will not start. How do you fix this problem and prevent it from happening again?
- A. Delete the files under the spool directory, remove the spooler resource, add external storage, and re-create the spool on a different drive.
 - B. You cannot resolve this issue. Once the quorum drive is filled, all nodes must be evicted and the server cluster re-created.
 - C. Reformat the quorum drive and apply disk quotas to prevent the spooler from filling the drive again.
 - D. Perform an ASR restore on the nodes.
- A.** Answer A is correct because a server cluster cannot start if the quorum drive is full. Restructuring the shared storage is the only way to ensure that the problem does not reoccur.
- B, C, D.** Answer B is a false statement. Removal of files will allow the server cluster to start. Answer C is incorrect because reformatting the quorum drive will destroy the server cluster. Also, disk quotas would not necessarily resolve the issue. Answer D is incorrect because an ASR restore would not resolve the out-of-disk-space issue and allow the server cluster to restart.
9. You are configuring a large, single quorum device server cluster consisting of eight nodes and a dozen shared storage cabinets with 30 logical drives among them. The storage cabinet that contains the quorum drive also contains eight other logical drives and is connected to the last port on your 32-port Fibre Channel switch. While running the Wizard to create your first node, you cannot see any of the drives in the quorum drive's cabinet, including the quorum drive. Which of the following is a possible cause of the problem?
- A. The maximum number of logical drives recognizable in by a server cluster configuration has been exceeded.
 - B. The cabinet containing the quorum drive is not properly connected or powered on.
 - C. The cabinet containing the quorum drive must be relocated to a lower numbered Fibre Channel switch port.
 - D. The maximum number of storage devices recognizable by a cluster has been exceeded.
- B.** Given the choices, Answer B is the only possible cause of the missing quorum drive and other drives in that cabinet.

- A, C, D.** Answer A is incorrect because the number of logical drives specified (30) falls within the number of drives recognizable by Windows Server 2003 and the cluster service. Answer C is incorrect because the ordering of the Fibre Channel connected devices is unknown and irrelevant to Windows Server 2003. Answer D is incorrect because the number of storage devices is an issue only with the Fibre Channel configuration, not the configuration of the server cluster or the operating system.

Making Network Load Balancing Part of Your High-Availability Plan

10. You have installed an NLB cluster onto a 10/100 Mbps switch. Other devices, including some older 10 Mbps-only devices, are also attached to the switch. Your NLB hosts are configured for 100 Mbps and full duplex. Soon, you notice that communications with the 10 Mbps devices have failed. After troubleshooting, you discover that apparently the increased traffic on the switch is preventing the 10 Mbps devices from having sufficient bandwidth for reliable communications. What is the best fix for this problem?
- A. Change the operating mode of the NLB cluster to multicast and enable IGMP support.
 - B. Relocate all of the NLB hosts to a different virtual LAN (VLAN).
 - C. Relocate all of the 10 Mbps-only hosts to the same VLAN.
 - D. Install a firewall between the NLB hosts and the 10 Mbps-only devices and filter all NLB-oriented traffic.
- A.** Answer A is correct because the failure of the 10 Mbps-devices combined with increased traffic on the switch indicates a switch-flooding problem. Changing the NLB cluster mode to multicast with IGMP support can help resolve this issue by controlling the NLB heartbeat traffic and limiting it to those ports on the switch that are part of the NLB cluster.
- B, C, D.** Answers B and C are incorrect because this is a more complicated solution than enabling multicast and IGMP. You should not need to encounter such additional router management functions to enable NLB. Answer D is incorrect because the problem is resolvable through configuration of the NLB cluster and does not require additional software (the firewall). In addition, any requests from the 10 Mbps devices would go unresolved due to the firewall implementation.
11. You have configured an NLB cluster with 10 hosts. The default port rule has been changed from all possible ports to just port 80. No other port rules have been defined. You have configured each node with IIS and followed the appropriate procedures for installing and securing IIS. After clients begin using the cluster, you notice that clients requesting normal Web pages are being served equally across the cluster, but clients requesting

secured Web pages (SSL) and FTP sessions are all going to the host with priority 1. What is the best way to resolve this issue and to balance the SSL and FTP requests?

- A. Do nothing. SSL and FTP traffic cannot be load-balanced.
- B. Split the NLB cluster into three clusters and serve the SSL and FTP sessions from different clusters.
- C. Add new port rules for the SSL and FTP traffic.
- D. Change the default port rule back to encompass all possible ports

C. Answer C is correct because it would balance traffic on all relevant ports and provide a more secure configuration.

A, B, D. Answer A is incorrect because SSL and FTP traffic can be load-balanced if the cluster is configured properly. Answer B is incorrect because this would require an enormous amount of administrative effort and, if the same configuration steps were followed, would suffer the same problem. Answer D is technically possible, but does not provide the additional security benefit that Answer C does.

12. You are a consultant. You have been called in to troubleshoot a malfunctioning NLB cluster that serves IIS Web pages. The cluster in question consists of six hosts, but only four successfully join the cluster. Two of the hosts never successfully join. When the rest of the hosts are shut down and those two hosts are started up together, they successfully perform convergence and form a cluster. This two-host cluster, however, seems to favor certain types of incoming traffic on each host, rather than equally among the two hosts. What is the most likely reason for this behavior?

- A. The two malfunctioning hosts are configured with different cluster IP addresses and a different host name than the four correctly operating hosts.
- B. The two malfunctioning hosts are underpowered and cannot join the cluster due to poor performance.
- C. The two malfunctioning hosts are configured with different port rules than the four correctly operating hosts.
- D. The two malfunctioning hosts are configured with the same priority.

C. Answer C is correct because NLB will not allow convergence for a host or hosts that have different or different numbers of port rules. Since the two hosts do form their own cluster but the traffic pattern differs, it would appear that they are configured for the same port rules, but different rules from the other four hosts.

A, B, D. Answer A is incorrect because if the two hosts were configured with different IP addresses and a different host name from the other four hosts, you would end up with two NLB clusters. Answer B is incorrect because NLB does not perform any sort of performance test before a convergence. Answer D is incorrect because if the two hosts had identical host priorities they would not form a cluster of their own.

13. You are a consultant. You have been called in to troubleshoot a malfunctioning NLB cluster that is supposed to serve Web pages with IIS. The cluster contains four hosts, but only one host at a time will successfully form the cluster. Clients appear to have no problems connecting to any of the single-host cluster configurations. What is the most likely cause of the problem?
- A. The hosts are configured with duplicate priorities.
 - B. The hosts are configured with different port rules.
 - C. The hosts are configured with different cluster IP addresses.
 - D. The hosts are configured with duplicate cluster IP addresses.
- A.** Since each host forms the cluster individually but not together, it is likely that the hosts are configured with the same host priority.
- B, C, D.** Answer B is incorrect because the problem of duplicate port rules does not fit the behavior described. Answer C is incorrect because different IP addresses would cause clients to experience connection failures. Answer D is incorrect because hosts should be configured with duplicate cluster IP addresses. This would not yield the problem described.
14. One of your hosts in multiple-host NLB cluster requires maintenance. The cluster is heavily used and central to the profitability of your company. You want to bring the node down for service in the least disruptive way. How should you accomplish this goal?
- A. Use the drainstop option on the host needing maintenance.
 - B. Use the drainstop option on all the hosts in the cluster not needing maintenance.
 - C. Use the suspend option on the host needing maintenance.
 - D. Use the suspend option on all the hosts in the cluster not needing maintenance.
- A.** Answer A is correct because the drainstop option is used to finish servicing active requests on an active node without accepting any new connections.
- B, C, D.** Answer B is incorrect because every host except the one needing service would be shut down. Answer C is incorrect because the sessions being serviced by the host would time out and be lost. Answer D would cause every session not on the intended node to time out and be lost.
15. You have been asked to develop a design for an NLB cluster for an IIS-based Web site. The specifications given to you state that the Web application will be using server-side cookies to keep track of a visitor's session state. Which port-rule filtering mode should you configure to support the application?
- A. Single host
 - B. Multiple host/Affinity: None

- C. Multiple host/Affinity: Single
- D. Multiple host/Affinity: Class C
- C. Answer C is correct because multiple host/single affinity supports server-side cookies, meeting the given specifications.
- A, B, D. Answer A is incorrect because single-host filtering directs all traffic for a specific port to a specific host, depending on the port rule configuration. Answer B is incorrect because the Multiple host/Affinity: None mode would lead to lost state information as clients are redirected to other hosts in the cluster. Answer D is incorrect because the Multiple host/Affinity: Class C mode would appear fully functional to most clients, but could fail if clients changed IP addresses while their session state was being tracked (as happens with some ISPs).

Chapter 10: Planning, Implementing, and Maintaining Internet Protocol Security

Understanding IP Security (IPSec)

1. You have decided to deploy IPSec in your organization because you have several departments that are doing sensitive work and communicating across the Internet and other networks with a variety of persons in various organizations. There have been a few incidents where messages were sent instructing lower-level employees to perform certain tasks, purporting to be from their managers. However, investigation revealed that the managers did not send the messages; rather, they were sent by someone else, pretending to be the manager, who was attempting to sabotage the project. This experience has pointed out the need to provide authentication for the data packets that travel across the network so that the receiver of a message can be assured that it is genuine. It is equally important to ensure that the data in these messages doesn't get changed during transmission. Finally, you want to be sure that nobody other than the authorized recipient is able to read the message itself. You want the entire packet to be digitally signed, so that it will have maximum protection. Which of the following IPSec configuration choices will provide this?
 - A. Use AH alone.
 - B. Use ESP alone.
 - C. Use AH and ESP in combination.
 - D. IPSec cannot provide authentication, integrity, and confidentiality simultaneously.
 - C. Using AH and ESP in combination will provide maximum protection. AH signs the entire packet, and ESP provides the data confidentiality.
 - A, B, D. Answer A is incorrect because AH alone will provide authentication and integrity, and AH signs the entire packet, but does not provide data confidentiality.

Answer B is incorrect because, although ESP provides authentication, integrity, and data confidentiality, it does not sign the entire packet. Answer D is incorrect because, although neither of the protocols can do so alone, when AH and ESP are used together, IPSec can provide authentication, integrity, and data confidentiality simultaneously.

2. You have been hired as a consultant to help deploy IPSec for the network of a medium-size manufacturing firm that is developing a number of new products and must share sensitive data about its products over the network. As part of the planning process, you must determine the best authentication method to use with IPSec. What are the authentication methods that can be used with IPSec? (Select all that apply.)

- A. Kerberos v5
- B. Perfect Forward Secrecy (PFS)
- C. Shared secret
- D. Diffie-Hellman groups

A, C. Kerberos v5 (Answer **A**) is the default method used for IPSec authentication. A preshared key is a shared secret (Answer **C**) that can be used for IPSec authentication for interoperability in situations where one of the communicating computers does not support any other method, but Microsoft recommends that it be used only in testing situations and not on production networks. Digital certificates can also be used for authentication if you have a PKI that is functioning within the network.

B, D. Answer **B** is incorrect because PFS is used to enable the master key in IPSec. It is not used for authentication. Answer **D** is incorrect because Diffie-Hellman groups are used in IPSec for key management; they are not authentication methods.

Deploying IPSec

3. You are the network administrator for a company that has recently migrated some of its servers to Windows Server 2003 from Windows 2000. However, there are still a number of Windows 2000 servers and clients on the network. You want to use the enhanced security available on your network, and you have some interoperability issues you are concerned with pertaining to Windows Server 2003 and your Windows 2000 servers and clients. Which key method should you implement?

- A. Rivest-Shamir-Adleman (RSA)
- B. Diffie-Hellman group 1
- C. Diffie-Hellman group 2
- D. Diffie-Hellman group 2048

- C.** When concerned with interoperability issues between Windows 2000 and Windows Server 2003 machines, use Diffie-Hellman group 2 (Answer **C**) as the keying method.
- A, B, D.** Answer **A** is incorrect because RSA is an encryption algorithm; it is not a key-agreement protocol. Answer **B** is incorrect because the Diffie-Hellman group 1 key method is the least secure algorithm and the question states you wish to use enhanced security. Answer **D** is incorrect because, when dealing with interoperability issues between Windows 2000 and Windows Server 2003, you should not implement Diffie-Hellman 2048, because although it is the strongest keying method, it is new with Windows Server 2003 machines and is not supported by previous Microsoft operating systems.
4. You are a network administrator for a medium-sized medical office and you have recently deployed IPsec on the network in response to the physician/owner's concerns about confidentiality of patient information. However, it appears that IPsec might not be working correctly on a particular client computer. You need to view the local routes assigned to this particular client on the network using the IPsec Policy Agent. How does the IPsec Policy Agent function in IPsec? (Select all that apply.)
- A. Surveys the policy for configuration changes
 - B. Routes the assigned IPsec policy information to the IPsec driver
 - C. Uses the IP Security Policy Agent console to manage IPsec policies
 - D. For nondomain member clients, retrieves local IPsec policy information from the Registry
- A, B, D.** The IPsec Policy Agent surveys the policy for configuration changes (Answer **A**), routes assigned IPsec policy information to the IPsec driver (Answer **B**), and retrieves local IPsec Registry information for nondomain member clients (Answer **D**).
- C.** Answer **C** is incorrect because the IPsec Policy Agent is not a console that it used to manage IPsec policy information, and it doesn't use such a console.

Managing IPsec

5. You are the network administrator for a large law firm. You have been tasked with the duty of deploying IP security for all network communications in the departments and divisions that handle sensitive data. You have delegated individual departments to your junior administrators. You now need to verify that IPsec has been deployed and configured properly on your Human Resources and Payroll computers. Which tools can be used to perform this function? (Select all that apply.)
- A. IPsec Security Policy Monitor console
 - B. netsh command
 - C. Certificates snap-in
 - D. Resultant Set of Policy (RSOP)

- A, B.** Using the IPsec Security Policy Monitor console (Answer **A**) will allow you to monitor IPsec on the network and to verify that computers are making the expected hard associations. The netsh utility (Answer **B**) can be used at the command prompt with various switches to view configurations and monitor IPsec policies.
- C, D.** Answer C is incorrect because the Certificates snap-in cannot be used to view IPsec policy configurations. Answer D is incorrect because RSoP is used to check Group Policy for existing policy settings that can be applied.
6. You have deployed IPsec on your company's network and it has been working well, except for one thing. You've tried modifying some of the IPsec policy rules using netsh commands in the ipsec context, but each time you do so, the rules work only until you reboot the server, and then they seem to disappear. You want to make changes to the IPsec policy rules that are permanent and do not change when the server is rebooted. Which netsh command could you use?
- A. netsh ipsec dynamic set config
 B. netsh ipsec dynamic
 C. netsh interface ip
 D. netsh interface ipv6 isatap
- A.** The netsh ipsec dynamic set config command (Answer **A**) is the valid command to use to make rules permanent, relating to your IPsec even after a reboot.
- B, C, D.** Answer B is incorrect because the netsh ipsec dynamic command can be used to make the appropriate rule changes; however, after the IPsec service is stopped and restarted or the server is rebooted, the changes will be lost because they are not permanent changes; you must use the netsh ipsec dynamic set config option to make permanent changes. Answer C is incorrect because this command is used to change the netsh utility to the interface ip context to configure the TCP/IP protocol. Answer D is incorrect because netsh interface ipv6 isatap is used to configure the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), which is used for communications between IPv6 and IPv4 nodes in an IPv4 site. It has nothing to do with IPsec policy.

Addressing IPsec Security Considerations

7. You are the network administrator for a medium-sized company that provides accounting services to a number of different clients. To avoid having clients' financial information disclosed to the wrong parties, you are planning to implement IPsec on your network. You want your employees to be able to communicate securely both within the company and across the WAN with employees in your branch offices. You have recently hired a junior administrator who has his MCSE in Windows NT and 2000. You give him the task of implementing IPsec in your organization. The first thing he tells you is that because your

smaller branch office uses NAT, that site will not be able to use IPSec. What is your response?

- A. You already knew this, and intend to change that site from a NAT connection to a routed connection to accommodate this.
 - B. He is mistaken; IPSec has been able to work with NAT since Windows 2000.
 - C. He is mistaken; IPSec did not work with NAT in Windows 2000 but it does in Windows Server 2003.
 - D. You know IPSec is not compatible with NAT “out of the box,” but you can install a third-party program that will make it compatible.
- C.** IPSec and NAT were not compatible in Windows 2000, but because Windows Server 2003 has added a feature called NAT traversal, you can now use IPSec (in ESP transport mode) and NAT together; therefore Answer **C** is correct.
- A, B, D.** Answer A is incorrect because Windows Server 2003’s NAT traversal allows IPSec and NAT to work together, so there is no need to change the translated connection to a routed one for this purpose. Answer B is incorrect because IPSec was not able to work with NAT in Windows 2000. Answer D is incorrect because third-party software is not necessary for Windows Server 2003’s IPSec and NAT to work together.
8. You have been hired as network security specialist for a new startup company that has recently installed a new Windows Server 2003 network. The network was originally set up by a group of consultants, and they implemented IPSec for network communications so that communications with their secure servers could be protected. You are reviewing and evaluating the IPSec policies. Although several policies have been created, none of them seem to be effective. What do you conclude the consultants forgot to do after creating the policy?
- A. Authorize the policy in Active Directory
 - B. Assign the policy in the IP Security Policy Management console
 - C. Edit the policy after creating it
 - D. Enable the policy in the IP Security Monitor console
- B.** A policy cannot be used until it has been assigned. To assign a policy, you must right-click it in the right pane of the IP Security Policy Management console (Answer **B**) and select Assign.
- A, C, D.** Answer A is incorrect because there is no mechanism for authorizing the policy in Active Directory. Answer C is incorrect because, although you can edit a policy after creating it, not doing so would not cause it to not be applied if you had assigned it. Answer D is incorrect because you do not assign policies in the IP Security Monitor console; it is a utility used for viewing IPSec statistics and information.

9. You have been tasked with the duty of implementing IPSec on your new Windows Server 2003 network to increase security. You have never worked with IPSec before and you have been reading up on it. You've decided that you want to use PFS, but you are concerned about the resource usage on the domain controller due to reauthentication. Which of the following types of PFS can you implement without putting an undue burden on the authenticating server?
- A. You can use master key PFS.
 - B. You can use session key PFS.
 - C. You can use either or both because PFS doesn't use any resources on the domain controller.
 - D. You can use neither because both types of PFS use considerable resources on the domain controller.
- B.** Session key PFS (Answer **B**), unlike master key PFS, does not force reauthentication and does not place nearly so great a burden on the domain controller.
- A, C, D.** Answer A is incorrect because master key PFS forces reauthentication of the master key keying material each time a new session key is required. This process uses a considerable amount of resources to function and could adversely affect the domain controller's performance. Answer C is incorrect because master key PFS does use considerable resources on the domain controller and therefore it is not an acceptable choice to implement when you need to take the domain controller's performance into consideration. Answer D is incorrect because only master key PFS, not session key PFS, puts a burden on the domain controller. Both of these options would not be a choice when implementing PFS types.
10. You are creating a project to implement IPSec using the IPv6 protocol. Part of your security plan states that you must maintain data confidentiality as part of your IPSec implementation. When developing your plan further, what must you remember about Microsoft's implementation of IPv6 that is included in Windows Server 2003?
- A. IPv6 does not support data encryption.
 - B. IPv6 does not support authentication.
 - C. IPv6 does not support integrity.
 - D. IPv6 does not support IPSec.
- A.** IPv6, as implemented in the Windows Server 2003 family, does not support the use of IPSec data confidentiality, which is obtained by ESP data encryption; therefore Answer **A** is correct.
- B, C, D.** Answer B is incorrect because IPv6 does support authentication as implemented with Windows Server 2003. Answer C is incorrect because IPv6 does support integrity as implemented with Windows Server 2003. Answer D is incorrect because

IPv6 does support IPSec. In fact, this was one of the major design goals of version 6 of the IP protocol.

11. You have been hired as a consultant to evaluate the IPSec deployment in a small music publishing company. Management is concerned that copyrighted material might be intercepted as it passes over the network and be stolen. You discover that the former network administrator who initially set up IPSec configured it to use the AH protocol only. You explain to the company manager that one of the things you recommend changing is to configure IPSec to use ESP. Why would you implement ESP in this situation? (Select all that apply.)
- A. ESP ensures data integrity and authentication.
 - B. ESP prevents capture of packets.
 - C. ESP provides confidentiality.
 - D. ESP encrypts the packets.
- C, D.** If you need to have confidentiality and encryption on the packets, you should use ESP; therefore Answers **C** and **D** are correct.
- A, B.** Answers A and B are incorrect because ESP does not provide for data integrity and authentication, nor does it prevent packets from being captured.
12. You are on an IT team that is planning the deployment of IPSec throughout a large enterprise network. You have been advised that cost-effectiveness and efficient use of personnel are two priorities, because the company does not want to hire additional IT staff to support the deployment. Of the authentication methods available, which has the lowest administrative overhead and is the most efficient if you wish to support the implementation on 10,000 client machines?
- A. Diffie-Hellman group 2048
 - B. Kerberos v5
 - C. Pre-shared keys
 - D. Digital certificates
- B.** Kerberos v5 (Answer **B**) has the lowest administrative overhead and is the easiest to support of the authentication methods.
- A, C, D.** Answer A is incorrect because Diffie-Hellman is a key-exchange protocol, not an authentication method. Answer C is incorrect because pre-shared keys must be entered into each client machine manually, creating a large amount of administrative overhead. Answer D is incorrect because digital certificates require the implementation of a PKI and setup and maintenance of certification authorities. After the PKI is set up, administrative overhead is lower than that involved with pre-shared keys, but initial overhead is very high.

Using RSoP for IPsec Planning

13. You have been hired to manage security for a medium-sized network. Your first project is to implement IPsec on the network to protect communications that travel across it. You have just assigned an IPsec policy to a client, and you need to view the precedence of IPsec policy assignments and which policies have been assigned to the client. Which logging mode would you use in RSoP?
- A. IPsec mode
 - B. RSoP mode
 - C. Logging mode
 - D. Planning mode
- C.** You would use the logging mode in RSoP (Answer C) for this purpose because it will show you which policies have taken precedence over others. It also shows detailed policy information such as filters, connection types, and tunnel endpoints.
- A, B, D.** Answers A and B are incorrect because they do not exist as mode types for RSoP. Answer D is incorrect because the planning mode can run queries to show administrators which policies are assigned to which users, as well as the names of the target client computer name, IP address, and domain controller assignment from the Windows Management Instrumentation (WMI).
14. You have IPsec configured and running on your network. You want to capture some IPsec packets to ensure that the data inside cannot be viewed. You want to capture packets being sent from a remote client to a remote server, using a server in the server room. Which of the following tools will you need to use in order to capture these packets?
- A. Network Monitor in Windows Server 2003
 - B. netsh commands in the ipsec context
 - C. The IP Security Monitor console
 - D. Systems Management Server (SMS)
- D.** To capture packets and view what is inside them, you need a network sniffer (protocol analyzer). The only tool in this list that will allow you to capture and view packets passing across machines on the network other than the one from which you are monitoring is the version of Network Monitor that is included in Microsoft's SMS console software, which can place the network card in promiscuous mode so that traffic not sent or received by the local computer can still be captured. Therefore Answer D is correct.
- A, B, C.** Answer A is incorrect because the Network Monitor included in Windows Server 2003 can capture packets, but only those sent to or from the local computer on which the Network Monitor is installed. Answer B is incorrect because the netsh com-

mand-line utility is used to apply various IPSec policies and cannot be used to view network traffic. Answer C is incorrect because the IPSec Monitor is used to view statistics and information about IPSec connections, but it does not allow you to view inside individual packets.

15. You want to use the RSoP tool in logging mode to build some reports on the existing policy settings of one of your client computers. You have used RSoP before in planning mode, but never in logging mode. You open the RSoP Wizard from the Active Directory Users and Computers console, as you've done before, but you notice that there is no mechanism for selecting the mode, and only planning mode seems to be available. What is the problem?
- A. The RSoP Wizard runs only in planning mode.
 - B. You should open the RSoP Wizard from Active Directory Sites and Services instead.
 - C. You should open the RSoP Wizard from the RSoP MMC instead.
 - D. You can select logging mode when you open the RSoP in Active Directory Users and Computers. You must have overlooked the option.
- C.** Answer C is correct. When you open the RSoP Wizard from either Active Directory Users and Computers or Active Directory Sites and Services, you can use only planning mode. To use logging mode, you must open a stand-alone RSoP MMC. This is done by selecting Start | Run, and entering mmc. Then select File from the menu, choose the Add/Remove Snap-in, then Add. Then you can scroll down the list and add the RSoP console by double-clicking the Resultant Set of Policy and selecting Add. After the console has been added, select the Close button and then select OK.
- A, B, D.** Answer A is incorrect because the RSoP Wizard can run in either planning or logging mode, but the available modes depend on how you open the Wizard. Answer B is incorrect because opening the Wizard from the Active Directory Sites and Services tool would not help; you would still have only planning mode available. Answer D is incorrect because there is no way to select logging mode when you use Active Directory Users and Computers to open the RSoP Wizard; only planning mode is available.

Chapter 11: Planning, Implementing, and Maintaining a Security Framework

Planning and Implementing Active Directory Security

1. You have instituted new security policies for the IT department. One important rule is to never log on as Administrator unless it is absolutely necessary. To enhance security, you want everyone to use their regular user accounts for everyday tasks so you can maintain

security as much as possible. A junior administrator comes to you and says he does not wish to log on to the server with an administrative account, but he needs to use a program that requires administrative privileges. What can he do?

- A. If running the program requires administrative privileges, he cannot run it unless he logs off and logs back on as Administrator.
- B. He can open the Computer Management console and use the Set password option.
- C. He can right-click the program he wants to run, select Properties, click the Advanced button, and configure the program to run without administrative privileges.
- D. He can right-click the program, choose the Run as command, and enter the Administrator account name and password.

D. Best security practice is to log on with a regular user account, and then use the Run as option (which uses the secondary logon service) to run programs that require administrative privileges.

A, B, C. Answer A is incorrect because, beginning with Windows 2000, Microsoft has provided a way to run programs with administrative credentials, even though you are not logged on as Administrator. Answer B is incorrect because the Set Password option in the Computer Management console is used to change your password; it does not affect the privileges with which you run a program. Answer C is incorrect because the Advanced button in the Properties sheet of a program allows you to compress or encrypt the program file and set archiving attributes. It does not provide any way for configuring the program to run without administrative privileges.

2. You have been hired as the network administrator for a small law firm. The first thing you want to do when you take over the job is increase the security on the network. You evaluate the current security level and find it lacking. You decide that you need to secure account passwords using strong encryption on domain controllers. Which utility should you use?

- A. System Key Utility
- B. Secedit
- C. MBSA
- D. SUS

A. The System Key Utility (syskey) will provide strong encryption techniques so account password information remains strong. This provides an extra line of defense against attacks by password-cracking software that targets the directory services for stored passwords.

B, C, D. Answer B is incorrect because Secedit is a command-line tool that is used to configure and analyze system security by comparing the current configuration with one or more templates. Answer C is incorrect because the Microsoft Baseline Security

Analyzer (MBSA) tool is used to analyze and correct security; it is not used to encrypt passwords on domain controllers. Answer D is incorrect because the Software Update Services (SUS) utility is used to apply administrator-approved security fixes and patches.

3. You have recently hired a new junior administrator to assist you in running the network for a medium-sized manufacturing company. You are explaining to your new assistant that AD objects are assigned security descriptors to allow you to implement access control. You tell your assistant that the security descriptor contains several different components. Which of the following are contained in the security descriptor for an object? (Select all that apply.)
- A. Discretionary access control list
 - B. System access control list
 - C. Dynamic access control list
 - D. Ownership information
- A, B, D.** The security descriptor contains the discretionary access control list (DACL), which has information about which groups and users are allowed or denied access to the object. The security descriptor also includes a system access control list (SACL), which specifies which events should be audited for this object if auditing is enabled. The third component of the security descriptor is the ownership information that identifies who owns the object.
- C.** This answer is incorrect because Windows Server 2003 does not support dynamic access controls, which allows access information to be changed “on the fly” and access granted based on the information at the time the request is made, instead of requiring that a user log off and back on before group membership changes take effect.
4. You are attempting to troubleshoot some problems with access that you think can be traced back to membership in multiple groups. You want to ensure that all administrative accounts are able to perform the tasks they need to accomplish, but you want to remove the built-in accounts from all groups to which they’ve been added by another administrator, and give them only the access they had by default. You are a little confused because you know that the built-in accounts already belong to some groups at installation, and you don’t want to remove them from groups they are supposed to belong to. To which groups does the Domain Administrator account belong in Windows Server 2003 by default? (Select all that apply.)
- A. Schema Admins
 - B. Enterprise Admins
 - C. Group Policy Creator Owners
 - D. Backup Operators

- A, B, C.** The Domain Administrator account has total control over every function of the domain and network by default. This account belongs to the Administrator, Schema Admins, Enterprise Admins, and Group Policy Creator Owners groups.
- D.** The Backup Operators group is used to give some users limited privileges so they can back up and restore data without having additional administrative permissions. The Domain Administrator account already has the proper permissions to back up and restore all data, so it doesn't need to be a member of the Backup Operators group.

Planning and Implementing Wireless Security

5. You want to allow wireless clients the ability to change their passwords after they authenticate on the network. Which method of authentication should you implement for these clients?
 - A. EAP-TLS
 - B. EAP
 - C. PEAP
 - D. EAP-MS-CHAP v2
 - D.** EAP-MS-CHAPv2 is the authentication method to use when you wish to allow clients to change their passwords after they have been authenticated on the network.
 - A, B, C.** EAP-TLS does not support the ability for clients to change their passwords, so Answer A is incorrect. EAP also does not support the ability for clients to change their passwords, so Answer B is incorrect. Answer C is incorrect because PEAP also does not support the ability of clients to change their passwords.

6. You are implementing a new wireless network and need to change the default settings for the equipment on the WLAN. What information should you change? (Select all that apply.)
 - A. SSID password
 - B. SSID network name
 - C. Domain Administrator password
 - D. Domain Administrator account should be renamed
 - A, B.** Because so many wireless manufacturers deliver their equipment preconfigured with the same SSID password and network name on all devices, it is imperative that you change this information before you use the wireless network in a live environment. It is also a good idea to disable SSID broadcasting so hackers won't be able to so easily discover the SSID, although they can still use a sniffer to capture packets being transmitted and determine the SSID from that.

- C, D.** Answer C is incorrect because the Domain Administrator password is not a default setting for the wireless equipment, although it is important that you change it frequently since the Domain Administrator account has full control over your network. Answer D is incorrect because the Domain Administrator name is a default setting for Windows Server 2003 but is not a default setting for WLAN equipment. However, renaming this account is a good idea, since hackers know of its existence. If you leave it at the default, hackers will have half the information they need (username and password) to gain control of your network.
7. You have a number of users who need to be able to roam through the building with their laptop computers and still stay connected to the network. Because of the nature of their work, it is important that they have relatively fast access for transferring a lot of very large data files over the network. You need to implement a wireless network that can connect devices up to 54 Mbps and a minimum of 24 Mbps. Which IEEE standard should you choose?
- A. 802.15
 - B. 802.11a
 - C. 802.11b
 - D. 802.1x
- B.** The 802.11a standard can handle up to 54 Mbps. It is used to connect schools and business with wireless technology and is especially appropriate in cases where faster access is required. However, its range is shorter than 802.11b, so you will need to place more access points closer together throughout the building.
- A, C, D.** 802.15 is used for Bluetooth technology. It can travel only about 10 meters and could not handle the minimum 24 Mbps, so Answer A is incorrect. 802.11b is a standard widely in use today, but it can handle only up to 22 Mbps, so Answer C is incorrect. Answer D is incorrect because 802.1x is a wireless security standard.
8. You have hired a consultant to help set up wireless access points on your network. He tells you that you should turn on WEP for the wireless network to help protect it from intruders. You tell him that you have heard that WEP has many flaws and you think additional security measures should be implemented. He assures you that WEP works fine. What do you tell him are some of the problems with WEP?
- A. WEP does not use encryption.
 - B. WEP uses a short (24 bit) initialization vector (IV).
 - C. WEP can use only a 40-bit key.
 - D. WEP uses a public key algorithm.

- B.** The 24-bit IV makes WEP especially vulnerable because, even when a longer key is used in conjunction with it, the short IV ensures that the key stream will be reused. A hacker can capture multiple packets, analyze them, and perform an XOR operation to discover the plaintext and break the encryption, or use software such as WEPCrack.
- A, C, D.** Answer A is incorrect because WEP does use encryption, but it uses weak encryption. The RC4 algorithm is a secret key (symmetric) method and the same key is shared among all clients. Answer C is incorrect because WEP can use either a 40- or 104-bit key, but it uses the same 24-bit IV, regardless of the key length. Answer D is incorrect because WEP does not use a public key (asymmetric) algorithm, which would be more secure. With a public key algorithm, there is no secret key shared among all the clients.

Monitoring and Optimizing Security

9. Your junior administrator wants to change the name of a user account, but he is worried that if he does so, the user will have problems accessing resources that she had previously been given permissions for. The administrator doesn't want to need to re-create all the group memberships for the newly named account. You tell him there is no need to worry; he can go ahead and change the name, and all the account properties will remain intact. What enables an account to retain its password, profile, group membership, user rights, and membership information?
- A. Group membership of the account
- B. Domain the account belongs as a member
- C. Password encryption method
- D. Security identifier (SID)
- D.** The SID enables an account to retain all of its information such as password, profile, group membership, and user rights. Even if the account is renamed, the SID does not change and the account still retains all of this information.
- A, B, C.** None of these answers has any effect on an account's ability to retain its network information. The group membership of the account affects only rights and permissions, the domain membership information gives the account access to domain specific information, and the method by which the password is encrypted has no bearing on the user account.
10. You suspect that one of your users has been trying to access data in a folder to which he is not supposed to have permission. You are trying to set auditing on this folder so you can see if there are any failed events in the log indicating that the user did try to open the folder. You enable object auditing in the domain's Group Policy Object. However, when you go to add this user to be audited for access to the folder, you find that the folder's property pages do not contain a Security tab. What could be the problem?

- A. Auditing is not set via the Security tab for folders because they don't have such a tab.
 - B. You cannot audit folder access for a particular user.
 - C. The folder is not on an NTFS partition.
 - D. You must share the folder before you can audit it.
- C.** NTFS is required for auditing. If a folder is on a FAT or FAT32 partition, you cannot set security permissions or configure auditing because no Security tab will appear in the folder's properties. You can move the folder to an NTFS partition or you can use the Convert command to upgrade the file system of the partition to NTFS without losing any data (however, this is a one-way process).
- A, B, D.** Answer A is incorrect because folders on NTFS partitions do have a Security tab in their properties, and this is where you configure auditing for them, using the Advanced button. Answer B is incorrect because you can audit folder access for a user or a group, as long as object auditing has been enabled in Group Policy and the partition is formatted with NTFS. Answer D is incorrect because you can set auditing on any folder on NTFS partition when object auditing is enabled; it does not need to be shared (however, if it isn't shared, remote users will not be able to access it; you'll only be auditing access by users logged on locally).

Planning a Change and Configuration Management Framework

11. You need to configure Kerberos policies because you want to force user logon restrictions. You go to the computer of the user on whom you want to enforce these policies and access the Local Security Policy. However, in the GPO Editor, you cannot find Kerberos policies in the Security Settings node under Computer Configuration, under Windows Settings. What is the problem?
- A. You are looking in the wrong section; Kerberos policies are located in the User Configuration node.
 - B. You cannot set Kerberos policies through the Local Security Policy console.
 - C. You must first raise the domain functional level before Kerberos can be used and this option will appear in the GPO.
 - D. Another administrator has deleted the Kerberos policies node from the GPO.
- B.** Kerberos policies can be set for domains only, not for local computers. You must edit a domain GPO to find the Kerberos policies option. To access these policies, expand Computer Configuration, then Windows Settings, then Security Settings.
- A, C, D.** Answer A is incorrect because the Kerberos policies are located in the node in which you are looking, but in a different GPO (one that is applied to a domain

instead of a local computer). Answer C is incorrect because the domain functional level does not affect the appearance of this option in the GPO; all Windows 2000 and 2003 domains use Kerberos authentication. Answer D is incorrect because an administrator cannot remove the Kerberos policies from a Local Security Policy GPO—it was never there to begin with.

12. You have been analyzing all of your security configuration information as part of a new project that requires you to provide a detailed report on your network's security to management. Toward that end, you need to evaluate the security database test.sdb at the command prompt. What command can you use to do this?
- A. `secedit /validate test.sdb`
 - B. `secedit /analyze test.sdb`
 - C. `secedit /configure test.sdb`
 - D. `secedit /export test.sdb`
- B.** The `secedit /analyze test.sdb` command is the appropriate command to use to analyze the test.sdb security database. You must use the valid switch after the command, as well as the name of the security database.
- A, C, D.** The `secedit /validate test.sdb` command is used to validate security settings with the `secedit` command, so Answer A is incorrect. The `secedit /configure test.sdb` command is used to configure the security database, so Answer C is incorrect. The `secedit /export test.sdb` command is used to export the security database, so Answer D is incorrect.
13. You want to set up auditing on several folders that contain important and sensitive information. There are other folders within the specified folders that contain less sensitive information, so you don't want to audit them, because you want to put as little overhead burden on the network as you can. What happens to subfolders and files within a parent folder if auditing has been enabled?
- A. Subfolders only are audited
 - B. Files only are audited; special access must be turned on for the folders to be audited
 - C. Subfolders and files are audited
 - D. No auditing is performed
- C.** By default, if auditing is turned on for a parent folder, all subfolders and files within that folder are audited as well. This option can be changed by using the **Apply Onto** box in the **Auditing Entry for File or Folder** dialog box and choosing **This folder only**.
- A, B, D.** Answer A is incorrect because subfolders and files are both audited when the parent folder has auditing enabled. Answer B is incorrect because no special access

needs to be turned on for folders to be audited when the parent folder has auditing enabled. Answer D is incorrect because, by default, the files and folders in the parent folder will automatically be audited.

14. A parent folder has auditing enabled. Two folders, Applications and Phone Listings, are listed under this parent folder. You need to have the Phone Listings folder audited but not the Applications folder. How can this be accomplished?
- A. It cannot; all subfolders are audited when the parent folder has auditing enabled.
 - B. Right-click the Applications folder, and click the **Properties** tab, select the **Security** tab, and click **Advanced**. Then select the **Auditing** tab and clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Include these with entries explicitly defined here**.
 - C. Right-click the **Phone Listings** folder, click the **Properties** tab, select the **Security** tab, and click **Advanced**. Then select the Auditing tab and clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Audit entries defined here**.
 - D. Right-click the **Phone Listings** folder, click the **Security** tab, and click **Advanced**. Then select the Auditing tab and clear the check box that is labeled **Inherit from parent the auditing entries that apply to child objects. Include these with entries explicitly defined here option**.
- B.** This is the correct procedure to use to turn off auditing for the Applications folder while leaving the Phone Listings folder with auditing inherited from its parent folder.
- A, C, D.** Answer A is incorrect because, although by default subfolders inherit the audit setting of the parent folder, this can be changed. Answers C and D are incorrect because neither of these procedures will provide the desired result.

Planning a Security Update Infrastructure

15. You need to install the Microsoft Software Update Services (SUS) within your domain to update security information on client computers. What are the minimum requirements that you should use for hardware for the server?
- A. Pentium III, 256MB RAM, NTFS with a minimum of 50MB for the installation folder and 6GB for SUS updates and Active Directory installed
 - B. Pentium III, 512MB RAM, NTFS with a minimum of 100MB for the installation folder and 6GB for SUS updates without Active Directory installed
 - C. Pentium III, 256MB RAM, NTFS with a minimum of 25MB for the installation folder and 6GB for SUS updates without Active Directory installed

- D. Pentium III, 512MB RAM, NTFS with a minimum of 50MB for the installation folder and 5GB for SUS updates and Active Directory installed
- B.** At a minimum, the hardware requirements should be Pentium III, 512MB RAM, NTFS with a minimum of 100MB for the installation folder and 6GB for SUS updates. If your hardware does not meet these requirements, you could have trouble running the software.
- A, C, D.** Answer A is incorrect because the memory level is too low, the install folder is too small, and Active Directory cannot be installed on a server on which you need to install SUS. Answer C is incorrect because the installation folder is too small and the amount of RAM is too low. Answer D is incorrect because there is not enough disk space for either the installation folder or the SUS updates.

Chapter 12: Planning, Implementing, and Maintaining a Public Key Infrastructure

1. You are setting up a procedure to keep documents exchanged between members of the R & D department secret. They will be sending these documents across the Internet to each other. Which PKI process will you need to employ to achieve this?
 - A. Confidentiality
 - B. Non-repudiation
 - C. Authentication
 - D. Data Integrity

A. PKI confidentiality is the encryption of data to keep only those without the proper credentials/authority from accessing the data; therefore, Answer A is correct.

Answer **B** is incorrect because non-repudiation provides the receiver with a guarantee that the sender cannot deny the origin of the data. Answer **C** is incorrect because authentication is the act of verifying the identity of the sender. Answer **D** is incorrect because data integrity guarantees that the data is unchanged from the time the sender sent or saved it.

Implementing Certification Authorities

2. You are the administrator for a large and very busy network and your bandwidth is nearing its limits. Your users are complaining about the time it takes to access the payroll server to update their hours. All users are required to have certificate authentication to access the server. What can you change in your current setup to help reduce network traffic and speed access to the payroll server?

- A. Configure the CA to use complete CRLs for replication.
 - B. Assign times for each user to update their payroll.
 - C. Use DES for the encryption method.
 - D. Configure the CA to use Delta CRLs.
- D.** Delta CRLs replicate only the new revocations to each CRL distribution point. This means a smaller file, hence less network traffic. This allows for more frequent but much smaller data transfers. Therefore, Answer D is correct.
- A, B, C.** Using complete CRLs is the default setting which means the entire CRL is replicated at the specified interval. If there is no new information the CRL is published regardless. Giving each user a time frame to update might work but the administrative overhead would be tremendous and there would be no guarantee the users would cooperate.
3. Your department has completed preliminary testing of a newly established PKI, and before actual deployment begins, you've been assigned the task of revoking the test certificates. So far, there is only a single enterprise CA installed, and Active Directory is of course in use. Which of the following steps should you take?
- A. In the **Certification Authority** console, expand the **Issued Certificates** container, and revoke all certificates by right-clicking each certificate and choosing **All | Revoke Certificate**.
 - B. In the **Certification Authority** console, expand the **Issued Certificates** container, and revoke all certificates by right-clicking each certificate and choosing **All | Revoke Certificate**. Right-click the **Revoked Certificates** container, and choose **All Tasks | Publish**.
 - C. Using the **Certificates** snap-in, expand the **Personal** container, and highlight the **Certificates** container found beneath. In the right pane of the console, right-click each certificate and choose **Add to Certificate Revocation List**.
 - D. Using the **Certificates** snap-in, expand the **Personal** container, and highlight the **Certificates** container found beneath. In the right pane of the console, right-click each certificate and choose **Add to Certificate Revocation List**. Right-click the **Trusted Root Certification Authority**, and choose **Publish to Directory**.
- B.** Clients only check the CRL, not the CA itself, for valid certificates. A revoked certificate must be distributed to the CRL, and the CRL must then be downloaded by the client, before the certificate is rendered completely invalid. Therefore, Answer **B** is correct.
- Answer **A** is incorrect because revoking a certificate is not sufficient to render a certificate invalid in a Windows PKI. If the certificate is not published to the CRL, clients will not be aware that the certificate has been revoked. Answer **C** is incorrect because

the Certificates snap-in cannot be used to revoke and publish certificates (there is no such command). In fact, the Certificates snap-in is used on the client machine, generally not the CA. Answer **D** is similarly incorrect.

4. You decide to implement a Windows Server 2003 based-PKI for your network, and because you want the most secure method of issuing and maintaining certificates, you decide to use a stand-alone server to issue a certificate to a subordinate, which in turn issues certificates to users. You take the root CA offline. Your users complain that they are unable to access some resources. After investigating the problem you discover that they can log on to the network and access everything except those resources protected by certificates. They also can connect to the servers by both name and IP address. What is preventing the users from gaining access to those resources?
- A. The root CA server is offline.
 - B. The subordinate CA is offline.
 - C. The certificates have been compromised.
 - D. The certificates are still pending.
- D.** By default the certificates on a stand-alone CA are issued manually by an administrator. Certificates will be in “pending” status until the administrator issues them; therefore, Answer **D** is correct.
- Answer **A** is incorrect because the root CA should be taken offline for increased security. Answer **B** is incorrect because, since the subordinate will be issuing the certificates to users this would not affect the capability to receive or verify a certificate. Answer **C** is incorrect because you have full connectivity using name and IP address, which verifies your DNS and network address settings and operation and access to the resource would not be affected even if the certificates had been compromised. Adding the compromised certificates to the CRL would be the appropriate action to take in this case.
5. You have a two-tier hierarchy for your certificate PKI. *OurRoot* is an enterprise root CA. *OurIssuer1* and *OurIssuer2* are *OurRoot* subordinates. These two CAs issue all the certificates for your company. *OurIssuer1* issues to the northern region and *OurIssuer2* issues to the southern region. An ex-employee appears to have obtained the issuing certificate for *OurIssuer2*. What steps would you take to prevent users from using certificates issued by the compromised server?
- A. Add the compromised certificate to the CRL from *OurRoot*.
 - B. Delete all certificates on *OurIssuer2* and reissue them.
 - C. Reinstall certificate services on *OurIssuer2*.
 - D. Add all certificates issued by *OurRoot* to the CRL.

- A.** The compromised certificate would be revoked and all certificates issued using that certificate would be revoked as a result. Don't forget to force a CRL update (**Revoked Certificates** | **All Tasks** | **Publish** in the CA snap-in).
- Answer **B** is incorrect because deleting all certificates and reissuing them would be extremely time consuming (every machine that had a certificate from *OurIssuer2* would have to be examined). Reissuing them would provide new certificates, but since the issuing certificate is compromised, this would not prevent anyone from issuing bogus certificates. Answer **C** is incorrect because reinstalling would not change the validity of current certificates or those issued by a compromised certificate. Answer **D** is incorrect because revoking all certificates issued by the root CA would also invalidate all the certificates issued by *OurIssuer1*. This would cause needless problems in the northern region.
6. As a member of the PKI design team in your company, you are charged with integrating one of your subsidiaries that already has a PKI with your office's PKI. The current proposal on the table has a second-tier CA located in your local PKI issuing certificates to a second-tier CA located on the subsidiary's PKI, and vice-versa. Both infrastructures are Windows Server 2003 based. Your company's security goals, however, mandate that only certain certificates be used on your PKI if they are issued from the subsidiary's CA, but all your CA's certificates need to be trusted by the subsidiary. What is your assessment?
- A. Both your office and the subsidiary will need to create a CTL that has a limited trust chain length on your side.
- B. The subsidiary's CA needs to be reconfigured as your CA's subordinate.
- C. A cross-trust needs to be created, and the type of acceptable certificates for your CA narrowed by using qualified subordination policies.
- D. This arrangement is not possible under Windows Server 2003. The company needs to implement a third-party PKI.
- C.** Windows Server 2003 is capable of using qualified subordination. This new feature enables restrictions to be placed on the level and depth of trust that exists in a cross-trust relationship; therefore, Answer **C** is correct.
- Answer **A** is incorrect because a Certificate Trust List (CTL) by itself is not capable of limiting the types of acceptable certificates on one partner. Also, the length of the trust chain is not the main concern in this situation. Answer **B** is incorrect because configuring the subsidiary's CA to be a subordinate of yours will not change the types of certificates that are trusted. It only changes the number of levels in the hierarchy. Answer **D** is incorrect because using qualified subordination properly, the situation described can be implemented without having to rely on a third-party company.

7. Your company has a partner with whom you need to communicate securely. You have an existing root CA and need to allow usage for partner-issued certificates as well. In which of the following ways can you accomplish this? Choose all that apply.
- A. Create a CTL.
 - B. Install an issuing CA at the partner's site.
 - C. Create a cross-trust hierarchy.
 - D. Install a partner's issuing CA at your site.
- A, C.** A CTL or cross-trust hierarchy will enable your CA to recognize a certificate issued by a CA that is not a part of your network; therefore, Answers A and C are correct.
- Answer **B** is incorrect because installing an issuing CA at your partner's site enables them to get a certificate from your hierarchy but doesn't enable your resources to recognize the partner-issued certificates. Answer **D** is incorrect because installing a partner's CA at your location enables you to get certificates for their network but doesn't enable them to access your PKI.
8. You are the administrator of an existing three-tier PKI including a stand-alone Root CA, three mid-level CAs, and twelve issuing CAs. You fear that your Root certificate has been compromised. What steps should you take to secure your infrastructure with the least amount of administrative effort?
- A. Add the twelve issuing CAs' certificates to the mid-level CAs' CRL.
 - B. Add the three mid-level CAs' certificates to the Root CA's CRL.
 - C. Add the Root CA's certificate to the three mid-level and twelve issuing CAs' CRL.
 - D. Create a new CA hierarchy and issue new certificates to all clients.
- D.** If your Root CA is compromised, all certificates that are in that hierarchy become compromised also. This is the reason that the logical and physical security of the Root is so important in any PKI – the cost of rebuilding after a successful attack on the Root can be enormous. Therefore, Answer D is correct.
- Answers **A** and **B** are incorrect because all certificate trust hierarchies begin with the Root CA. Revoking either the mid-level or the issuing certificates alone is insufficient in this case. If an issuing CA had been compromised, then adding its certificate to the upstream CA's CRL would work, and similarly if a mid-level CA had been compromised, then adding its certificate to the Root's CRL would be reasonable. However, there is no upstream CA to a Root. Answer **C** is incorrect because you cannot add an upstream, and therefore trusted, certificate to a CRL.

Planning Enrollment and Distribution of Certificates

9. You are attempting to request a certificate by using Internet Explorer, but fail to display the welcome screen of the Web site. You have typed in the address *http://mycertauthority/certsrv* and you've double-checked the name of the CA. Also, you have confirmed with the network administrator that the CA is configured with IIS, and the Web enrollment support option was chosen during the certificate services installation. What is the most likely cause of the problem?
- A. The CA is configured as a standalone.
 - B. IIS was installed after certificate services.
 - C. The EAP protocol has not been installed.
 - D. You are using a Windows 2000 Professional client.
- B.** If IIS is installed after certificate services, even if the Web enrollment support option is chosen, the appropriate virtual directories are not created. To remedy the situation, you can use the command line tool **certutil -vroot**. Therefore, Answer B is correct.
- Answer **A** is incorrect because as long as the CA is on the network and has IIS installed properly, Web enrollment by clients is possible. Answer **C** is incorrect because EAP is not necessary when using a browser to request a certificate. EAP is used primarily for smart card authentication. Answer **D** is incorrect because a Windows 2000 Professional client comes pre-installed with a version of Internet Explorer higher than 5.0, which is the minimum requirement.
10. The Ecstatic Llama Company wants your consulting firm to implement a two-tier private CA design made specifically for their PKI. Because the plans for ELC call for high security, the root CA will be designated as standalone and offline. Your job is to install an enterprise subordinate CA while maintaining the security needs of your client. What are the two best methods to accomplish this task? Choose two answers.
- A. In the Certification Authority console, configure the subordinate to use auto-enrollment and reboot the machine.
 - B. In the Certification Authority console, point the subordinate to use Active Directory and configure the subordinate to trust the root CA.
 - C. Put the root CA briefly online and use Web enrollment to obtain the root CA certificate, then take the root CA back offline.
 - D. Save the subordinate request as a PKCS #10 file, transport the file to the root CA, issue the certificate, and then transport the certificate back to the subordinate.
- C, D.** Answer **C** is correct because any client can use a browser to obtain certificates from a CA, even the root CA. As long as the root CA is online, running IIS, and has Web enrollment support installed, this method will succeed. Answer **D** is correct

because of Windows' capability to create a file containing a certificate request. PKCS #10 is the standard PKI request form, and after the file is saved to disk, the disk can be physically transported to the root CA. After the root CA generates the certificate, the certificate itself can be transported back to the subordinate.

Answer **A** is incorrect because auto-enrollment cannot be used to request certificates from a stand-alone CA. Also, the Certification Authority console would not function properly until the subordinate had actually been installed. Answer **B** is incorrect because there is no direct option to configure the use of Active Directory, and Active Directory cannot be used with a stand-alone machine. Also, there is no option in the Certification Authority console to “trust” another CA. Trust is established by validating other PKI entities' certificates.

11. You are the CA administrator for your branch office and want to have greater control over your certificate managers. Your plan is to have each manager manage certificates over a different Active Directory group, but you do not want to give any manager the capability to renew the CA's certificate. What is your best course of action?

A. In the **Certification Authority** snap-in, use the **Security** tab of the CA's property sheet to configure manager restrictions.

B. Using the **Certificate Templates** snap-in, right-click the **Certificate Templates** container, and choose **Properties**. On the **Security** tab, give the Certificate Managers group the *Issue and Manage Certificates* permission.

C. In the **Certification Authority** snap-in, use the **Certificate Managers Restrictions** tab of the CA's property sheet and choose the **Restrict certificate managers** option.

D. It cannot be done.

C. The options on the **Certificate Managers Restrictions** tab enable you to grant or deny each manager's capability to manage users, groups, and computers. Renewing the CA's certificate is a capability given only to the CA administrator with *Manage CA* permission. Therefore, Answer **C** is correct.

Answer **A** is incorrect because the **Security** tab allows only *Read, Issue and Manage Certificates, Manage CA, and Request Certificates* permissions. In fact, it is the *Issue and Manage Certificates* permission itself that defines the role of certificate manager. You cannot dictate what groups a manager has control over using this tab. Answer **B** is incorrect for similar reasons. Also, there is not an *Issue and Manage Certificates* permission on the **Security** tab of the **Certificate Template** container's property sheet. Answer **D** is incorrect because using the **Certificate Managers Restrictions** tab of the CA's property sheet enables you to properly restrict control.

12. As the network administrator for B & H Day Care Centers, you are attempting to configure a third-tier CA to issue a particular type of certificate. From the **Certificate**

Templates snap-in, you have duplicated an existing template and modified it to B & H's specifications. However, users are still unable to successfully install the certificate governed by the new template. You have checked the structure of the CA hierarchy and are comfortable that no intentional attacks have taken place. What first step can you take to ensure the proper distribution of the certificate?

- A. Launch the **Certificate Templates** snap-in, right-click the **Certificate Templates** container, and select **New | Certificate Template to Issue**. Select the new certificate template.
 - B. Launch the **Certificate Templates** snap-in and highlight the **Certificate Templates** container. In the right pane of the console, right-click the new certificate template, and choose **Properties**. From the **Publish** tab, select the **Publish to Directory** option.
 - C. From any PKI client's browser, point to **http://servername/certsrv**, where *servername* is the name of the CA that contains the new certificate template. Select the **Issue a Certificate Template** link.
 - D. Using an account with appropriate permissions, copy the new certificate template to the root CA's certificate store. From the root CA, enable the template by using the **Certificate Templates** snap-in.
- A.** Answer A is the only possible answer. After a template is created by duplicating another template and editing the copy, it will not become available to clients until the template is enabled. Answer A gives the proper way to enable a certificate template.
- B.** Answer B is incorrect because a **Publish** tab does not exist as described. The correct way to enable a certificate is by selecting **New | Certificate Template to Issue** as in Answer A. Answer C is incorrect because you cannot manage templates using the Web enrollment service. Answer D is incorrect because any CA may issue certificates. In a three-tier hierarchy, the root CA's only responsibility would be to issue certificates to subordinate CAs.

Implementing Smart Card Authentication in the PKI

13. You have been designated as the enrollment agent for the entire Pants, Inc. organization during the smart card deployment that has just been completed. Your supervisor has now assigned you the project of updating the company's VPN solution by configuring the current RRAS server to accept smart card remote access. However, when you log on to the server and attempt to configure it, you are unsuccessful. What is the most likely reason for the failure?
 - A. The Extensible Authentication Protocol (EAP) has not been installed.
 - B. You are not a member of the Administrators group.

- C. The Routing and Remote Access Service does not have the required application certificate.
- D. A smart card reader has not been installed on the server.
- B.** For security reasons, you must be an administrator to configure the server's RRAS. Just because you are an enrollment agent does *not* mean that you are an administrator also. Enrollment agents are simply users who have been granted the appropriate permissions to configure smart cards. Therefore, Answer **B** is correct.
- Answer **A** is incorrect because EAP is installed as part of the server's configuration process – it is not required to be installed beforehand. Answer **C** is incorrect because, although it is true that a machine certificate is required on the server, an application certificate is not. Answer **D** is incorrect because a smart card reader is required only on the remote client machines. The server does not need a reader to receive authentication requests over the VPN.
14. Your company uses smart card authentication for its local network. You are an administrator and have been directed to install a new domain controller in the main office. You install Windows Server 2003 on the new hardware and begin the *dcpromo* process. When the install process asks you for authentication, what will you need to supply to finish the promotion?
- A. Username and password
- B. Smart card and PIN
- C. Username and PIN
- D. Smart card and password
- A.** To promote a server to a domain controller, you must provide Kerberos or NTLM authentication information *prior* to receiving a certificate from the issuing CA; therefore, Answer **A** is correct.
- Answer **B** is incorrect because a smart card and PIN are required after the server has been promoted and is operational. Answers **C** and **D** are incorrect because a username requires a password and smart cards use PINs as the second part of the two-part authentication process.
15. You are the administrator of a small network, and you have recently assigned yourself as an enrollment agent for your firm's new smart card system by making sure that you have Read and Enroll permissions on the Smart Card Logon template's **Security** tab. However, when you begin testing the implementation, you discover that you are unable to fully complete a request for a certificate on behalf of another user. You are using Internet Explorer on the enrollment station computer. Which of the following, if true, could be reasons for the failure? Choose all that apply.
- A. The smart card manufacturer's CSP has not been installed on the enrollment station.

- B. IIS has not been installed on the enrollment station.
 - C. The Write permission has not been assigned to your account.
 - D. Neither the Smart Card Logon nor the Smart Card User templates have been enabled on the CA.
 - E. You logged on to the enrollment station using your administrator account.
- A and D.** Answer A is correct because many smart card manufacturers use a proprietary CSP, or use one that has not been pre-installed on Windows Server 2003. A smart card cannot be enrolled until an appropriate CSP has been installed on the enrollment station. Answer D is correct because these are the collective templates that a CA uses to issue certificates that have been requested by the enrollment agent. At least one of the templates must be enabled, or the CA will be unable to issue the necessary certificate.
- Answer **B** is incorrect because the enrollment station is only the *client* requesting the certificate, and therefore only needs to employ a browser. The *CA* responsible for issuing the certificate must have IIS installed and Web enrollment support enabled. Answer **C** is incorrect, because only the Read and Enroll permissions are required for an enrollment agent to request certificates for another user. Also, you are an administrator and will already have the needed permission levels assigned to you. Answer **E** is incorrect for this same reason – the administrator is fully capable of behaving as an enrollment agent.

- : (colon), 216
- ! (exclamation mark), 107
- . (period), 345
- ? (question mark), 107
- _ (underscore character), 360–361, 396
- 2G (second-generation), 804
- 32-bit CPU, 563
- 3DES. *See* Triple Data Encryption Standard (3DES)
- 4GB tuning (4GT), 562
- 64-bit CPU, 563
- 6bone, 193
- 6to4 tunneling, 192
- 802.11 standards
 - authentication methods in, 806–807
 - product support of, 532
 - security of, 489
 - types of, 501
 - for wireless encryption, 504–505
 - wireless network types and, 803–804
 - wireless security and, 801–803
- 802.11a standard, 501
- 802.11b standard, 501
- 802.11g standard, 501
- 802.11i standard, 505, 816
- 802.11x standard, 804
- 802.1x standard, 489, 504–505, 802–803, 804

A

A records

- for delegating authority, 347
- of resource record, 342
- update with DNS/DHCP interaction, 387–389

ABR (area border router), 232

Acceptable Use Policy (AUP), 17–18, 46

access, 128

See also remote access strategy

access control, 58

access control entry (ACE), 784, 786

Access Control List (ACL)

- AD security guidelines for, 786

- DAACL/SACL in, 783

- WINS security and, 450–451

access point (AP), 807

access servers, 318

access token, 800

accidental threats, 91–92

account lockout duration setting, 826

Account Lockout Policy

- Group Policy to enforce, 785

- settings, 826

- for user account security, 797

account lockout threshold setting, 826

Account Policies, 94

account security

- built-in accounts, 796

- computer accounts, 797–798

- security principals, 795

- user account, 796–797, 798–800

- user authentication, 800

accounting with IAS, 309

accounts, disabling, 117–118

ACE (access control entry), 784, 786

ACL. *See* Access Control List (ACL)

Active Directory (AD)

- Active Directory–integrated zones, 375–377

- based IPsec policies, 747–749

- CAs and, 882–883

- client configuration for SUS updates, 844–845

- configuration planning and, 4

- DNS, hardware requirements for, 194

- DNS relationship to, 361–363

- domain controllers and, 58–60

- functional levels, 83–90

- Group Policy and, 746–747

- integrated with DNS, 64

- IPsec Policy Agent and, 724, 725

- network planning and, 13

- permissions, 787–788

- replication, 376, 377

- RSOP and, 766

- securing domain controllers, 121–122

- security features with, 81–83

- structure, 41–42

- supporting with BIND, 397–398

Active Directory (AD) security, 782–800

- account security, 795–800

- cross-domain relationships, 791–792

- cross-forest relationships, 793–795

- domain controllers, physically securing, 790

- guidelines for, 786

- permission types, 787–790

- permissions supported by, 783–784

- scenarios/solutions for, 785–786

- Schema Admins group, securing, 790
- static access control, 782–783
- summary of, 849
- Active Directory Domains and Trusts
 - checking domain function level in, 506–507
 - for external trust creation, 793–794
 - for forest trust creation, 794–795
 - function of, 82
 - raising domain/forest functionality, 90
 - raising domain functional level, 84–85, 508
 - raising forest functional level, 88
- Active Directory Installation Wizard (DCPROMO), 59, 363
- Active Directory-integrated zone
 - advantages of, 375–377
 - for DNS server, 373–374
 - footprinting and, 405
 - in high-level DNS security, 409–410
 - summary of, 463
 - troubleshooting, 455–456
 - updates, 348
 - zone replication security with, 382
 - zone transfers with BIND, 395
- Active Directory-integrated zone replication scope
 - changing, 380, 382
 - creating partition, 381
 - options of, 379–380
- Active Directory Sites and Services, 82
- Active Directory Users and Computers
 - to access domain/OU settings, 110
 - enabling remote access in, 493–495
 - function of, 82
 - for user account settings, 799–800
- AD. *See* Active Directory (AD)
- ad hoc mode, 801
- adapter settings, 666
- adapters. *See* network adapters
- Add/Edit Port Rule dialog box, 695
- Add or Remove Programs
 - for Certificate Services installation, 72–75
 - for Web server configuration, 67–68
- Address Pool tab, 295
- Address Resolution Protocol (ARP), 162
- addressing component, 214
- adjacency, 230
- Adleman, Leonard, 864
- administrative access, 669
- administrative model, remote access, 492–495
- administrative password, 814
- Administrator account
 - described/disabling, 796
 - disabling for security, 118
 - renaming, 796–797
 - security of, 851
- administrators, CA, 896–897
- Administrators group, 126
- Advanced Encryption Standard (AES), 868
- affinity, 680
- aging, 391–392
- AH. *See* Authentication Header (AH)
- AH tunnel mode. *See* tunnel mode
- AirSnort, 813
- algorithms
 - DES/3DES, 761–762
 - Diffie-Hellman, 724, 864
 - hash, 716
 - IKE and, 723
 - IPSec encryption, 715
- all communications (mixed network)
 - described, 663
 - illustrated, 664
 - security of, 667–668
- alternate configuration, 166–167
- Always On power scheme, 662
- American Registry for Internet Numbers (ARIN), 290
- Analyzing Configuration window, 672–673
- ANDing, 174–175
- ANI (Automatic Number Identification), 317
- announcements, 228
- Anonymous group, 789
- Anonymous Users group, 851
- antivirus software
 - for server security, 117
 - turning off for SUS installation, 838–839
- AP (access point), 807
- APIPA (Automatic Private IP Addressing), 166–167, 491
- application certificates, 870
- application directory partition
 - for Active-Directory-integrated zones, 376
 - AD-integrated zone replication scope and, 380–382
 - to reduce replication traffic, 383
- Application layer, OSI model, 238
- Application layer, TCP/IP, 161
- Application log, 584

- application memory tuning, 562
- application security, NLB, 691
- application servers
 - adding to Windows Server 2003, 76–77
 - defined, 57
 - function of, 75
 - securing, 130
 - Web server configuration, 67–68
- application services, 23–24
- applications, 701
- Approval Log, SUS, 845
- area border router (ABR), 232
- ARIN (American Registry for Internet Numbers), 290
- ARP (Address Resolution Protocol), 162
- ASBR (autonomous system boundary router), 233
- ASR. *See* Automated System Recovery (ASR)
- assets
 - determining value of, 92–93
 - protecting with security requirements, 93–94
 - security cost *vs.* benefit, 114
- ATA interface, 564–565
- attribute sets, 787
- attributes, 58
- audio services, 26
- Audit Policy function, 785
- auditing
 - centralized with IAS, 309
 - files/folders, 820–821, 822, 852
 - Registry keys, 821–822
 - Security log settings for, 823
 - security, turning on, 818–820
 - summary of, 848
 - viewing results of, 822
- Auditor role, 897
- AUP (Acceptable Use Policy), 17–18, 46
- authentication
 - with 802.1x standard, 803
 - authorization *vs.*, 329
 - described, 863
 - with domain controller, 58
 - EAP authentication, 804–805
 - IAS servers for, 532
 - Internet Authentication Service, 308–318
 - Kerberos authentication, 81
 - for mail server security, 128–129
 - overview of, 715–716
 - with pre-shared keys, 763–764
 - protocols, 810–812
 - with Public Key Infrastructure, 70
 - smart card in PKI, 897–906
 - with SQL Server, 127, 128
 - for wireless networks, 806–810
 - See also* Public Key Infrastructure (PKI)
- Authentication Data field, 720, 721
- Authentication Header (AH)
 - defined, 712
 - function of, 258
 - overview of, 721–722
- authentication methods
 - of IAS server, 314–317
 - for remote access, 508–512
 - restricting access by, 524–525
- authenticator, 804–806, 807
- authoritative answer, 351
- authoritative response, 473
- authoritative server
 - in DNS name resolution process, 351–352
 - DNS server placement, 372
 - name servers, 373–374
 - zone transfer and, 347–348
- authorization
 - authentication *vs.*, 329
 - of IAS, 317
 - of remote access, 516–520
- auto-enrollment
 - of certificates, 497, 895–896
 - PKI, 868
 - user certificates and, 911
- Automated System Recovery (ASR)
 - alternatives to, 614–615
 - backups, 120–121, 657
 - described, 612–613
 - overview of, 626–627
 - processes, 613–614
 - recovery with, 628–629
 - restoring with, 615, 617–618
 - Wizard, 615–617
- Automatic Number Identification (ANI), 317
- automatic partner configuration, 429–430
- Automatic Private IP Addressing (APIPA), 166–167, 491
- Automatic Updates software
 - required for SUS, 838
 - settings, 115–117
 - for SUS client configuration, 843

Automatic Wireless Wizard Configuration window, 808
 autonomous system boundary router (ASBR), 233
 availability. *See* high availability
 availability, network, 15

B

b-node. *See* broadcast node (b-node)
 backbone router, 233
 backup
 ASR, 613–614
 of cluster servers, 657
 domain controllers created from, 83
 media for, 604–605
 network performance and, 43
 overview of, 593–594, 626
 restoring from, 605–607
 scheduling, 605, 607–612
 strategy, 628
 for SUS update procedure, 842
 tools for, 120–121, 602–604
 types of, 596–599
 what to backup, 600–602
 Windows Backup, 594–596
 of WINS database, 451–452
 backup domain controller (BDC), 61
 Backup Operator role, 897
 Backup Utility. *See* Windows Backup
 Backup Wizard, 603–604
 bandwidth, 198, 491
 bandwidth allocation protocol (BAP), 492
 banner ads, 340
 BAP (bandwidth allocation protocol), 492
 baseline security, 94–112
 default security settings, 109–112
 secure baseline installation parameters, 103–108
 security templates/tools, 94–103
 summary of, 138
 baselining, 573, 626
 batch file, 448
 BDC (backup domain controller), 61
 #BEGIN_ALTERNATE tag, 420
 Bellman-Ford algorithms, 225
 Berkeley Internet Name Domain (BIND)
 DNS server support of, 362, 363
 GSS-TSIG and dynamic updates, 391
 supporting AD with, 397–398

 troubleshooting host name resolution and, 456
 zone replication and, 377–378
 zone transfers with, 395–397
 binary numbers
 converting to decimals, 176–177
 host address determination and, 174–175
 BIND. *See* Berkeley Internet Name Domain (BIND)
 binding order, 569, 666
 block ciphers, 762
 Block Inheritance flag, 38
 Bluetooth, 803, 806
 /boot parameter, 136
 bottlenecks
 disk, 564–568
 memory, 561–562
 network component, 568–570
 overview of, 560–561, 627–628
 processor, 563–564
 bridges, 240–242
 broadcast, 418–419
 broadcast addresses, 213
 broadcast network, 231
 broadcast node (b-node)
 described, 419, 465
 WINS proxy agent and, 443
 buffers, 569
 bug fixes, 115–117
 bugs, 115
 built-in logon authentication, 784
 burst handling
 for DoS attack protection, 450
 for WINS, 446, 467
 bus architecture, 563

C

cabling
 in test labs, 31
 upgrade considerations, 10, 28–29
 cache pollution, 406
 cached IPSec policy, 748
 caching name servers, 374
 callback security, 513
 Calling Line Identification (CLI), 317
 Canonical name (CNAME), 343
 CAs. *See* certification authorities (CAs)
 case-sensitivity, 359
 CDPs (CRL Distribution Points), 886–887

- central processing unit (CPU), 80, 425, 445
 - See also* processors
- centralization, 13–14
- certificate revocation list, 911
 - See also* Delta Certificate Revocation Lists (Delta CRLs)
- Certificate Services
 - for CA creation, 71–75
 - installation of, 873–875
 - process of, 872–873
- certificate templates
 - enrollment/distribution of, 887–892
 - version 2 of Server 2003, 868
- certificate trust list (CTL), 883
- certificates
 - application, 870
 - CA security and, 129
 - data in, 71
 - digital, 868–870
 - machine, 870
 - placing CA on VLAN for, 503
 - requests, 892–895
 - revocation, 886–887
 - smart cards for remote access, 514
 - user, 497, 870, 896
 - uses for, 69, 70–71
- Certificates MMC snap-in, 497–499
- certificates, PKI
 - auto-enrollment deployment, 895–896
 - enrollment/distribution of, 887
 - requests, 892–895
 - role-based administration, 896–897
 - templates, 887–892
- certification authorities (CAs)
 - certificate revocation, 886–887
 - Certificate Services and, 71–75
 - certificates, 70–71
 - function of, 69
 - hierarchy, planning, 881–884, 911
 - needs analysis, 881
 - overview of, 862, 870–872, 907
 - placing on VLAN, 503
 - Public Key Infrastructure and, 69–70
 - securing, 129
 - security, 885
 - types of, 881–883
- certification authorities (CAs), PKI
 - configuring, 876–880
 - implementing, 875–876
- Certification Authority snap-in, 876, 907
- Certification Revocation List (CRL), 71
- chaining, 72
- Challenge Handshake Authentication Protocol (CHAP), 509–511
- change and configuration management framework, 830, 850
- change-only replication, WINS, 428
- CHAP (Challenge Handshake Authentication Protocol), 509–511
- child domain
 - delegating authority to, 347
 - DNS and AD, 361–362
 - records in stub zone, 365–366
- classful addressing, 173–175
- Classless Interdomain Routing (CIDR)
 - IP address ranges listed with, 213
 - overview of, 180–181
 - supported protocols, 202
- CLI (Calling Line Identification), 317
- client
 - access, 669
 - configuring to retrieve updates with SUS, 843–844
 - defining subtype on, 809–810
 - restricting remote access by configuration, 524
 - support of VPN protocols, 496
- client access only (public network), 663, 667–668
- client compliant encryption level, 131
- Client IPSec policy, 732–733
- client-server connection, 301
- Clone Principal tool, 31
- Cluster Administrator tool, 653–654, 676–677
- cluster configuration log file security, 669
- cluster groups, 642–643
- cluster IP address, 679
- Cluster IP Addresses window, 694
- cluster nodes
 - failure, recovery from, 657
 - hub-and-spoke replication model and, 436
 - for WINS performance, 445
- cluster service account, 668–669
- Cluster.exe, 654, 655–656
- clustering
 - described, 15
 - server fault tolerance with, 624
 - See also* Network Load Balancing (NLB); server clustering

- clusters
 - ASR backups on, 614
 - backup of, 602
 - data arrangement and, 566
- CMAK. *See* Connection Manager Administration Kit (CMAK)
- CN (common name), 73
- CNAME (Canonical name), 343
- colon (:), 216
- command-line utilities
 - backups with, 604
 - for maintaining/monitoring DNS servers, 416–417
 - for scheduling, 197
 - with Windows Server 2003, 82
- command prompt, 219
- common name (CN), 73
- compatws template, 95
- components
 - CA, installation of, 894
 - hot swappable, 625
 - IPSec Policy Agent, 724–725
 - network, 568–570
 - of PKI, 867–868
- computer accounts security, 797–798
- computer certificates, 497–499
- computer clock synchronization, 825–826
- conditional forwarding
 - design configuration for, 384–386
 - for disjointed namespace, 365
 - function of, 374–375
 - server configuration for, 370
- confidentiality, 863–864
- Configure Your Server Wizard
 - for application server configuration, 76–77
 - for domain controller installation, 59
 - server roles applied with, 54
 - steps of, 55–57
 - for Web server configuration, 67
 - Web server role not offered by, 139
- Connection Manager
 - CMAK, using, 319–324
 - defined, 513
 - in NAQC, 524
 - Quarantine control and, 514
 - security issues, 324–325
 - summary of, 326, 327
- Connection Manager Administration Kit (CMAK)
 - configuration options of, 328
 - custom actions, 323–324
 - custom help, 324
 - defined, 513–514
 - installing/running, 319–320
 - security issues, 324–325
 - service profiles, 323
 - using, 320–323
 - VPN support by, 324
- connections
 - controlling remote connections, 525–528
 - dial-in remote access, 488, 489–495
 - remote access, managing, 513–514
 - restricting remote access by, 521–523
 - VPN remote access, 488–489, 495–500
 - wireless remote access, 500–505
- connectivity devices, 236–245
- consistency checking, 448–449
- constrained delegation, 800
- contention, disk access, 562
- context strings, 234
- controller. *See* disk controller
- convergence
 - of NLB cluster, 680, 687
 - of RIP routers, 228
- convergence time
 - factors that affect, 427
 - replication models and, 435
- convert.exe, 120
- copy backup, 596
- cost, 16, 114
 - See also* Total Cost of Ownership (TCO)
- counters. *See* performance counters
- counting to infinity, 228–229
- CPU. *See* central processing unit (CPU)
- CRL (Certification Revocation List), 71
- CRL Distribution Points (CDPs), 886–887
- CRL updates, 129
- cross-domain relationships, 791–792
- cross-forest relationships, 793–795
- cross trust, CA, 871, 883
- cryptography
 - overview of, 866–867
 - PKI, 864–866
- cryptology, 863
- CTL (certificate trust list), 883
- custom actions, 323–324
- custom help file, 324
- custom security templates, 131–134

D

DACL (discretionary access control list), 783, 784
data

- backup, 600
- confidentiality, 717
- drive arrangement of, 566–568
- encryption level, 512–513
- integrity, 496, 716
- security of server cluster, 669
- storage/retrieval, 21–23
- transit security, 714

Data Encryption Standard (DES), 715, 761–762

Data Link layer, OSI model

- function of, 237
- illustrated, 239

Layer 2 switches operate at, 244

data modification attacks, 405

data points, Event Viewer, 587

data source name (DSN), 128

Data Sources (ODBC) applet, 128

data stream, 257–258

database compaction, 448

Database description packet, 230

database servers, 68, 127–128

Day-and-Time-Restrictions attribute, 523

DC (domain component), 73

DC security template, 95

DCOM (Distributed Component Object Model),
683–684

DCPROMO (Active Directory Installation Wizard),
59, 363

DCs. *See* domain controllers (DCs)

DDNS. *See* Dynamic DNS (DDNS)

debug logging, 414–415

decentralization, 13–14

default cluster group, 667

default gateway, 222

default host, 678–679

default route, 217

default security settings, 109–112

default security.inf file, 828

default settings, wireless network devices, 813–815

default static route, 250–251

deliberate threats, 91

Delta Certificate Revocation Lists (Delta CRLs)

- CDPs and, 886–887
- PKI, 868
- Server 2003 and, 887

demand-dial connection, 304–306

demand-dial interface

- adding, 261–262
- configuring, 304–306
- in NAT installation, 293

Demand-Dial Interface Wizard, 261–262

demand-dial routing, 260, 261–262

demilitarized zone (DMZ), 258

Denial of Service (DoS) attacks

- on DNS server, 406–407
- on WINS, 449–450
- on WINS server, 126
- as wireless security threat, 813

Department of Defense (DOD) networking model,
452–453

deployment

- of IPSec, 711, 726–728
- testing and, 29–30

DES (Data Encryption Standard), 715, 761–762

design, network, 36–38, 39

destination address, 212, 217

DFS (Distributed File Service), 22, 63

DHCP. *See* Dynamic Host Configuration Protocol
(DHCP)

DHCPACK, 182

DHCPOFFER, 182

DHCPREQUEST, 182

dial-in access design, 489–495

- incoming port needs, 491–492
- IP addresses, allocating, 490–491
- list of, 489
- remote access by policy, 494–495
- remote access by user, 493–494
- summary of, 530

dial-in connection

- advantages of modems, 488
- callback security for, 513
- controlling IP address, 528
- restricting access by connection type, 522–523
- summary of, 529

Dialed Number Identification Service (DNIS), 317

dictionary attack, 807

differential backup, 598–599

Diffie-Hellman groups, 713, 762–763

Diffie-Hellman key-exchange algorithm, 724

Diffie, Whitfield, 864

digital certificates, 868–870

digital signatures

- CAs and, 907
 - RSA and, 865–866
 - security and, 910–911
- Dijkstra algorithm, 231
- direct memory access (DMA), 569
- directory, 58
- Directory Service log, 585
- disabled filtering mode, 679
- discretionary access control list (DACL), 783, 784
- disjointed DNS namespace
 - DNS configurations for, 361–362
 - features that support, 365–366
- disk controller
 - to drive ratio, 568
 - technologies, 564–565
- Disk Defragmenter, 566–568
- disk partitions, 120
- disk quotas
 - described, 22
 - e-mail and, 46
 - for mail servers, 129
- disk resource security, 669
- diskette drives, 616
- disks
 - controller/drive ratio, 568
 - controller technology of, 564–565
 - data access on, 568
 - data arrangement on, 566–568
 - drive life expectancy, 565–566
 - fault tolerance solutions for, 620–624
 - hot spare drives, 624
 - requirements for Windows OSs, 80
 - for server cluster, 661
 - shared cluster, 659–665
- distance-vector routing protocol, 225, 226–229
- distinguished name, 73
- Distributed Component Object Model (DCOM), 683–684
- Distributed File Service (DFS), 22, 63
- distribution groups, 86–87
- distribution of certificates, 887–897
- DLLS (dynamic link libraries), 233–235
- DMA (direct memory access), 569
- DMZ (demilitarized zone), 258
- DNIS (Dialed Number Identification Service), 317
- DNS. *See* Domain Name Service (DNS)
- DNS Console Monitoring tab, 413
- Dnscmd utility, 417
- DNSLint utility, 417, 454
- DnsUpdateProxy group, 390–391, 472
- DOD (Department of Defense) networking model, 452–453
- #DOM tag, 420
- domain
 - of Active Directory, 361–362
 - applying security template to, 109–110
 - authentication, 31
 - computer account security and, 797–798
 - defined, 58
 - functional level, upgrading, 32
 - zone *vs.*, 461–462, 472
- Domain Admins group, 519–520
- domain component (DC), 73
- domain controllers (DCs)
 - Active Directory and, 58–59
 - AD-integrated zone replication scope and, 379–380
 - AD-integrated zones and, 375–377
 - auditing, 819–820
 - created from backups, 83
 - defined, 57
 - defining subtype on, 808–809
 - DNS service and, 363
 - DnsUpdateProxy group and, 390
 - functional levels and, 83–90
 - functions of, 58
 - IPSec and, 712
 - operation master roles, 59–62
 - password requirements for, 119
 - physically securing, 790
 - root CAs and, 885
 - securing, 121–122
 - security templates and, 95–96, 97
 - tracks function level, 507
- domain functional levels
 - described, 83–87
 - raising, 90
 - remote access security and, 505–508
- domain local group scope, 792
- domain name
 - installing DNS service and, 353–354
 - supporting multiple namespaces, 363–369
- domain name master, 60
- Domain Name Service (DNS)
 - cache, 340–341, 455–456
 - client suffix search list, 403–404
 - databases, backup of, 602

- domain namespace, 344–345
- domains *vs.* zones, 345–348
- forwarding, 383–387
- function of, 341
- installing DNS service/configuring reverse lookup zones, 353–357
- monitoring DNS servers, 412–417
- name resolution process, 348–352
- namespace, designing, 357–369
- new features of, 472
- query, 455
- resource records and, 342–344
- reverse lookup zones, 352–353
- security issues, 404–412
- server deployment, 369–377
- settings, 32
- summary of, 461–464
- troubleshooting host name resolution, 453–457
- updates with DHCP, 387–392
- Windows Server 2003 DNS interoperability, 392–404
- zone replication, 377–383
- Domain Name Service (DNS) client
 - in DNS name resolution process, 348–352
 - troubleshooting host name resolution, 454–455
- Domain Name Service (DNS) namespace
 - security and, 410, 411
 - split DNS configuration, 398–399
 - summary of, 461–462
- Domain Name Service (DNS) namespace, designing
 - considerations for, 357–358
 - DNS and AD, 361–363
 - host naming conventions/limitations, 359–361
 - multiple namespaces, supporting, 363–369
 - parent domain name, choosing, 358–359
- Domain Name Service (DNS) records
 - aging/scavenging of, 391–392
 - security for, 389–391
 - updates with DHCP, 387–389
- Domain Name Service (DNS) security, 404–412
 - DoS attacks, 406–407
 - footprinting, 405
 - in general, 404
 - guidelines for, 410–412
 - redirection, 406
 - securing DNS deployment, 407–408
 - security levels, 408–410
- Domain Name Service (DNS) server
 - Active Directory and, 361–363
 - defined, 57
 - in DNS name resolution process, 348–352
 - domains *vs.* zones, 347–348
 - forwarding, 383–387
 - function of, 341
 - host naming conventions and, 360
 - installing DNS service/configuring reverse lookup zones, 353–357
 - monitoring, 412–417
 - multiple namespaces, 363–369
 - name resolution with, 64–65
 - securing, 125, 126
 - security issues, 404–412
 - split DNS configuration, 398–399
 - troubleshooting host name resolution, 453–457
 - zone replication planning and, 377–383
- Domain Name Service (DNS) server deployment, 369–377
 - number of servers, 369–371
 - placement, 372
 - roles, 373–377
 - server capacity, 371–372
- domain namespace
 - designing DNS namespace, 357–369
 - domains *vs.* zones, 345–348
 - structure of, 344–345
 - summary of, 461–462
- domain naming master, 382
- domain rename utility (rdom.exe), 86
- domain security ID (SID), 60
- domain tree, 361
- domain-wide master roles, 60–62
- domainlet, 669
- domains
 - cross-domain relationships, 791–792
 - cross-forest relationships, 793–795
 - smart cards and, 898
 - structure of, 31, 32
 - trust relationships between, 851
 - zones *vs.*, 345–348
- DoS attacks. *See* Denial of Service (DoS) attacks
- drainstop option, 687
- drive-by, 813
- drive-letter assignments, 661
- driver logging, IPsec, 756
- drivers, 725–726
- drives. *See* disks

- DSN (data source name), 128
 - duplex setting, 570
 - duplexing, 622
 - dynamic access control, 783
 - dynamic content, 127
 - Dynamic DNS (DDNS)
 - DHCP interaction with, 387–392
 - security for, 389–391
 - updates, 362–363
 - Dynamic Host Configuration Protocol (DHCP)
 - to assign IP address, 290
 - databases, backup of, 601
 - DNS server and, 463–464
 - DNS updates with, 387–392
 - security for, 389–391
 - troubleshooting, 182–183
 - Dynamic Host Configuration Protocol (DHCP) server
 - defined, 57
 - described, 154–155
 - hardware requirements, 194
 - for IP addressing, 490
 - role of, 63–64
 - securing, 125–126
 - WINS client configuration and, 440–441
 - dynamic IP address, 63–64
 - dynamic link libraries (DLLs), 233–235
 - dynamic mode commands, 749
 - dynamic mode policy, 749
 - dynamic records, 423
 - dynamic registration, 454–455
 - dynamic routing, 220–222, 245
 - dynamic updates
 - BIND support of, 397–398
 - DNS performance counters for, 416
 - redirection attack and, 406
 - troubleshooting, 457
- E**
- e-mail
 - mail servers for, 68–69
 - network performance and, 46
 - network planning and, 17–18
 - securing mail servers, 128–129
 - EAP. *See* Extensible Authentication Protocol (EAP)
 - EAP-TLS. *See* Extensible Authentication Protocol-Transport Layer Security (EAP-TLS)
 - eavesdropping, 813
 - EDNSO (Extension Mechanisms for DNS), 351, 412
 - EFS. *See* Encrypted File System (EFS)
 - electromagnetic (EM) field, 801
 - electronic communications. *See* e-mail
 - emergency repair disk (ERD), 612, 842
 - Encapsulating Security Payload (ESP)
 - defined, 713
 - function of, 258
 - overview of, 719–720
 - tunnel mode, 772
 - encapsulation, VPN protocols, 307
 - Encrypted File System (EFS)
 - described, 22
 - function of, 63
 - international use of, 910
 - for securing file server, 122–124
 - encryption
 - backups and, 595
 - certificates and, 70–71
 - data encryption level for remote access, 512–513
 - with EFS, 122–124
 - IPSec, 715, 761–762
 - protocols for VPN security, 307–308
 - with Public Key Infrastructure, 70
 - reversible encryption for password storage, 825
 - strength, controlling, 527
 - for terminal servers, 131
 - VPN protocols for, 496–497
 - with VPN remote access, 488–489
 - with WEP, 802
 - WEP weakness, 815
 - wireless, 504–505
 - End User License (EULA), 843
 - #END_ALTERNATE tag, 420
 - Enforce Inheritance flag, 38
 - enforce password history option, 824
 - Enforce User logon restrictions setting, 825
 - enrollment, certificate, 887–897
 - enrollment station, smart card, 899
 - enterprise CAs
 - defined, 72
 - overview of, 882, 908
 - environmental threats, 91, 114
 - ERD (emergency repair disk), 612, 842
 - error detection, 687
 - Error event type, 585
 - ESP. *See* Encapsulating Security Payload (ESP)
 - Ethernet

- switches *vs.* hubs on, 570
- Token Ring networks and, 569–570
- EULA (End User License), 843
- event header, Event Viewer, 587–588
- Event Log, 94
- event log files
 - with Audit Policy function, 785
 - overview of, 626
 - settings, 823
 - storage of, 590–592
- Event Logging tab, 415
- Event Viewer
 - IPSec troubleshooting with, 754–755
 - for monitoring NLB, 687
 - server monitoring with, 584–592
- exclamation mark (!), 107
- explicit permissions, 788
- extended rights, 787
- extensible architecture, IAS, 309
- Extensible Authentication Protocol (EAP)
 - authentication process, 804–805
 - described, 810
 - EAP-MD5, 511
 - EAP-MD5 CHAP, 316
 - EAP-MS-CHAPv2, 810
 - EAP over LAN (EAPOL), 805
 - EAP Over RADIUS, 512
 - EAP-RADIUS, 317
 - enabling authentication, 315–316
 - Remote Access Server, enabling on, 905
 - types for IAS authentication, 315
 - types supported by Windows Server 2003, 511–512
 - for wireless security, 802, 803
- Extensible Authentication Protocol–Transport Layer Security (EAP–TLS)
 - computer certificate to use, 503
 - defined, 511
 - described, 316–317, 810
 - support of, 512
- Extension Mechanisms for DNS (EDNSO), 351, 412
- external DNS infrastructure, 411
- external root certificate of authority, 871
- external trusts, 793–794
- extinction interval, 423, 447
- extinction timeout, 447

F

- failback, 643
- failover

- defined, 436, 641
- described, 643
- multiple interconnections and, 664
- ring, 651–652
- server cluster deployment and, 647–653
- Failure Audit event type, 585
- fast zone transfers
 - with BIND, 395–396
 - support for, 377–378
- fault tolerance
 - of DNS infrastructure, 462
 - of dynamic routing, 221, 222
 - of e-mail services, 17
 - forwarders and, 386
 - with hot spare drives, 624
 - in hub-and-spoke replication model, 436
 - Internet solutions, 619–620
 - of NetBIOS names, 417–418
 - network planning and, 15
 - network solutions, 619
 - number of DNS servers for, 369, 371
 - overview of, 627
 - planning for, 618–619, 629
 - RAID disk solutions, 620–624
 - in ring replication model, 435
 - in RIP networks, 229
 - of server cluster nodes, 642
 - server solutions, 624–625
 - with Windows Server 2003, 45
- feature test description, 34
- Federal Information Processing (FIPS) compliant encryption level, 131
- Feistel, Horst, 864
- Fibre Channel
 - Fibre Channel-based controller, 659
 - interface, 565
 - for server clustering, 643–644
- file encryption keys, 122
- File Replication Service log, 585
- file servers
 - defined, 57
 - function of, 62–63
 - securing, 121–124
- File Share resource, 642
- File System, 94
- File Transfer Protocol (FTP), 66
- files, auditing, 822
- filtering
 - filter lists/actions, 744–746
 - firewall packet, 762

- modes for port rules, 679–680
 - packet filtering, 268–269, 279
 - records with WINS, 424
 - FIPS (Federal Information Processing) compliant encryption level, 131
 - firewall filters, 499–500
 - firewalls
 - for DoS attack protection, 407
 - IPSec/IKE traffic and, 723
 - NAT/IPSec traffic on, 711, 772
 - packet filtering with, 268–269, 762
 - for server cluster, 667
 - for Web server security, 127
 - See also* Internet Connection Firewall (ICF)
 - five nines, 618–619
 - flat namespace, 418
 - Flexible Single Master of Operations (FSMO) roles
 - described, 59–60
 - importance of, 62
 - infrastructure master, 61
 - PDC emulator, 61
 - relative ID master, 60
 - floppy disk drive, 616
 - folders, auditing, 820–821, 822, 852
 - footprinting, 405
 - /force parameter, 136
 - foreign security principal, 793
 - forest
 - cross-domain relationships, 791–792
 - cross-forest relationships, 793–795
 - described, 361
 - functional level, upgrading, 32
 - functional levels, 87–90
 - operations master roles and, 59–62
 - root domain, 86, 361
 - smart card certificates and, 898
 - trust relationships between domains, 851
 - forest trusts
 - in AD security scenarios/solutions, 785
 - creating, 794–795
 - cross-forest relationships, 793
 - forest functional level and, 89
 - forest-wide operations master roles, 60
 - forward lookup record, 400–401
 - forward lookup zones
 - configuring, 353–354, 355–356
 - defined, 352
 - update with DNS/DHCP interaction, 387–389
 - forward-only server, 374, 384
 - forwarding. *See* conditional forwarding
 - forwarding address, 218
 - forwarding servers
 - described, 374–375
 - planning for, 383–387
 - FQDN. *See* fully qualified domain name (FQDN)
 - fragmentation, 566–568
 - freeloading, 813
 - front-end/back-end architecture, 681–682
 - FSMO roles. *See* Flexible Single Master of Operations (FSMO) roles
 - FTP (File Transfer Protocol), 66
 - full backup, 596–597, 599
 - Full Control permissions, 786
 - fully qualified domain name (FQDN)
 - changing, 353–354
 - period in, 345
 - troubleshooting host name resolution and, 454
 - functional levels
 - domain functional levels, 83–87, 792
 - forest functional levels, 87–90
 - remote access security and, 505–508
- ## G
- gateways
 - default gateway, 222
 - IPSec/IKE traffic and, 723
 - multiple gateways, 223–225
 - Gemplus, 899
 - General tab, RIP Properties dialog box, 253–254
 - Generic Script resource, 642, 662
 - Generic Security Service TSIG (GSS-TSIG), 391
 - Generic Service resource, 642
 - geography, network, 11
 - Global Catalog
 - defined, 60
 - forest trusts and, 89
 - infrastructure master and, 61
 - global group scope, 792
 - glue records, 365–366
 - GPMC (Group Policy Management Console), 4, 35–36
 - GPO. *See* Group Policy Object (GPO)
 - GPRresult command-line utility, 42
 - gpupdate command, 135, 136
 - graphics services, 26
 - Gray, Jim, 618

- green check mark, 107
 - group policies
 - applying, 109–112, 746–747
 - settings, 134–136
 - Group Policy
 - Account Lockout Policy settings with, 797
 - in AD security, 785–786
 - assigning/applying policies in, 746–747
 - auditing in, 819–822
 - configuring clients for SUS updates with, 844–845
 - defining subtype with, 808–809
 - modeling report, 4–9
 - RSoP and, 767
 - security templates and, 827
 - setting security policies with, 823–824
 - Group Policy Management Console (GPMC), 4, 35–36
 - Group Policy Object Editor
 - to apply security templates, 109–112
 - for custom security templates, 140
 - described, 99
 - Group Policy Object (GPO)
 - applying security templates with, 109–112
 - deploying security templates with, 134, 135
 - GPMC and, 35–36
 - object-based access control for, 818
 - settings, network testing and, 31
 - Group Policy Tool, 752
 - group scope, 792
 - groups
 - AD permissions on, 788
 - authorizing remote access by, 518–520
 - limiting membership to, 790
 - nesting, 86–87, 791
 - NetBIOS names, 418
 - restricting remote access by, 521
 - security groups, nesting, 86–87
 - growth, network, 28–29
 - GSS-TSIG (Generic Security Service TSIG), 391
 - Guest account, 118, 796
 - guest authorization, 317
- ## H
- h-node (hybrid node), 419, 465
 - hackers
 - DNS security against, 404
 - security updates and, 831
 - wireless network security and, 852–853
 - hard drive. *See* disks
 - hardware
 - acceleration, 760–761
 - bottlenecks, 561–570
 - network planning and, 29
 - network testing and, 30
 - requirements, analyzing, 193–194
 - requirements for SUS, 838
 - requirements for Windows OSs, 80
 - routing problems and, 274
 - for server clustering, 658–662
 - server fault tolerance and, 624–625
 - smart card hardware, 514
 - upgrades, 29, 43
 - WINS performance and, 445
 - Hardware Compatibility List (HCL), 658
 - hardware router, 289
 - hash functions, 716
 - HBAs (host bus adapters), 659
 - HCL (Hardware Compatibility List), 658
 - Health Insurance Portability and Accountability Act (HIPAA), 26
 - heartbeat
 - defined, 641
 - multiple interconnections and, 664
 - in NLB cluster, 680
 - in NLB process, 681
 - security and, 668
 - Heisenberg Principal, 573
 - Hellman, Martin, 864
 - hello packet, 230
 - HelpAssistant account, 796
 - helper files, 233–234
 - Helper service, IPv6, 192
 - HFNetChk tool, 831
 - hierarchy, CA
 - overview of, 871–872
 - planning, 883–884
 - high availability
 - Automated System Recovery, 612–618
 - backup/recovery strategy, 593–594
 - backups, 602–612
 - bottlenecks and, 560–570
 - described, 560
 - fault tolerance, planning for, 618–625
 - server monitoring with Event Viewer, 584–592
 - server monitoring with service logs, 593

- server monitoring with System Monitor, 570–580
 - System Monitor console, creating, 580–584
 - Windows Backup, 594–602
 - See also* Network Load Balancing (NLB); server clustering
 - high encryption level, 131
 - high-level DNS security, 409–410
 - HIPAA (Health Insurance Portability and Accountability Act), 26
 - hiscdc template, 96
 - hiscws template, 96
 - HomeRF, 806
 - hop count, 218
 - hops, 430
 - host bus adapters (HBAs), 659
 - host list feature, 691
 - host name resolution, 337–417
 - DNS basics, 341, 344–353
 - DNS/DHCP interaction, 387–392
 - DNS namespace, designing, 357–369
 - DNS security issues, 404–412
 - DNS server deployment, 369–377
 - forwarding, planning for, 383–387
 - host names, 338–339
 - hosts file, 339–341
 - installing DNS service/configuring reverse lookup zones, 353–357
 - key points about, 469–470
 - monitoring DNS servers, 412–417
 - necessity of, 337
 - NetBIOS over TCP/IP, 338
 - resource records, 342–344
 - summary of, 461–464
 - troubleshooting, 453–457, 468
 - Windows Server 2003 DNS interoperability, 392–404
 - zone replication, planning, 377–383
 - host naming
 - basics of, 338–339
 - conventions/limitations, 359–361
 - host names, 338–339
 - hosts file, 339–341
 - kinds of, 337
 - NetBIOS over TCP/IP, 338
 - resolution with hosts file, 339–340
 - Host Parameters window, 696
 - host route, 217
 - hosts
 - convergence/heartbeats, 680
 - IP address configuration for, 689–690
 - in NLB cluster, number of, 702
 - NLB error detection, 687
 - NLB host, 678–679
 - in NLB process, 681
 - NLB traffic distribution and, 679–680
 - security, 691
 - on TCP/IP networks, 171–172
 - hosts file, 339–341
 - hot spare drives, 624
 - hot-spare node, 649
 - hot-standby server/N+1 deployment option, 649–651
 - hot swapping, 15
 - hotfixes
 - with Software Update Services, 837–847
 - with Windows Update Web site, 115–117
 - “How New Delegation of Authentication Options Improve Security in Windows Server 2003” (Shinder), 800
 - HTTP (Hypertext Transfer Protocol), 66
 - hub-and-spoke replication model
 - described, 428, 435–436
 - summary of, 466
 - hubs
 - function of, types of, 239–240
 - UPSs for, 625
 - human intervention, 91
 - hybrid node (h-node), 419, 465
 - hybrid replication model, 437
 - Hypertext Transfer Protocol (HTTP), 66
 - hyperthreading, 563
- I**
- I/O ports, 259
 - IAB (Internet Architecture Board), 160
 - IAS. *See* Internet Access Server (IAS); Internet Authentication Service (IAS)
 - IAS management console, 312–313
 - IAS software development kit (SDK), 313
 - ICANN (Internet Corporation for Assigned Names and Numbers), 357, 358
 - ICF (Internet Connection Firewall), 162
 - ICMP (Internet Control Message Protocol), 162, 215
 - ICMP (Internet Control Message Protocol) Router Discovery, 159
 - ICMP tab, 295
 - ICS. *See* Internet Connection Sharing (ICS)
 - ICV (integrity check value), 720

- idle timeout, 525, 526–527
- IE (Internet Explorer), 96
- iesacIs template, 96
- IETF (Internet Engineering Task Force), 710, 770
- IGMP (Internet Group Messaging Protocol), 429, 430
- IGMPv3, 165
- IGPs (interior gateway protocols), 225
- IIS. *See* Internet Information Services (IIS) 6.0
- IKE. *See* Internet Key Exchange (IKE)
- in-addr.arpa, 352
- in-band DoS attack, 406
- in-house applications, 130
- #INCLUDE tag, 420
- Incoming Forest Trust Builders group, 795
- incoming ports, 491–492
- incremental backup, 597–598, 599
- incremental zone transfer (IXFR), 347–348, 363
- indexing service, 22–23
- InetOrgPerson object class, 86, 87
- Information event type, 585
- information flow factors, 11–12
- Information Technology (IT) management structure, 14
- Infrared Data Association (IrDA), 806
- infrastructure. *See* Public Key Infrastructure (PKI)
- infrastructure master, 61
- infrastructure mode, 801
- infrastructure planning
 - group policy modeling report, 4–9
 - network design, 9–11
 - overview of, 2–3, 40, 43–45
 - strategies, 3
 - tools for, 3–4
- initialization vector (IV), 815
- input filters, 268–269
- installation, network, 10
- integrity check value (ICV), 720
- integrity, data, 864
- intelligent hub, 240
- interconnect
 - adapter settings, 666
 - configuring interconnect networks, 663
 - interface, configuring, 664–665
 - multiple interconnections, 664
 - server cluster security, 668
 - TCP/IP settings, 666–667
- interface
 - for default route, choosing, 250–251
 - demand-dial interface, 261–262, 293, 304–306
 - field in route entry, 218
 - network, 260, 291, 292
- interface metric. *See* metric
- interior gateway protocols (IGPs), 225
- internal cluster communications only (private network), 662–663, 668
- internal DNS infrastructure, 411–412
- internal DNS root zone, 367–368
- internal domain name, 363–365
- internal domain namespaces, 368–369
- internal network, 358–359
- internal router, 232
- International Organization for Standardization (ISO), 160
- Internet
 - access, DNS security and, 404
 - connecting LAN to, 289–300
 - fault tolerance solutions, 619–620
 - TCP/IP and, 152–153
- Internet Access Server (IAS)
 - for authentication, 532
 - for wireless connections, 501, 503–504
- Internet Architecture Board (IAB), 160
- Internet Authentication Service (IAS), 308–318
 - access server support, 318
 - advantages of, 308–309
 - authentication methods, 314–317
 - authorization methods, 317
 - configuring with wireless networking, 811–812
 - management of, 309–313
 - outsourced dialing, 318
 - questions about, 328, 329
 - summary of, 327
 - for wireless authentication, 802
- Internet-based VPNs
 - communications in, 301–302
 - configuring, 302–303
 - reason to use, 301
- Internet Connection Firewall (ICF), 162
- Internet Connection Sharing (ICS)
 - activating, 297–298
 - adding custom service, 299–300
 - configuring, 298–299
 - limitations of, 297
- Internet connectivity strategy, 288–325
 - connecting LAN to the Internet, 289–300
 - overview of, 288

- using Connection Manager, 318–325
 - using Internet Authentication Service, 308–318
 - virtual private networks, implementing, 300–308
- Internet Control Message Protocol (ICMP), 162, 215
- Internet Control Message Protocol (ICMP) Router Discovery, 159
- Internet Corporation for Assigned Names and Numbers (ICANN), 357, 358
- Internet Engineering Task Force (IETF), 710, 770
- Internet Explorer (IE), 96
- Internet Group Messaging Protocol (IGMP), 429, 430
- Internet Information Services (IIS) 6.0
 - on application server, 130
 - application server and, 75
 - installing, 67–68
 - setting up, 139
 - Web server protocols, 66
 - Web server security, 126–127
- Internet Key Exchange (IKE)
 - audit disabling, 755
 - defined, 713
 - detailed tracing, 757–758
 - dynamic mode commands and, 749
 - IPSec and, 721–722
- Internet layer, TCP/IP, 162
- Internet name resolution, 367
- Internet Protocol Next Generation (IPng), 215
- Internet Protocol Security (IPSec)
 - components, 724–725
 - for data integrity/sender authentication, 496–497
 - deployment, 726–728, 770
 - diagnostics with netsh, 750
 - driver, 725–726
 - for encryption during transmission, 124
 - filtering, 710–711, 728
 - IP Security Policy Management MMC Snap-in, 728–731
 - IPv6 and, 726
 - with L2TP, 532
 - managing, 728, 771
 - modes, 717–718
 - monitoring, 749–751
 - netsh command-line utility, 731–732
 - overview of, 710–712
 - performance and, 569
 - policies, AD based, 747–749
 - policies, assigning/applying in Group Policy, 746–747
 - policies, custom, 734–746
 - policies, default, 732–734
 - policy precedence, 752
 - policy security levels, 727–728
 - process, 713–717
 - protocols, 718–724
 - RSoP for planning, 765–768, 771
 - securing IP packets with, 257–258
 - security considerations, 761–764, 771
 - service, 739
 - summary of, 769–770
 - terminology/concepts, 712–713
 - test lab, 746
 - troubleshooting, 751–761
 - for VPN security, 307–308
 - when not to use, 712
 - for wireless security, 816
 - for zone replication security, 382
- internet research, 23
- Internet Security Association and Key Management Protocol (ISAKMP)
 - defined, 713
 - IPSec and, 258, 721–722
- Internet Service Provider (ISP), 290
- Internet Software Consortium (ISC), 393
- Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX), 152
- InterNic.net Web site, 359
- interrupts, 563–564
- Intervals tab, 447
- intranet
 - domain namespace choices for, 363–365
 - Web server for, 65–66
- IP address
 - allocating for remote access, 490–491
 - configuring for NLB Manager, 689–690
 - DNS installation and, 354
 - DNS servers and, 64–65
 - formats/types of, 213
 - host names and, 339
 - issued by DHCP server, 63–64
 - multihomed WINS server and, 439–440
 - multiple gateways for, 223
 - name resolution with hosts file, 339–340
 - NAT and, 214–215
 - NetBIOS over TCP/IP and, 338
 - for new server cluster, 673–674
 - NLB traffic distribution and, 679–680
 - for PPP connections, 528

- routed connections and, 289–290
 - routing tables and, 217–218
 - settings for interconnects, 666–667
 - as software address, 277
 - troubleshooting NetBIOS name resolution and, 459
 - for Windows Server 2003 as router, 245
 - WINS name registration, 422
 - WINS server and, 65
 - See also* Transmission Control Protocol/Internet Protocol (TCP/IP) infrastructure
 - IP Address Assignment window, 265
 - IP Address resource, 642
 - IP addressing
 - public class, 203
 - requirements, analyzing, 171–172
 - for routed connections, 290
 - strategy, 201
 - troubleshooting, 181–183
 - IP packet filters, 528
 - IP properties, 294
 - IP routing, 270–276
 - See also* routing strategy
 - IP Security Monitor MMC snap-in, 750, 753–754
 - IP Security Policy Management MMC Snap-in described, 711
 - IPSec, viewing with, 752
 - overview of, 714
 - using, 728–731
 - IP Security Policy Wizard, 735–744
 - IP version 4 (IPv4), 215–216
 - IP version 6 (IPv6)
 - 6bone, 193
 - 6to4 tunneling, 192
 - described, 215–216
 - Helper service, 192
 - installation of, 184–189
 - IPsec6.exe, 190–191
 - Netsh commands with, 189–190
 - overview of, 165–166
 - PING/Tracert parameters, 191–192
 - Teredo, 193
 - transitioning to, 183–184
 - utilities, 184
 - ipconfig, 160
 - ipconfig /all, 454, 457
 - ipconfig /registerdns, 454, 455
 - IPng (Internet Protocol Next Generation), 215
 - IPSec. *See* Internet Protocol Security (IPSec)
 - IPSec Policy Agent, 724–725
 - IPsec6.exe, 190–191
 - ipseccmd.exe, 760
 - ipsecmn command, 750, 753–754
 - ipsecpol.exe, 731
 - IPv4 (IP version 4), 215–216
 - IPv6. *See* IP version 6 (IPv6)
 - IPX/SPX (Internetwork Packet Exchange/Sequenced Packet Exchange), 152
 - IrDA (Infrared Data Association), 806
 - ISAKMP. *See* Internet Security Association and Key Management Protocol (ISAKMP)
 - ISC (Internet Software Consortium), 393
 - ISO (International Organization for Standardization), 160
 - ISP (Internet Service Provider), 290
 - IT (Information Technology) management structure, 14
 - iterative query, 348–349, 473
 - IV (initialization vector), 815
 - IXFR (incremental zone transfer), 347–348, 363
- ## J
- Jetpack utility, 448
- ## K
- KDC (Key Distribution Center), 81, 898
 - Kerberos
 - alternate implementations of, 799
 - authentication in Windows 2000, 81
 - for cross-domain relationships, 791
 - for domain controllers, 122
 - IP Security Policy Wizard and, 739
 - policy settings, 825–826
 - Key Distribution Center (KDC), 81, 898
 - key recovery agent, 884
 - keys
 - archival/recovery, 868, 884
 - exchange settings, 743–744
 - key pairs, 864
 - lifetimes, 735
 - pre-shared, IPSec and, 763–764
 - in Public Key Infrastructure, 69–70
 - WEP keys, 807, 813, 815, 816
 - known-plaintext attack, 807

L

- L2TP. *See* Layer 2 Tunneling Protocol (L2TP)
- lame delegation
 - DNSLint utility for, 417
 - stub zones and, 366
 - troubleshooting host name resolution and, 456
- LAN. *See* Local Area Network (LAN)
- LAN Manager, 96, 97
- language preferences, 846
- LAS (Local Security Authority), 800
- Last Known Good boot, 614
- lastLogonTimestamp, 86
- law, 26–27
- Layer 2 switches, 244
- Layer 2 Tunneling Protocol (L2TP)
 - described, 307
 - filter types for, 269
 - IPSec encryption and, 715
 - IPSec with, 532
 - pre-shared keys and, 763–764
 - selecting for VPN remote access, 496–497
- Layer 3 switches, 244
- Layer 4 switches, 244
- layered network architecture
 - of Server 2003, 153
 - TCP/IP model numbering, 163
- lease time, 182
- least-cost path, 231
- Length field, AH header, 721
- licensing
 - SharePoint requirements, 21
 - Terminal Services, 24
 - for test/training lab, 30, 42
- lights-out configuration, 885
- Link-State Acknowledgement packet, 230
- link-state advertisements (LSAs), 225–226, 230–231
- link-state database (LSDB), 230–231
- link-state protocol, 225–226
- Link-State Request packet, 230
- Link-State Update packet, 230
- linked value replication, 89
- LMHOSTS file
 - function of/use of, 420–421
 - for NetBIOS name resolution, 419–420
 - summary of, 465
 - troubleshooting NetBIOS name resolution and, 457, 458
- load balancing
 - DFS and, 22
 - load-balanced configuration, 619
 - troubleshooting host name resolution and, 456
- load weight, 679
- Local Area Network (LAN), 289–300
 - routed connections, 289–290
 - routing option, 248–249
 - summary of, 326
 - translated connections, 290–300
- local computer IPSec policy, 748–749
- Local Policies, defined, 94
- Local Security Authority (LSA), 800
- Local Security Policy
 - of computer, 134–135
 - configuring clients for SUS updates with, 844
 - function of, 827
 - GPO settings override, 140
- Local Security Policy GPO, 823–824
- log files
 - of NLB Manager, 687–688
 - System Monitor, 197
 - System Overview, 574–576
 - WINS, protecting, 450–451
- logging
 - debug, 414–415
 - driver, 756
 - level, 269–270
 - mode, 765, 767–768
- Logging tab, 270
- logical print queue, 24–25
- /logoff parameter, 136
- logon. *See* authentication
- logon rights, 827
- low encryption level, 131
- low-level DNS security, 408–409
- Ipseccmd.exe utility, 731
- LSAs (link-state advertisements), 225–226, 230–231
- LSDB (link-state database), 230–231

M

- m-node. *See* mixed node (m-node)
- MAC. *See* Media Access Control (MAC)
- machine certificates, 870, 896
- mail servers
 - defined, 57
 - function of, 68–69

- securing, 128–129
- use of certificates, 70–71
- main mode SA, 714, 722–723
- majority node set (MNS) server cluster model, 646–647
- Manage Documents permission, 124, 125
- Manage Printers permission, 124, 125
- Manage Your Server tool
 - for application server configuration, 76–77
 - managing server roles with, 54–58
- management
 - of IPSec, 728–749
 - priorities in network design, 14–16
- management model, 12
- master drive, 564–565
- master key, 743–744
- maximum lifetime for service ticket setting, 825
- maximum lifetime for user ticket, 825
- maximum lifetime for user ticket renewal, 825
- maximum password age option, 824
- maximum session time, 525–527
- maximum tolerance for computer clock synchronization, 825–826
- MBSA. *See* Microsoft Baseline Security Analyzer (MBSA)
- mbsacl.exe, 832–833
- mean time between failures (MTBF), 565–566
- Media Access Control (MAC) address
 - MAC bridges, 240
 - NetBIOS over TCP/IP for, 338
 - restricting remote access by, 525
- Media Access Control (MAC) filtering, 816
- media, backup, 604–605
- Media Services, 26
- medium-level DNS security, 409
- member server, 375–377
- memory (RAM)
 - bottlenecks, 561–562
 - buffers, 569
 - planning DNS server capacity, 371
 - requirements for Windows OSs, 80
- message authentication code, 720
- Message Exchanger (MX), 342
- metric
 - automatic determination of, 166–167
 - defined, 279
 - determining for default gateway, 167–170
 - disabling automatic, 170
 - field in route entry, 218
 - multiple gateways and, 223
- Microsoft
 - authentication protocols, 862–863
 - DoS attack on DNS servers of, 407
 - push/pull partnership recommendation of, 433–434
 - server cluster hardware compatibility, 658
- Microsoft Baseline Security Analyzer (MBSA)
 - function of, 831
 - installing, 832–833
 - to scan for security problems, 852
 - setting up Windows XP client for wireless networking, 833–837
- Microsoft certificate services, 872–875
- Microsoft Challenge Handshake Authentication Protocol version 1 (MS-CHAP v1), 509–511
- Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), 314, 511
- Microsoft Outlook, 831
- Microsoft Point-to-Point Encryption (MPPE), 307
- Microsoft POP3 Service, 128–129
- Microsoft SharePoint. *See* SharePoint
- Microsoft Software Update Services (SUS). *See* Software Update Services (SUS)
- Microsoft SQL Server. *See* SQL Server
- Microsoft Windows operating systems. *See* specific Windows operating systems
- Microsoft Windows Update site
 - downloading updates from, 840–841
 - Manage Your Server tool and, 57
 - SUS service connects to, 838
 - synchronizing SUS server with, 839–840
 - updates from, 115
- migrate on setting
 - for redirection attack protection, 450
 - static WINS entries and, 438
- minimum password age option, 824
- minimum password length option, 825
- mirrored IPSec filters, 745
- mirroring. *See* RAID 1
- mirroring plus striping. *See* RAID 0+1
- mixed mode
 - authentication with, 127
 - to native mode, switching, 32
- mixed node (m-node)
 - described, 419, 465
 - troubleshooting NetBIOS name resolution and, 458
- MNS (majority node set) server cluster model, 646–647
- modem

- adding demand-dial interface, 261
- needs for dial-in access, 491–492
- outsourced dialing with, 318
- for remote access, 488
- restricting access by connection type, 522–523
- modes, IPSec
 - ESP/AH and, 719
 - transport, 718
 - tunnel, 717–718
- module operator, 866–867
- monitoring
 - IAS server, 313
 - IPSec, 749–751
 - NLB, 687–689
 - SUS server updates, 846–847
 - updates, 831
- monitoring DNS servers, 412–417
 - command-line tools for, 416–417
 - debug logging, 414–415
 - event logging, 415
 - in general, 412
 - with Performance console, 415–416
 - testing DNS server configuration, 413
- monitoring security
 - auditing, 818–823
 - object-based access control for, 818
 - summary of, 848, 849
 - Wireless Monitor, 817–818
- Monitoring tab, 413
- monitors, 625
- MPPE (Microsoft Point-to-Point Encryption), 307
- mrinfo tool, 272, 273
- MS-CHAP v1 (Microsoft Challenge Handshake Authentication Protocol version 1), 509–511
- MS-CHAP v2 (Microsoft Challenge Handshake Authentication Protocol version 2), 314, 511
- MTBF (mean time between failures), 565–566
- multi-master replication, 59
- multicast addresses, 213
- multicast routing, 272, 273
- multihomed computer, 290
- multihomed WINS servers, 439–440, 459
- multilink connections, 491–492
- multiple gateways, 223–225
- multiple host filtering mode, 680
- multiple namespaces, 363–369
- multiprocessing, 563
- multiprotocol environments, 153–156

- multithreading, 563
- mutual authentication, 511
- MX (Message Exchanger), 342

N

- N-node failover pairs deployment option, 648–649
- name checking, 396
- name mappings, 444
- name release update, 427
- name resolution
 - with DNS server, 64–65
 - feature of NAT, 214, 294
 - overview of, 336
 - of server cluster, 643
 - with WINS server, 65
- name resolution strategy, host name resolution, 337–417
 - DNS basics, 341, 344–353
 - DNS/DHCP interaction, 387–392
 - DNS namespace, designing, 357–369
 - DNS security issues, 404–412
 - DNS server deployment, planning, 369–377
 - forwarding, planning for, 383–387
 - host names, 338–339
 - hosts file, 339–341
 - installing DNS service/configuring reverse lookup zones, 353–357
 - monitoring DNS servers, 412–417
 - necessity of, 337
 - NetBIOS over TCP/IP, 338
 - resource records, 342–344
 - summary of, 461–464
 - Windows Server 2003 DNS interoperability, 392–404
 - zone replication, 377–383
- name resolution strategy, NetBIOS name resolution, 417–452
 - LMHOSTS file, 420–421
 - multihomed WINS servers, 439–440
 - need for, 417–418
 - NetBIOS names, 418
 - process of, 418–420
 - split WINS registrations, preventing, 444
 - static WINS entries, 438–439
 - summary of, 464–468
 - WINS basics, 421–423
 - WINS client configuration, 440–444
 - WINS database backup/restoration, 451–452

- WINS improvements, 424
- WINS performance issues, 444–449
- WINS replication, planning for, 427–437
- WINS security issues, 449–451
- WINS server deployment, planning, 424–426
- name resolution, troubleshooting, 452–460
 - in general, 452–453
 - host name resolution, 453–457
 - key points about, 471
 - NetBIOS name resolution, 457–460
 - summary of, 468–469
- Name Server (NS) records
 - delegating authority to child domain, 347
 - disjointed namespaces and, 365, 366
 - in DNS name resolution process, 350
 - internal DNS root zone deployment and, 367
 - of resource record, 342–343
 - in reverse lookup zone, 356
- namespaces
 - internal domain namespace guidelines, 368–369
 - split DNS configuration, 398–399
 - supporting multiple, 363–369
 - See also* Domain Name Service (DNS) namespace
- NAQC. *See* Network Access Quarantine Control (NAQC)
- NAS-Port-Type attribute, 521–522
- NAT. *See* Network Address Translation (NAT)
- NAT/Basic Firewall tab, 295
- NAT traversal, 723
- National Security Agency (NSA), 864
- native mode, 32
- natural disasters, 114
- NBMA (non-broadcast multiple access), 231
- nbstat command, 442–443, 457
- nbstat -RR command, 426, 442
- ND (Neighbor Discovery), 183–184
- NDIS (Network Driver Interface Specification), 148–149, 163
- NDS (Netware Services Directory), 153
- negative answer, 351–352
- Neighbor Discovery (ND), 183–184
- nested contexts, 236
- nesting
 - group nesting, 791
 - security/distribution groups, 86
- net start command, 448
- net stop command, 448
- NetBEUI, 338
- NetBIOS
 - described, 152
 - disabling, 198
 - name, defined, 418
 - names in Windows Server 2003, 139
 - node types, 419, 465
 - scopes, 402
 - security issues, 449–450
- NetBIOS name resolution, 417–452
 - host name resolution and, 337
 - key points about, 470–471
 - LMHOSTS file, 420–421
 - multihomed WINS servers, 439–440
 - need for, 417–418
 - NetBIOS names, 418
 - process of, 418–420
 - split WINS registrations, preventing, 444
 - static WINS entries, 438–439
 - summary of, 461, 464–469
 - support with WINS, 399–400
 - troubleshooting, 453, 457–460
 - WINS basics, 421–423
 - WINS client configuration, 440–444
 - WINS database backup/restoration, 451–452
 - WINS improvements, 424
 - WINS performance issues, 444–449
 - WINS replication, 427–437
 - WINS security issues, 449–451
 - with WINS server, 65
 - WINS server deployment, 424–426
- NetBIOS over TCP/IP (NetBT)
 - described, 338
 - troubleshooting name resolution and, 453
 - WINS client configuration and, 440–441
 - WINS server security and, 126
- NetBT. *See* NetBIOS over TCP/IP (NetBT)
- netdiag tool, 751
- netdom trust command, 793, 794
- netlogon service, 388
- netlogon.dns file, 363
- netmask, 217–218
- netsh command-line utility
 - commands, using, 233–236
 - controlling with, 731–732
 - IPSec driver logging and, 756
 - IPSec monitoring with, 749–750
 - overview of, 189–190
 - for troubleshooting routing, 273

- for WINS server configuration, 426
- netsh dynamic mode policy, 749–750
- netsh ipsec command, 750
- Netware Services Directory (NDS), 153
- network
 - cluster network configuration, 662–667
 - components, performance and, 568–570
 - destination in route entry, 217
 - fault tolerance solutions, 619
 - subnetting, 177–180
 - test, 30–33
 - types, 231–232
- Network Access Quarantine Control (NAQC)
 - described, 514
 - to restrict client access, 329
 - restricting access by client configuration, 524
- network adapters
 - multiple, 154
 - multiple, for NLB, 689, 690
 - for Windows Server 2003 as router, 245
- Network Address Translation (NAT)
 - components, 214–215
 - configuring NAT connection, 295–296
 - on firewalls, 772
 - Internet Connection Sharing, 297–300
 - managing, 294
 - NAT server, installing, 291–293
 - NAT server tasks, 296–297
 - questions/answers about, 328, 329
 - Teredo and, 193
- Network Driver Interface Specification (NDIS), 148–149, 163
- network ID route, 217
- network interface adapter, 701
- network interface card (NIC)
 - duplex setting of, 570
 - fault tolerance of, 619
 - performance and, 568, 569, 570
- network interface controllers, 658–659
- Network Interface layer, TCP/IP, 162–163
- network interfaces
 - minimizing number of, 260
 - on NAT server, 291, 292
- Network Interfaces node, 304–306
- network latency, 647
- Network layer, OSI model
 - function of, 237–238
 - Layer 3 switches operate at, 244
 - network that requires, 239
- Network Load Balancing (NLB), 678–698
 - best practices, 689–691
 - creating NLB cluster, 691–698
 - display command, 689
 - function of, 678
 - Internet and, 619
 - managing clusters, 682–687
 - monitoring, 687–689
 - overview of, 640
 - questions/answers about, 701–702
 - relationship to clustering, 681–682
 - summary of, 699, 700
 - terminology/concepts, 678–681
- network mask, 217
- Network Monitor
 - described, 195–196
 - for IPSec Protocol determination, 719
 - IPSec troubleshooting with, 759–760
 - for network planning, 4
 - System's Management Server, 759
 - for troubleshooting routing, 271
- Network Name resource, 642
- Network News Transfer Protocol (NNTP), 66, 299–300
- network priorities, 677
- Network Priority property, 665
- network services, 31
- network testing. *See* test environment
- network topology
 - DNS server deployment and, 369
 - replication models and, 434
 - simplifying for routing security, 259–262
 - WINS replication and, 427–428
- New Delegation Wizard, 347, 357
- New Server Cluster Wizard, 654, 663, 671–676
- Next Header field, 720, 721
- next-hop interface, 217
- next-hop IP address
 - defined, 217
 - in route entry, 218
 - See also* gateway
- NIC. *See* network interface card (NIC)
- NLB. *See* Network Load Balancing (NLB)
- NLB Manager
 - accessing, 683
 - creating NLB cluster with, 691–698
 - drainstop option, 687

- monitoring with, 687–688
 - NLB.exe utility and, 684
 - remote management with, 683–684
 - security with, 691
 - tasks with, 682–683
 - NLB query, 688–689
 - NLB.exe utility
 - for administrative tasks, 682
 - command-line parameters used with, 684–687
 - for status information, 688
 - NNTP (Network News Transfer Protocol), 66, 299–300
 - node-to-node communication, 664–665
 - nodes
 - cluster nodes, 436, 445, 657
 - server cluster, 641–642
 - non-broadcast multiple access (NBMA), 231
 - non-repudiation, 864
 - nonauthoritative response, 371, 473
 - nonclustered network (disabled), 663
 - nonrecursive servers, 375
 - nonroutable transport protocols, 150
 - normal backup, 596–597, 599
 - notation, CIDR, 180, 203
 - notify list, 378
 - NS records. *See* Name Server (NS) records
 - NSA (National Security Agency), 864
 - NSLookup
 - footprinting with, 405
 - for monitoring DNS servers, 417
 - reverse lookup zones used by, 352
 - troubleshooting host name resolution with, 454, 455
 - NTDS.dit file, 59
 - NTFS
 - for application server security, 130
 - partitions, 129
 - security with, 119–120
 - volumes, EFS encryption and, 123–124
 - volumes on file server, 122
 - NTFS permissions
 - described, 22
 - for server cluster security, 669
 - settings, 788–789
- O**
- Oakley key-determination protocol, 713, 724
 - object-based access control, 818
 - object classes, 86
 - object permissions, 787
 - objects
 - in directory, 58
 - Windows Server 2003 AD and, 82
 - offline files, 22
 - offload functions, hardware, 760–761
 - one-way initiation, 306
 - open authentication, 806
 - Open Shortest Path First (OSPF)
 - benefits over RIP, 229–230
 - configuring RRAS for, 255–257
 - as link-state protocol, 225–226
 - network types supported by, 231–232
 - process of, 230–231
 - router roles, 232–233
 - routing problems with, 275–276
 - Open Systems Interconnection (OSI) model
 - layers of, 237–239
 - overview of, 149
 - protocols and, 160–161
 - troubleshooting name resolution in, 452–453
 - operating system (OS)
 - backup, 120–121
 - choosing for server security, 79–81, 137
 - Connection Manager security and, 324–325
 - functional levels, 83–90
 - IPSec and, 710
 - security features, 81–83
 - support of VPN protocols, 496
 - operations master roles
 - described, 59–60
 - importance of, 62
 - infrastructure master, 61
 - PDC emulator, 61
 - relative ID master, 60
 - Option 044 WINS/NBNS Servers, 441
 - Option 046 WINS/NBT Node Type, 441
 - organizational needs
 - centralization *vs.* decentralization, 13–14
 - growth, planning for, 28–29
 - information flow factors, 11–12
 - for IPSec, 727
 - legal/regulatory considerations, 26–27
 - management model/organizational structure, 12
 - management priorities, 14–16
 - overview of, 11, 40–41
 - for remote access, 487

- TCO, calculating, 27–28
 - user priorities, 17–26
 - organizational policies, 93
 - organizational structure, 12
 - Organizational Units (OUs)
 - AD and, 41–42
 - applying security template to, 109–110
 - defined, 73
 - Full Control permissions and, 786
 - IPSec policies and, 747
 - network planning and, 13
 - security policies set at, 135
 - OSI model. *See* Open Systems Interconnection (OSI) model
 - OSPF. *See* Open Shortest Path First (OSPF)
 - OUs. *See* Organizational Units (OUs)
 - out-of-band DoS attack, 406–407
 - Outlook, Microsoft, 831
 - output filters, 268–269
 - outsourced dialing, 318
 - ownership, 784
- P**
- p-node (peer node), 419, 465
 - packet event logging, 755–756
 - packet filtering
 - choices for, 268–269
 - firewall, 762
 - methods of RRAS, 279
 - packet header structure, 230
 - Packet Signature and Encryption field, 721
 - Packet Signature with the AH Header field, 721
 - packets
 - described, 149, 569
 - IPSec, 710–711, 728
 - in NAT process, 296–297
 - Padding field, ESP trailer, 720
 - Padding Length field, ESP trailer, 720
 - page files, 561–562
 - PAP (Password Authentication Protocol), 509–511
 - parent domain
 - delegating authority to child domain, 347
 - DNS and AD, 361–362
 - name, choosing, 358–359
 - parity block, 623
 - partitions, 120, 602
 - See also* application directory partition
 - passive hubs, 240
 - password
 - adding to InetOrgPerson accounts, 86
 - administrative password for WAP, 814
 - in custom security template, 132, 134
 - with DSN, 128
 - for IAS server, 312
 - user account security, 798–799
 - Password Authentication Protocol (PAP), 509–511
 - password-based authentication methods, 509–511
 - Password Expiration problem, 835–836
 - password must meet complexity requirement, 825
 - password policies
 - applying to all clients, 852
 - Group Policy to enforce, 785
 - options of, 824–825
 - passwords
 - Connection Manager security and, 325
 - PDC emulator and, 61
 - strong passwords, 118–119
 - patches
 - importance of, 831
 - with Software Update Services, 837–847
 - from Windows Update Web site, 115–117
 - pathping command, 272
 - PDC (primary domain controller) emulator, 61
 - PEAP (Protected Extensible Authentication Protocol), 810–811
 - peer node (p-node), 419, 465
 - perfect forward secrecy (PFS), 743
 - performance
 - baseline for server cluster, 657
 - network planning and, 16
 - optimizing network, 198–199
 - of WINS, 444–449
 - Performance console, 415–416
 - performance counters
 - commonly referenced, 572–573
 - data, assessing, 576–578
 - log files, 574–578
 - System Monitor, 196
 - for WINS, 445
 - Performance Logs and Alerts function, 573–574, 578
 - Performance Monitor, 45
 - perimeter network, 257, 258
 - period (.), 345
 - Permcopy.exe, 830–831
 - permissions

- Active Directory, 786, 787–788
 - for database server security, 127–128
 - NTFS, 22, 669, 788–789
 - printer, 124–125
 - with rootsec template, 96
 - share, 789–790
 - supported by AD, 783–784
 - for terminal servers, 130
 - user rights *vs.*, 826
- persistent connections, 306, 431
- personal identification number (PIN), 897
- PFS (perfect forward secrecy), 743
- phone lines, 488
- phone number, 525
- Physical Disk resource, 642
- Physical layer, OSI model, 237, 239
- physical printer, 24
- physical security
 - of domain controllers, 790
 - of NLB, 691
 - of print servers, 124
 - of server cluster, 667
 - of servers, 113–114
- PIN (personal identification number), 897
- PING
 - ICMP and, 162
 - IPv6 parameters, 191–192
 - testing TCP/IP connections with, 271, 279
- PKCS (Public Key Cryptography Standard), 864
- PKI. *See* Public Key Infrastructure (PKI)
- plan
 - network, documenting, 36–38
 - test network, 29–30
- planning mode, 765, 768
- Point-to-Point network, 232
- Point-to-Point Protocol (PPP), 314, 488
- Point-to-Point Tunneling Protocol (PPTP)
 - described, 307
 - packet filters, 269
 - for VPN remote access, 496–497
- pointer record (PTR)
 - for DNS server, adding, 356
 - of resource record, 342
 - for reverse lookup zones, 352
 - update with DNS/DHCP interaction, 387–389
- policies, IPsec
 - AD based, 747–749
 - assigning/applying in Group Policy, 746–747
 - custom, 734–746
 - default, 732–734
 - managing, 772
- policy
 - assignment information, 752
 - enabling remote access by, 494–495
 - See also* password policies; remote access policies; security policies
- POP3 (Post Office Protocol), 68–69
- port rules
 - filtering modes, 679–680
 - in NLB cluster creation, 694–695, 698
- Port Rules window, 694
- ports
 - DNS ports and security, 412
 - of hubs, 239
 - port switching, 243
- positive answer, 351
- Possible Owners property, 648, 649–650
- Post Office Protocol (POP3), 68–69
- power-management features, 662
- power sources, redundant, 625
- PPP (Point-to-Point Protocol), 314, 488
- PPTP. *See* Point-to-Point Tunneling Protocol (PPTP)
- pre-shared keys, 716, 763–764
- #PRE tag, 420
- predefined templates, 95–97
- Preferred Owners property
 - failover ring order and, 651
 - setting, 649, 650
- Presentation layer, OSI model, 238
- primary domain controller (PDC) emulator, 61
- primary master server, 373
- primary server, 347
- Print permission, 124, 125
- print queue, 24–25
- print servers
 - defined, 57
 - securing, 122, 124–125
- print services, 24–25
- printer permissions, 124–125
- printer pool, 24
- printer servers, 62
- priority, 564
- privacy, 802
- private DNS namespace, 357
- private IP addresses
 - address blocks defined as, 214

- uses of, 174–175
 - private key
 - in Public Key Infrastructure, 69–70
 - of Server 2003, 865
 - private root zone, 408
 - privileges, 827
 - processor affinity, 564
 - processors
 - performance and, 563–564
 - requirements for Windows OSs, 80
 - project collaboration, 19–21
 - Properties dialog box, 295–296
 - property set, 788
 - Protected Extensible Authentication Protocol (PEAP), 810–811
 - protocol field, 218
 - protocols
 - authentication, 810–812
 - multiprotocol environments, 153–156
 - network testing and, 31
 - NLB support of, 689
 - nonroutable transport protocols, 150
 - requirements, identifying, 149–151
 - routable, 152
 - Server 2003 supported, 569
 - supported by Windows, 153
 - VPN protocols, 306–307, 496–497
 - Web server protocols, 66
 - See also* routing protocols; specific protocol names
 - protocols, IPSec
 - AH, 721–722
 - IPSec ESP, 719–720
 - ISAKMP/IKE, 722–724
 - overview of, 718
 - primary, 712
 - proxy servers
 - internal DNS root zone with, 367–368
 - IPSec/IKE traffic and, 723
 - redundancy with, 620
 - PTR. *See* pointer record (PTR)
 - public DNS namespace, 357
 - public key, 69–70
 - Public Key Cryptography Standard (PKCS), 864
 - Public Key Infrastructure (PKI)
 - CAs, implementing, 875–887, 908–909
 - CAs, overview of, 870–872
 - certificates, enrollment/distribution of, 887–897, 909–910
 - components of, 867–868
 - cryptography, 864–867
 - described, 69–70
 - digital certificates, 868–870
 - function of, 867
 - for L2TP, 497
 - Microsoft certificate services, 872–875
 - overview of, 863–864
 - purpose of, 907
 - Server 2003 certificate-based, 862–863, 908
 - Smart Card authentication, 897–906
 - Public Key Interoperability, 871
 - public keys, 865
 - pull-only replication, 433
 - pull replication partnership
 - convergence time and, 427
 - push replication *vs.*, 474
 - settings, 432–433
 - troubleshooting, 460
 - pull request, 430
 - push notification, 430–431, 434
 - push-only replication partnership, 432
 - push partnerships
 - manually starting push notification, 431–432
 - process of, 430
 - settings, 431
 - push/pull replication partnership
 - convergence time and, 427
 - described, 433–434
 - push replication partnership
 - convergence time and, 427
 - pull replication *vs.*, 474
 - troubleshooting, 460
- ## Q
- qualified subordination, 868
 - query
 - cache pollution and, 406
 - forwarding planning and, 383–386
 - IPSec-related, 767–768
 - iterative, 348–349, 473
 - recursive, 348–349, 413, 416, 473
 - question mark (?), 107
 - quorum drive, 661, 675
 - quorum resource, 644
- ## R
- RA (registration authority), 70
 - radio frequency (RF), 801
 - RADIUS

- access server, 318
- client, configuring WAP as, 503–504
- IAS of, 308–318
- remote access policies and, 515
- RRAS server and, 512
- server in EAP authentication, 804–805
- for wireless authentication, 501
- RAID, 566, 620
- RAID 0, 620
- RAID 0+1, 623–624
- RAID 1, 621–622
- RAID 5, 622–623
- RAID array, 445
- RAID controller, 659
- RAM (random access memory). *See* memory (RAM)
- random deployment option, 652–653
- ranges, private address, 214
- RC4 encryption algorithm, 802, 815
- RDC (Remote Desktop Connection), 753
- Read Group Membership permission, 784
- readers, smart card, 899
- recovery
 - key, 868
 - overview of, 593–594
 - root CAs and, 885
 - strategy, 628
- recovery agent, 884
- Recovery Console, 120
- recursion
 - disabling for DNS security, 404, 406
 - disabling for DoS attack protection, 407
 - by DNS server, 348
 - nonrecursive servers, 375
 - troubleshooting host name resolution and, 455
- recursive query
 - defined, 473
 - in DNS name resolution process, 348–349
 - of DNS server, monitoring, 416
 - test, 413
- red X, 107
- redirection attack
 - causes of, 406
 - prevention with static entries, 439
 - on WINS, 450
- redundancy
 - fault tolerance and, 618
 - proxy server, 620
 - with server hardware, 624–625
- referral answer, 351
- referral zone, WINS, 403–404
- registration authority (RA), 70
- Registry
 - editing, 755
 - function of, 94
 - keys, auditing, 821–822
- regulations, 26–27
- relative ID master, 60
- relative IDs (RIDs), 60
- Remote Access Policies, 494–495
- remote access policies
 - creating, 515–528
 - with IAS, 309
 - included in NAQC, 524
 - for router-to-router VPN, 306
 - summary of, 531
 - for VPN connection, 500
 - for wireless connections, 502–503
- remote access policies, creating
 - authorizing remote access, 516–520
 - controlling remote connections, 525–528
 - policies/profiles, 515–516
 - restricting remote access, 520–525
- remote access profile
 - controlling remote connections, 525–528
 - function of, 515
- Remote Access Quarantine Agent service (RQS.EXE), 524
- remote access strategy
 - analyzing organizational needs, 487
 - analyzing user needs, 487
 - authentication methods for, 508–512
 - callback security for, 513
 - connections, managing, 513–514
 - data encryption level for, 512–513
 - dial-in access design considerations, 489–495
 - domain functional level and, 505–508
 - in general, 486
 - Network Access Quarantine control, 514
 - remote access policies, creating, 515–528
 - remote access types to allow, 487–489
 - smart cards for, 514
 - summary of, 529–531
 - VPN design considerations, 495–500
 - wireless remote access design considerations, 500–505
- remote access types, 487–489

- remote access/VPN server role, 57
- remote administration, 668
- remote connections, 525–528
- Remote Desktop Connection (RDC), 753
- remote management, NLB, 683–684, 691
- Remote Procedure Call (RPC), 683–684
- removable storage, 22
- renaming tool, 85–86
- rendom.exe (domain rename utility), 86
- renewal interval, 447
- replication
 - of domain controllers, 59
 - linked value replication, 89
 - PDC emulator and, 61
 - WINS server deployment and, 424
- replication partnership
 - accepting with WINS, 424
 - configuration, 428–434
 - summary of, 466–467
 - troubleshooting, 459–460
- replication, WINS, 427–437
 - change-only replication, 428
 - convergence time factors, 427
 - multihomed WINS server and, 440
 - push *vs.* pull replication, 474
 - replication models, 434–437
 - replication partnership configuration, 428–434
 - summary of, 466–467
 - troubleshooting, 459–460
- report, backup, 604
- Request Security IPSec policy, 733
- requester, 807
- requests, certificate, 892–895
- Require Security IPSec policy, 733
- Reservations button, 295
- reset account lockout counter after setting, 826
- resource cluster groups, 642–643
- resource records (RRs)
 - components of, 342–344
 - DNS namespace design and, 357–358
 - DNS server capacity and, 371
 - domains *vs.* zones, 345–348
 - function of, 341
 - multiple namespaces and, 363–364
 - secure updates and, 389–390
 - update with DNS/DHCP interaction, 387–389
- resources, physical, 194
- Respond Only IPSec policy, 732–733
- restore
 - ASR, 613, 614
 - from backup, 605–607
 - of cluster servers, 657
 - WINS database, 452
 - See also* recovery
- Restore to Alternate Location feature, 601
- Restore Wizard, 603–604
- Restricted Groups, 94
- restriction of remote access, 520–525
- Resultant Set of Policy (RSoP)
 - defined, 713
 - for Group Policy modeling, 4
 - IPSec planning with, 765–768
 - XP IPSec policies and, 752
- reverse lookup records, 401–402
- reverse lookup zones
 - creating, 356
 - described, 352
 - security considerations for, 353
 - update with DNS/DHCP interaction, 387–389
- reversible encryption, 825
- revocation, certificate, 886–887
- RF (radio frequency), 801
- RIDs (relative IDs), 60
- ring replication model, 434–435
- RIP. *See* Routing Information Protocol (RIP)
- RIP version 1 (RIPv1), 226, 227
- RIP version 2 (RIPv2), 226–227, 252–255
- risk, 91–92
- Rivest, Ron, 802, 864
- rogue router, 227
- rogue servers, 126
- rogue WLANs, 812
- role-based administration, 896
- roles, 27
 - See also* server roles
- root CAs
 - capabilities of, 911
 - overview of, 872
 - security and, 129, 885
 - at top of hierarchy, 72
- root hints file, 349–350
- root zone, 367–368
- rootsec template, 96
- round robin, 456
- routable protocols, 238
- route add command, 251–252

- route command, 170
- route entry, 217–218
- route table, 168–169
- routed connections
 - advantages of, 289
 - hardware/software routers, 289–290
 - IP addressing for, 290
 - summary of, 326
- router
 - components of, 259
 - defined, 222
 - function of, 244–245
 - hardware/software routers, 289–290
 - setting up Windows Server 2003 as, 245–257
- router-to-router VPNs
 - connection types for, 303–304
 - on demand/demand-dial connections, 304–306
 - described, 263
 - persistent connections, 306
 - remote-access policies, 306
 - Windows Server 2003 as, 267–268
- routes
 - minimizing number of, 260
 - types of, 216–217
- routing, 150
- Routing and Remote Access console
 - managing NAT from, 294
 - for troubleshooting routing, 271
 - VPN connections with pre-shared keys, 763
- Routing and Remote Access Server Setup Wizard
 - configuring VPN server with, 302–303
 - installing NAT with, 292–293
 - for Windows 2003 Server as static router, 248–249
- Routing and Remote Access Service (RRAS)
 - configuring OSPF, 255–257
 - configuring RIPv2 on router, 252–255
 - configuring Windows 2003 Server as static router, 246–251
 - EAP and, 905
 - IAS integration with, 309
 - packet-filtering methods of, 279
- Routing and Remote Access Service (RRAS) server
 - activating IAS authentication for, 310–312
 - assigning IP addresses with, 490–491
 - authentication methods for, 512
 - restricting authentication methods in, 525
 - routing problems and, 274–275
 - supports multiple functions, 328
- Routing Information Protocol (RIP)
 - as distance-vector protocol, 225
 - OSPF benefits over, 229–230
 - problems with, 275, 276
 - RIP router process, 228–229
 - RIP v1/v2, 226–227
- routing options, 236–245
 - bridges, 240–242
 - connectivity devices, selecting, 236–237
 - hubs, 239–240
 - OSI model review, 237–239
 - routers, 244–245
 - switches, 242–244
- routing protocols
 - distance-vector/link-state protocols, 225–226
 - dynamic routing and, 221
 - minimizing number of, 260–262
 - Open Shortest Path First, 229–233
 - problems with, 274–276
 - Routing Information Protocol, 226–229
- routing security, 257–270
 - IPSEC security features/process, 257–258
 - logging level, 269–270
 - network topology, simplifying, 259–262
 - packet filtering/firewalls, 268–269
 - router-to-router VPNs, 263–268
 - routing components, requirements for, 259
 - summary of, 278
- routing strategy
 - evaluating routing options, 236–245
 - gateways, 222–225
 - IP addresses, 213
 - IP version 6, 215–216
 - NAT components, 214–215
 - netsh commands, 233–236
 - routing concept, 212
 - routing protocols, 225–233
 - routing tables, 216–220
 - security, 257–270
 - static *vs.* dynamic routing, 220–222
 - troubleshooting IP routing, 270–276
 - Windows Server 2003 as router, 245–257
- routing tables
 - configuration problems, 276
 - defined, 216
 - route entry component parts, 217–218
 - rows, 220
 - types of routes, 216–217

- viewing, 219
- rows, 220
- RPC (Remote Procedure Call), 683–684
- RQS.EXE, 524
- RRAS. *See* Routing and Remote Access Service (RRAS)
- RRs. *See* resource records (RRs)
- RSA Labs, 864
- RSA technology, 864–866
- RSoP. *See* Resultant Set of Policy (RSoP)
- Run As command
 - for administrative server tasks, 785
 - to view routing table, 219

S

- S/MIME (Secure/Multipurpose Internet Mail Extensions), 881
- SA. *See* Security association (SA)
- SACL (system access control list), 783, 784
- Safe Mode boot, 614
- scalability, 16, 309
- scale of nines, 618
- scanning, 834
- scavenging
 - of DNS records, 391–392
 - of WINS records, 446–447
- schedule
 - deployment, 29
 - network planning and, 18–19
 - test, 34–35
- schema
 - disabling objects, 89
 - function of, 60
 - securing, 790
- Schema Admins group, 790
- schema master, 60
- Schlumberger smart cards, 899
- scope
 - AD-integrated replication scope, 379–382
 - group scope, 792
 - of services, SLAs and, 27
- SCSI-based controllers, 659
- SCSI (Small Computer System Interface) interface, 565, 643–644
- secedit /analyze command, 100–101, 828–829
- secedit /configure command
 - function of, 100, 828
 - syntax for/parameters of, 101
- secedit /export command
 - function of, 100, 828
 - syntax for/parameters of, 102
- secedit /GenerateRollback command
 - function of, 100, 829
 - syntax for/parameters of, 102–103
- secedit /import command
 - function of, 100, 828
 - syntax for/parameters of, 102
- Secedit utility
 - for applying security settings, 139
 - commands, 100–103
 - function of, 99
 - for security templates, 828–829
 - for template settings, 140
- secedit /validate command
 - function of, 100, 829
 - syntax for/parameter of, 102
- second-generation (2G), 804
- secondary server, 347–348, 373
- secure dynamic updates
 - BIND support of, 397
 - enabling, 389
 - GSS-TSIG and, 391
- Secure/Multipurpose Internet Mail Extensions (S/MIME), 881
- Secure Password Authentication (SPA), 128–129
- Secure Server IPSec policy, 733
- securedc template, 97
- securews template, 97
- security
 - AD structure and, 44
 - CA, planning, 885
 - configurations, deploying, 134–136
 - of Connection Manager, 324–325
 - for DDNS and DHCP, 389–391
 - DNS namespace design and, 357–358
 - DNS security issues, 404–412, 463
 - IPSec, 761–764
 - with Layer 4 switches, 244
 - levels for IPSec, 727–728
 - NAT limitations, 215
 - network authentication and, 45
 - network planning and, 15–16
 - for NLB cluster, 690–691
 - with private root zone, 367
 - protocols and, 150
 - remote access security, 505–514, 531

- reverse lookup zones and, 353
- RSoP and, 766–767
- of server clusters, 667–669
- VPN encryption protocols, 307–308
- of WINS, 449–451
- wireless encryption and, 504–505
- for zone replication, 382
- See also* authentication; baseline security; routing security; server security; wireless security
- security, AD. *See* Active Directory (AD) security
- Security association (SA)
 - data transit and, 714
 - defined, 713
 - IPSec, 258
 - IPSec driver and, 725–726
 - main mode, 714
 - overview of, 770
 - process, 713–714
- Security Configuration and Analysis
 - analyzing computer with, 103–108
 - to apply security templates, 109
 - for custom security templates, 131, 140
 - for security templates, 827–828
 - tasks performed with, 98
- security descriptors, 782–783
- security framework, 782–847
 - Active Directory security, 782–800
 - change and configuration management framework, 830
 - monitoring/optimizing security, 817–829
 - security update infrastructure, 830–847
 - summary of, 848
 - wireless security, 801–816
- security groups, 86–87
- security identifier (SID)
 - filtering, 793
 - relative ID master and, 60
 - in user authentication, 800
- Security log, 584
- Security log settings, 823
- Security Parameters Index (SPI), 720, 721
- security policies
 - account lockout policies, 826
 - Kerberos policies, 825–826
 - password policies, 824–825
 - security templates, 827–829
 - setting, 823–824
 - settings, 134–135
 - user rights, 826–827
- security principals, 60, 795
- Security properties, IAS, 310–311
- security requirements
 - configurations for, 93–94
 - identifying, 91–93
- security settings, enforcing, 109–112
- Security Settings extension to Group Policy, 827
- Security tab, 254–255
- security templates
 - applying, 109–112
 - custom, creating, 131–134
 - custom, tools for, 139–140
 - planning secure baseline installation parameters, 103–108
 - summary of, 137
 - tools for, 827–829
- security templates and tools, 94–103
 - Group Policy Object Editor, 99
 - policies/settings, 94–95
 - predefined templates, 95–97
 - Secedit utility, 99–103
 - Security Configuration and Analysis tool, 98
- Security Templates MMC snap-in
 - for creating/editing templates, 94–95
 - for custom security templates, 131–134, 140
- security update infrastructure, 830–847
 - Microsoft Baseline Security Analyzer, 831–837
 - Microsoft Software Update Services, 837–847
 - Permcopyp.exe, 830–831
 - security updates, importance of, 831
 - Subinacl.exe, 830
 - summary of, 850
- segment switching, 242–243
- segments, network, 149–150
- sender authentication, 496–497
- Sequence Number field, 720, 721
- Serial Line Internet Protocol (SLIP), 488
- server cluster
 - creating new, 653–654, 670–677
 - defined, 641
 - groups, resource types, 642–643
 - name resolution, 643
 - node failure, recovering from, 657
- server cluster deployment options, 647–653
 - consideration of, 647
 - failover ring, 651–652
 - hot-standby server/N+1, 649–651

- N-node failover pairs, 648–649
- random, 652–653
- server cluster models, 644–647
 - majority node set, 646–647
 - model most frequently used, 701
 - single node, 644–645
 - single quorum device, 645–646
- server cluster nodes
 - described, 641–642
 - in failover ring, 651–652
 - failure, recovering from, 657
 - in hot-standby server/N+1 deployment option, 649–651
 - N-node failover pairs, 648–649
 - number of, 701
 - in random deployment option, 652–653
 - security of, 667–669
 - of single node model, 644–645
 - of single quorum device model, 645–646
- server clustering, 641–677
 - administration, 653–656
 - cluster models, 644–647
 - cluster network configuration, 662–667
 - cluster node failure, recovering from, 657
 - creating new cluster, 670–677
 - deployment options, 647–653
 - hardware issues, 658–662
 - Network Load Balancing *vs.*, 678
 - overview of, 640
 - questions/answers about, 701
 - relationship to NLB, 681–682
 - security, 667–669
 - summary of, 699–700
 - terminology/concepts, 641–644
- Server IPsec policy, 733
- server log files, 593
- Server Message Block (SMB), 646
- server principal name (SPN), 800
- server roles, 54–77
 - application servers, 75–77
 - application servers, securing, 130
 - certificate authorities, 69–75
 - certificate authorities, securing, 129
 - database servers, 68
 - database servers, securing, 127–128
 - DHCP, DNS, WINS servers, 63–65
 - DHCP, DNS, WINS servers, securing, 125–126
 - domain controllers, 58–62
 - domain controllers, securing, 121–122
 - file and printer servers, 62–63
 - file servers, securing, 121–124
 - mail servers, 68–69
 - mail servers, securing, 128–129
 - Manage Your Server tool, 54–58
 - print servers, securing, 124–125
 - security issues of all server roles, 113–121
 - security requirements and, 93–94
 - summary of, 137
 - terminal servers, 78
 - terminal servers, securing, 130–131
 - Web servers, 65–68
 - Web servers, securing, 126–127
- server room, 114
- server security, customizing, 113–136
 - for application servers, 130
 - for certificate authorities, 129
 - custom security templates, 131–134
 - for database servers, 127–128
 - deploying security configurations, 134–136
 - for DHCP, DNS, WINS servers, 125–126
 - for domain controllers, 121–122
 - for file servers, 121–124
 - for mail servers, 128–129
 - for print servers, 124–125
 - security issues of all server roles, 113–121
 - summary of, 138
 - for terminal servers, 130–131
 - for Web servers, 126–127
- server security strategy
 - configurations for security requirements, 93–94
 - in general, 78–79
 - operating system, choosing, 79–90
 - security requirements, identifying, 91–93
 - summary of, 137, 138
- servers
 - for Internet Authentication Protocol, 309–310
 - monitoring with System Monitor tool, 570–580
 - placement/performance of, 197–198
 - smart cards and, 898
 - upgrades, 43
 - virtualization, 625
- Service Level Agreement (SLA), 26–27
- service locator record (SRV), 343, 362
- service logs, 593
- service packs, 115–117
- service profiles

- with CMAK, 320–323
 - options of, 328
 - preventing editing of, 324
 - secure distribution of, 325
- service set identifier (SSID), 801–802, 814
- Service Settings dialog box, 300
- service ticket, 81, 825
- services
 - adding custom service for ICS, 299–300
 - configuring for ICS, 298–299
 - disabling unneeded, 117
 - See also* specific service
- Services and Ports tab, 295
- Session layer, OSI model, 238
- session time, maximum, 525–527
- setup security template, 97, 103
- sever cluster node, 643–644
- sexual harassment, 26
- Shamir, Adi, 864
- share permissions, 788, 789
- shared cluster disks, 659
- shared-key authentication, 807
- shared secret, 312
- shared secret key cryptographies, 864
- SharePoint, 20–21
- Shinder, Debra Littlejohn, 800
- Shiva Corporation, 509
- Shiva Password Authentication Protocol (SPAP)
 - disabling, 509–511
 - for IAS authentication, 314
- shortest path first (SPF), 225
- show helper command, 235
- SID. *See* security identifier (SID)
- signature files, 117
- signatures. *See* digital signatures
- Simple Mail Transport Protocol (SMTP), 66, 68–69
- simple query test, 413
- single host filtering mode, 679
- single-instancing, 786
- single node server cluster model, 644–645
- single point of failure, 407
- single quorum device server cluster model
 - described/illustrated, 645–646
 - N-node failover pairs deployment option, 648–649
- SLA (Service Level Agreement), 26–27
- slave drive, 565
- sliding window, 198
- SLIP (Serial Line Internet Protocol), 488
- Small Computer System Interface (SCSI) interface, 565, 643–644
- smart cards
 - authentication in PKI, 897
 - authentication, process of, 898
 - EAP-TLS supports, 317
 - implementing/using, 900–903
 - logon, deploying, 898–899
 - overview of, 897–898
 - PKI and, 908
 - readers, 899
 - for remote access strategy, 514
 - for remote access VPNs, 903–905
 - Terminal Server logon with, 906
 - Windows 2000 support of, 81
 - Windows logon with, 899
- SMB (Server Message Block), 646
- SMS (Systems Management Server), 4, 759
- SMTP (Simple Mail Transport Protocol), 66, 68–69
- SOA record. *See* Start of Authority (SOA) record
- soft association, 764
- software
 - network testing, 30–31
 - performance testing, 46–47
- software router, 290
- Software Update Services (SUS), 837–847
 - configuring clients with Group Policy, 844–845
 - configuring clients with Local Security Policy, 843–844
 - installing, 838–839
 - parts of, 852
 - setting options, 845–846
 - using, 839–843
- Software Update Services (SUS) server component
 - downloading updates, 840–841
 - function of, 852
 - required for SUS, 838
 - setting options for, 845–847
 - synchronizing, 839–840
- source address, 212
- SPA (Secure Password Authentication), 128–129
- spam, filtering, 17
- SPAP. *See* Shiva Password Authentication Protocol (SPAP)
- Special Permissions option, 789
- speed-buffering bridge, 242
- speed, wireless equipment, 501
- SPF (shortest path first), 225
- SPF tree, 231

- SPI (Security Parameters Index), 720, 721
- spindle count, 566
- split-brain
 - majority node set and, 647
 - quorum resource to prevent, 644
- split DNS configuration
 - described, 398–399
 - for DNS security, 411
- split horizon, 229
- split horizon with poison reverse, 229
- split seek, 568
- split WINS registrations, 444, 467
- SPN (server principal name), 800
- spoofing, 812
- SQL Server
 - function of, 68
 - security features of, 127–128
 - username/password in, 128
- SRV (service locator record), 343, 362
- SSID. *See* service set identifier (SSID)
- stack, protocol, 149
- stand-alone CAs
 - CA security and, 885
 - overview of, 882–883
 - use of, 72
- Start of Authority (SOA) record
 - of resource record, 343–344
 - in reverse lookup zone, 356
 - troubleshooting host name resolution and, 456
 - zone transfer and, 378–379
- stateful filtering, 751
- static access control, 782–784
- static address pool, 490
- static IP address, 666
- static IP route, 251–252
- static mappings
 - for redirection attack protection, 450
 - static WINS entries, 438–439, 467
 - summary of, 465
 - troubleshooting, 458
- static router, 246–251
- static routing, 220–222, 245
- static WINS entries, 438–439
- statistics, IPSec, 753–755
- stealth servers, 374, 411
- storage, data, 21–23
- storage device
 - node connected to, 643
 - for server cluster, 659–662
 - single quorum device and, 645
- streaming media server, 57
- streaming media services, 26
- striping. *See* RAID 0
- striping with parity. *See* RAID 5
- strong passwords
 - elements of, 118–119
 - Group Policy to enforce, 785
- stub zone
 - for child domain authority, 347
 - for disjointed namespace, 365–366
 - zone replication planning and, 383
- subdomain, 364–365
- Subinacl.exe, 830
- subnet masks
 - custom, 179–180
 - with private addressing, 214
 - standard, 178–179
- subnets
 - ANDing/binary numbering, 175–177
 - CIDR and, 180–181
 - classful addressing, 173–175
 - schemes, creating, 173
 - subnetting networks, 177–180
- subordinate CAs, 72, 872
- subtype
 - defining on client computer, 809–810
 - defining on domain controller, 808–809
- Success Audit event type, 585
- superseded templates, 890
- supplicant, 804–806
- SUS. *See* Software Update Services (SUS)
- switches
 - authenticating with IAS, 318
 - segment/port switching, 242–243
 - types of, 244
 - UPSs for, 625
- switching hub, 240, 243
- symmetric key encryption, 864
- /sync parameter, 136
- Synchronization Log, 841–842
- syskey (System Key Utility), 786
- system access control list (SACL), 783, 784
- System Key Utility (syskey), 786
- System log, Event Viewer, 584–585
- System Monitor
 - console, creating, 580–584

- described, 195, 196–197
 - log data, viewing, 576–578
 - to monitor IAS, 313
 - overview of, 626
 - Performance console for monitoring DNS server, 415–416
 - for servers, using, 570–580
 - System Overview counter log, 574–576
 - system performance comparisons with, 578–579
 - System Overview counter log, 574–576
 - system requirements, 79–80
 - System Services, 94
 - system state data, 600–601
 - Systems Management Server (SMS), 4, 759
- ## T
- tapes, backup, 605
 - /target:{*computer* | *user*} parameter, 136
 - task management, centralized, 18–19
 - Task Manager, 580
 - Task Scheduler, 19
 - TCO (Total Cost of Ownership), 10, 27–28
 - TCP/IP. *See* Transmission Control Protocol/Internet Protocol (TCP/IP)
 - TCP port 53, 412
 - TCP (Transmission Control Protocol), 162
 - TCP1323Opts-enabled computer, 198
 - teamed configuration, 658
 - templates, 887–892
 - See also* security templates
 - Temporal Key Integrity Protocol (TKIP), 816
 - temporary files, 123
 - Teredo, 193
 - Terminal Server, 906
 - terminal servers
 - defined, 57
 - function of, 78
 - securing, 130–131
 - Terminal Services
 - benefit of, 78
 - centralized, 23–24
 - for server cluster nodes, 668
 - test environment
 - building, 30–33
 - implementing, 34–36
 - overview of, 29–30, 39
 - planning, 30–33
 - test lab, 746
 - testing, performance, 46
 - TGT (ticket-granting ticket), 81, 898
 - thin-client technology, 78
 - third-party certification authorities (CAs), 870–871
 - threats, 91–92
 - three-tier model, CA, 881–882
 - thresholds, counter, 578
 - ticket granting authority, 791
 - ticket-granting ticket (TGT), 81, 898
 - ticket, Kerberos, 81
 - time restriction, 523
 - timestamp, 422
 - TKIP (Temporal Key Integrity Protocol), 816
 - TLD (top-level domain) name, 358–359
 - Token Ring networks, 569–570
 - tombstoned record, 423
 - tombstoning, 473–474
 - tools
 - backup, 602–604
 - planning, 3–4
 - for troubleshooting, 271–273
 - See also* security templates and tools
 - top-level domain (TLD) name, 358–359
 - topologies, Server 2003 supported, 568
 - topology
 - network testing and, 31
 - test lab, 33
 - topology, network
 - performance and, 569–570
 - planning, 193–194, 201
 - Total Cost of Ownership (TCO), 10, 27–28
 - Tracert utility
 - IPv6 parameters, 192
 - to test TCP/IP connections, 273, 279–280
 - tracing, IKE, 757–758
 - traffic distribution, 679–680
 - traffic, network
 - management planning, 194–195, 202
 - monitoring, 195–197
 - translated connections
 - described, 290
 - Internet Connection Sharing, 297–300
 - Network Address Translation, 291–297
 - summary of, 326
 - translating (or translational) bridge, 241
 - translation component, NAT, 214
 - Transmission Control Protocol/Internet Protocol (TCP/IP)
 - basics, 160–164
 - checking configuration/testing connections, 279–280

- configuration problems, 276
 - configuring properties for DNS installation, 354
 - connection testing with PING, 271
 - enhancements in Server 2003, 164
 - hardware acceleration, disabling, 760–761
 - manual configuration, 155–160
 - overview of, 148
 - settings for interconnects, 666–667
 - Tracert to test connections, 273
 - Transmission Control Protocol/Internet Protocol (TCP/IP) infrastructure
 - addressing requirements, 171–172
 - bandwidth requirements, 198
 - basics, 160–164
 - IP addressing, troubleshooting, 181–183
 - IPv6, transitioning to, 183–193
 - manual configuration, 155–160
 - multiprotocol environments, 153–156
 - network topology, 193–194
 - network traffic management, 194–197
 - performance, optimizing, 198–199
 - protocol suite, 151–153
 - Server 2003 network protocols, 148–151
 - Server 2003 TCP/IP features, 164–170
 - subnetting scheme, creating, 173–181
 - Transmission Control Protocol (TCP), 162
 - transparent bridge, 240, 241
 - Transport layer
 - function of, 238
 - Layer 4 switches operate at, 244
 - TCP/IP, 161–162
 - transport mode
 - ESP/AH and, 719
 - IPSec, 711, 718
 - triggered updates
 - logging choices for, 253–254
 - RIP routers and, 229
 - Triple Data Encryption Standard (3DES)
 - defined, 713
 - encryption for user account, 799
 - IPSec and, 761–762
 - of Server 2003, 868
 - VPN connections and, 715
 - on Windows 2000, 734
 - troubleshooting IP addressing, 181–183
 - troubleshooting IP routing, 270–276
 - common routing problems, 274–276
 - summary of, 278
 - troubleshooting tools, 271–273
 - troubleshooting, IPSec
 - with IKE tracing, 757–758
 - netdiag for Server 2003, 751
 - with Network Monitor, 759–760
 - overview of, 751
 - with packet event logging, 755–757
 - Policy Assignment Information, viewing, 752
 - statistics, viewing, 753–755
 - TCP/IP, IPSec hardware acceleration, 760–761
 - troubleshooting name resolution, 452–460
 - in general, 452–453
 - host name resolution, 453–457
 - key points about, 471
 - NetBIOS name resolution, 457–460
 - summary of, 468–469
 - trust relationships
 - in AD security scenarios/solutions, 785–786
 - cross-domain relationships, 791–792
 - cross-forest relationships, 793–795
 - forest trusts, 89
 - between multiple domains, 851
 - Windows domain upgrades and, 42
 - tunnel mode
 - described, 772
 - ESP/AH and, 719
 - IPSec, 308, 717–718
 - tunneling protocol, 301, 306–307
 - two-way initiation, 306
 - two-way transitive trust, 365
- ## U
- UCS-2, 360
 - UDP packet, 351
 - UDP port 53, 412
 - UDP (User Datagram Protocol), 162
 - UNC (Universal Naming Convention), 338
 - underscore character (_), 360–361, 396
 - Unencrypted Password, 314
 - unicast addresses, 213
 - unicast messages, 421
 - Uniform Resource Locator (URL), 62
 - Uninterruptible Power Supply (UPS), 114, 625
 - unique NetBIOS names, 418
 - universal group scope, 792
 - Universal Naming Convention (UNC), 338
 - universal security groups, 87
 - Universal Transformation Format-8 (UTF-8), 360
 - UNIX, 399–400
 - updates
 - convergence time and, 427

- from Microsoft Windows Update Web site, 115–117
 - with Software Update Services, 837–847
 - WINS change-only replication and, 428
 - See also* dynamic updates; security update infrastructure
 - upgrades
 - hardware, 29
 - server, 43
 - testing, 34
 - UPS (Uninterruptible Power Supply), 114, 625
 - URL (Uniform Resource Locator), 62
 - US-ASCII, 359–360
 - user
 - access, 118
 - authentication, 308–309, 800
 - authorization, 784
 - authorizing remote access by, 516–518
 - enabling remote access by, 493–494
 - locations, network planning and, 11
 - logon restrictions, 825
 - managing remote access by, 532
 - priorities, network planning and, 17–26
 - remote access needs, 487
 - restricting remote access by, 521
 - user account
 - network testing and, 31
 - security, 796–797
 - security scenarios, 798–800
 - user certificates, 497, 870, 896
 - User Datagram Protocol (UDP), 162
 - user rights
 - assigning, 826–827
 - for backups, 594–595
 - IPSec policies and, 742
 - user ticket, 825
 - UTF-8 (Universal Transformation Format-8), 360
- ## V
- validated writes, 788
 - verification interval, 447
 - version ID
 - NetBIOS name resolution and, 460
 - WINS replication and, 428, 473
 - video services, 26
 - virtual LAN (VLAN), 503
 - virtual memory (VM), 561–562
 - virtual private network (VPN)
 - adding demand-dial interface and, 261
 - configuring VPN connection from client computer, 266–267
 - connections, CMAK's support of, 324
 - data encryption level for, 512–513
 - Internet-based VPNs, 301–303
 - IPSec encryption and, 715
 - packet filtering for, 269
 - protocols, 306–307, 328, 496–497
 - remote access, 488–489, 529
 - remote access with smart cards, 903–905
 - router-to-router VPNs, 263, 267–268, 303–306
 - security, 307–308
 - summary of, 326, 327
 - tunnels, 382
 - Windows Server 2003 VPN Server, installing, 263–265
 - virtual private network (VPN) remote access design, 495–500
 - access policies for, 500
 - computer certificates, installing, 497–499
 - firewall filters, configuring, 499–500
 - in general, 495
 - summary of, 530
 - VPN protocol selection, 496–497
 - virtual private network (VPN) server
 - firewall filters for, 499–500
 - role, 57
 - virtual server
 - in hot-standby server/N+1 deployment option, 650–651
 - in N-node failover pairs, 648, 649
 - in single node server cluster, 644–645
 - in single quorum device server cluster, 646
 - virtualization, 625
 - See also* clustering
 - virus, 117
 - VLAN (virtual LAN), 503
 - VM (virtual memory), 561–562
 - VMware, 30
 - VMware Workstation 4.0, 353
 - volume shadow copy
 - backups and, 43
 - described, 22
 - system file backup with, 599
 - VPN. *See* virtual private network (VPN)
 - vulnerabilities, 115

W

- /wait: *Value* parameter, 136
- WAN (Wide Area Network) links, 370, 532
- WAPs. *See* wireless access points (WAPs)
- Warning event type, 585
- warranties, hardware, 29
- Web applications, 75
- Web Edition servers, 58
- Web servers
 - certificate requests from, 894–895
 - configuration, 67–68
 - function of, 65–66
 - protocols, 66
 - securing, 126–127
 - setting up, 139
 - use of certificates, 70
- Web sites
 - for Automatic Updates software, 843
 - for constrained delegation article, 800
 - for DCOM white paper, 684
 - for Dijkstra algorithm, 231
 - for domain name registrars, 358
 - InterNic, 359
 - on IPv6, 184
 - for MBSA download, 832
 - on Public Key Interoperability, 871
 - for reverse lookup zones security, 353
 - Server 2003 networking/communication enhancements, 164
 - for VMware, 30, 353
 - for Windows Server 2003 download, 353
- WEP. *See* Wired Equivalent Privacy (WEP)
- Whois application, 359
- Wi-Fi Protected Access (WPA), 505, 816
- Wide Area Network (WAN) links, 370, 532
- Windows 2000
 - 3DES algorithm use on, 734
 - disk defragmentation and, 567
 - forest functional level, 88
 - functional levels and, 84
 - IPSec and, 723
 - IPSec monitoring on, 750
 - IPSec policy precedence viewing tools, 752
 - IPSec policy/statistics viewing tools, 753
 - ipsecmon command, 750, 753
 - modes, 83
 - NAT and, 723
 - netdiag tool and, 751
- Windows 2000 Advanced Server, 80, 139
- Windows 2000 Datacenter, 80
- Windows 2000 Mixed, 506
- Windows 2000 mixed mode
 - features in, 83–84
 - group nesting in, 86–87
- Windows 2000 Native, 506
- Windows 2000 native functional level, 84, 86–87
- Windows 2000 Server, 80, 81
- Windows Authentication
 - authentication with, 127
 - remote access policies for, 515
 - for RRAS server, 512
- Windows Backup utility
 - overview of, 594–602, 626
 - using, 602–604
- Windows Calculator, 176–177
- Windows cluster, 436, 466
- Windows-Groups attribute, 521
- Windows Internet Naming Service (WINS)
 - BIND secondaries and, 396–397
 - client configuration, 440–444
 - database backup/restoration, 451–452, 601–602
 - for fault tolerant NetBIOS name resolution, 418
 - improvements to, 424
 - interoperability with Windows Server 2003 DNS, 399–404
 - logging, 450
 - multihomed servers, 439–440
 - name registration, renewal, release, 421–423
 - performance issues, 444–449
 - proxy agent, 443–444
 - referral zone, 403–404
 - replication, 427–437
 - security issues, 449–451
 - server deployment, 424–426
 - split registrations, 444
 - static entries, 438–439
 - summary of, 464–468
 - troubleshooting NetBIOS name resolution, 457–460
 - Users group, 451
- Windows Internet Naming Service (WINS) database
 - backup and restore, 451–452
 - compaction, 448
 - consistency checking, 448–449
 - scavenging records, 446–447
 - security of, 450–451
 - split WINS registrations and, 444
- Windows Internet Naming Service (WINS) records
 - backup/restore, 451–452

- consistency checking, 448–449
- scavenging, 446–447
- Windows Internet Naming Service (WINS)
 - replication, 427–437
 - change-only replication, 428
 - convergence time factors, 427
 - replication models, 434–437
 - replication partnership configuration, 428–434
- Windows Internet Naming Service (WINS) server
 - defined, 57
 - deployment, 424–426
 - function of, 65
 - multihomed, 439–440
 - multiple server addresses, 443
 - name registration, renewal, release, 421–423
 - for NetBIOS name resolution, 419–420
 - performance issues, 198, 444–449
 - securing, 125, 126
 - split WINS registrations, 444
 - summary of, 465–469
 - troubleshooting NetBIOS name resolution, 457, 458–460
 - WINS replication, 427–437
- Windows Load Balancing Service (WLBS), 678, 688–689
- Windows Management Instrumentation (WMI), 62, 683
- Windows Media Services. *See* Media Services
- Windows NT, 83–84
- Windows NT 4, 31, 567
- Windows NT 4.0 Enterprise Edition, 641
- Windows NT 4.0 Service Pack 6, 116
- Windows NT Server 4
 - AD can not be used on, 81
 - security templates and, 96, 97
 - system requirements for, 80
- Windows Server 2003
 - certificate-based PKI, planning, 862–875
 - default configuration, 626
 - Disk Defragmenter, 566–568
 - DNS interoperability, 392–404
 - DNS service/reverse lookup zones, 353–357
 - Event Viewer of, 584–592
 - functional levels, 83–90
 - infrastructure planning, 2–11
 - IPSec policy viewing tools, 752, 753
 - multiple object selection in, 139
 - NAT traversal, 723
 - netdiag tool, 751
 - network protocols, 148–151, 200–201
 - networking features, 155
 - organizational needs, analyzing, 11–29
 - PKI of, 865
 - planning/design process, 36–38
 - RSoP console, 765
 - security features of, 81–83
 - server clustering with, 641
 - server roles on, 54
 - smart cards and, 898–899
 - System Monitor, 196–197
 - TCP/IP features, 164–170
 - test environment, developing, 29–36
 - topologies supported by, 568
- Windows Server 2003 as router, 245–257
 - adding static IP route, 251–252
 - checklist for, 245
 - configuring as static router, 246–251
 - OSPF, configuring, 255–257
 - RIP version 2, configuring, 252–255
- Windows Server 2003 Datacenter Edition, 80
- Windows Server 2003 Enterprise Edition, 80
- Windows Server 2003 forest functional level, 88–89
- Windows Server 2003 functional level, 86–87, 506
- Windows Server 2003 Interim, 506
- Windows Server 2003 interim forest functional level, 88
- Windows Server 2003 interim level, 84
- Windows Server 2003 level, 84
- Windows Server 2003 Standard Edition, 80
- Windows Server 2003 VPN Server, 263–265
- Windows Server 2003 Web Edition, 80
- Windows Server Catalog, 658
- Windows Sockets. *See* Winsock (Windows Sockets)
- Windows Terminal Services. *See* Terminal Services
- Windows Update site. *See* Microsoft Windows Update site
- Windows XP
 - AD-based IPSec policies and, 748
 - IP Security Monitor on, 750
 - IPSec configuration on, 731
 - IPSec policy precedence viewing tools, 752
 - IPSec policy/statistics viewing tools, 753
 - netdiag tool and, 751
 - netsh command-line utility and, 731
 - viewing monitoring information on, 760
- Windows XP client
 - MBSA on, 833–837
 - setting up for wireless networking, 808
- WINS. *See* Windows Internet Naming Service (WINS)

WINS-R record, 401–402

WinSock (Windows Sockets)

- host names and, 338–339
- TCP/IP and, 153
- to troubleshoot name resolution, 453

Wired Equivalent Privacy (WEP)

- described, 504
- function of/implementations of, 802
- security of, 815–816
- settings, 814–815
- weaknesses, 815, 853
- in wireless network authentication, 806–807

Wired Equivalent Privacy (WEP) keys

- AirSnort to capture, 813
- changing regularly, 816
- weakness of, 815
- in wireless network authentication, 807

wireless access points (WAPs)

- configuring as RADIUS clients, 503–504
- configuring wireless networking and, 808
- default settings and, 814
- EAP authentication and, 804
- IAS support of, 318
- multiple, 503
- RADIUS authentication of, 501
- vulnerability of, 852–853

Wireless Configuration service, 817

wireless local area network (WLAN), 804, 806

wireless metropolitan area network (WMAN), 804

Wireless Monitor, 807, 817–818

wireless network interface card (NIC), 808

wireless personal area network (WPAN), 803

wireless remote access

- design considerations, 500–505
- summary of, 529
- when to use, 489

wireless security, 801–816

- 802.11 specifications, 801–803
- authentication methods, 806–810
- authentication protocols, 810–812
- EAP authentication, 804–806
- issues, 812–816
- summary of, 848, 849
- vulnerability of wireless network, 852–853
- wireless network types, 803–804
- wireless networking technologies, 806

wireless wide area network (WWAN), 804

Wireless Zero Configuration service, 801

WLAN (wireless local area network), 804, 806

WLBS (Windows Load Balancing Service), 678, 688–689

WMAN (wireless metropolitan area network), 804

WMI (Windows Management Instrumentation), 62, 683

WPA (Wi-Fi Protected Access), 505, 816

WPAN (wireless personal area network), 803

write caching, 659, 660

WWAN (wireless wide area network), 804

X

X.509 standard

- CAs and, 875
- for certificates, 71, 910

Z

.zone, 350

zone

- authoritative name servers for, 373–374
- domain *vs.*, 345–348, 461–462, 472
- replication, 370, 377–383

zone of authority, 346, 370

zone transfer

- with BIND, 395–397
- of DNS server, monitoring, 416
- process, 347, 373
- securing against footprinting, 405
- security, 366
- summary of, 462–463
- troubleshooting, 456
- zone replication, planning for, 377–383



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®