

Microsoft®

EXAM 70-642

Covers Windows
Server 2008

R2

Configuring Windows Server® 2008 Network Infrastructure



Tony Northrup
J.C. Mackin

SECOND EDITION

Training Kit

Exam 70-642: Windows Server 2008 Network Infrastructure, Configuring

OBJECTIVE	CHAPTER	LESSON
1. CONFIGURING ADDRESSING AND SERVICES		
1.1 Configure IPv4 and IPv6 addressing.	Chapter 1	Lessons 2 and 3
1.2 Configure Dynamic Host Configuration Protocol (DHCP).	Chapter 4	Lessons 1 and 2
1.3 Configure routing.	Chapter 5	Lesson 1
1.4 Configure Windows Firewall with Advanced Security.	Chapter 6 Chapter 8	Lesson 1 Lesson 1
2. CONFIGURING NAMES RESOLUTION		
2.1 Configure a Domain Name System (DNS) server.	Chapter 2	Lesson 2
2.2 Configure DNS zones.	Chapter 3	Lessons 1 and 3
2.3 Configure DNS records.	Chapter 3	Lesson 1
2.4 Configure DNS replication.	Chapter 3	Lesson 2
2.5 Configure name resolution for client computers.	Chapter 2	Lesson 3
3. CONFIGURING NETWORK ACCESS		
3.1 Configure remote access.	Chapter 7	Lessons 2 and 3
3.2 Configure Network Access Protection (NAP).	Chapter 8	Lesson 2
3.3 Configure DirectAccess.	Chapter 7	Lesson 4
3.4 Configure Network Policy Server (NPS).	Chapter 7	Lesson 1
4. CONFIGURING FILE AND PRINT SERVICES		
4.1 Configure a file server.	Chapter 11	Lessons 1 and 2
4.2 Configure Distributed File System (DFS).	Chapter 11	Lesson 2
4.3 Configure backup and restore.	Chapter 11	Lesson 3
4.4 Manage file server resources.	Chapter 11	Lesson 2
4.5 Configure and monitor print services.	Chapter 12	Lesson 1
5. MONITORING AND MANAGING A NETWORK INFRASTRUCTURE		
5.1 Configure Windows Server Update Services (WSUS) server settings.	Chapter 9	Lessons 1 and 2
5.2 Configure performance monitoring.	Chapter 10	Lessons 1 and 2
5.3 Configure event logs.	Chapter 10	Lesson 1
5.4 Gather network data.	Chapter 10 Chapter 6	Lessons 1, 2, and 3 Lesson 1

Exam Objectives The exam objectives listed here are current as of this book's publication date. Exam objectives are subject to change at any time without prior notice and at Microsoft's sole discretion. Please visit the Microsoft Learning website for the most current listing of exam objectives: <http://www.microsoft.com/learning/en-us/exam.aspx?ID=70-642&locale=en-us>.

**Self-Paced Training Kit
(Exam 70-642):
Configuring
Windows Server® 2008
Network Infrastructure
(2nd Edition)**

Tony Northrup
J.C. Mackin

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

Copyright © 2011 by Tony Northrup and J.C. Mackin

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

Library of Congress Control Number: 2011924391

ISBN: 978-0-7356-5160-9

9 10 11 12 13 14 15 16 17 QG 8 7 6 5 4 3

Printed and bound in the United States of America.

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

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

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

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

Acquisitions Editor: Jeff Koch

Developmental Editor: Karen Szall

Project Editor: Carol Dillingham

Editorial Production: Online Training Solutions, Inc.

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

Copy Editor: Victoria Thulman; Online Training Solutions, Inc.

Indexer: Jan Bednarczuk; Online Training Solutions, Inc.

Cover: Twist Creative • Seattle

Contents

Introduction	xvii
System Requirements	xvii
Using the CD	xix
Acknowledgments	xxii
Support & Feedback	xxii
Preparing for the Exam	xxiv

CHAPTER 1 Understanding and Configuring TCP/IP **1**

Before You Begin	1
Lesson 1: Introducing Windows Networking	2
What Are Network Layers?	2
Exploring the Layers of the TCP/IP Networking Model	5
Configuring Networking Properties in Windows Server 2008 R2	14
Lesson Summary	39
Lesson Review	39
Lesson 2: Understanding IPv4 Addressing	40
The Structure of IPv4 Addresses	40
Understanding Routing and Default Gateways	50
Understanding IPv4 Address Ranges	51
What Is Subnetting?	61
Advantages of Subnetting	63
The Subnet ID	65
Creating Equally Sized Subnets	66
Using Variable-Length Subnet Masks	67

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

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

www.microsoft.com/learning/booksurvey/

Enumerating Subnets in an Address Space	68
Verifying Subnet Ownership and Configuration	76
Lesson Summary	84
Lesson Review	84
Lesson 3: Understanding IPv6 Addressing	86
Introducing IPv6 Addresses	86
Understanding IPv6 Address Types	87
IPv6 Transition Technologies	92
Working with IPv6 Subnets	96
Lesson Summary	103
Lesson Review	104
Chapter Review	105
Chapter Summary	105
Key Terms	106
Case Scenarios	106
Suggested Practices	107
Take a Practice Test	107

CHAPTER 2 Configuring Name Resolution 109

Before You Begin.	109
Lesson 1: Understanding Name Resolution in Windows Server 2008 Networks	111
Name Resolution Methods in Windows	111
What Is Link Local Multicast Name Resolution?	112
What Is NetBIOS Name Resolution?	116
What Is DNS Name Resolution?	120
DNS Components	122
Understanding How a DNS Query Works	124
Understanding How Caching Works	131
Lesson Summary	134
Lesson Review	135
Lesson 2: Deploying a DNS Server.	136
Deploying a DNS Server on a Domain Controller	137
Deploying a DNS Server on a Stand-Alone or Member Server	139

Deploying a DNS Server on a Server Core Installation of Windows Server 2008 R2	140
Configuring a Caching-Only DNS Server	142
Configuring Server Properties	143
Configuring DNS Socket Pooling	151
Configuring DNS Cache Locking	151
Lesson Summary	156
Lesson Review	156
Lesson 3: Configuring DNS Client Settings.	158
Specifying DNS Servers	158
Specifying a Computer Name and DNS Suffixes	160
Configuring a Suffix Search List	162
Configuring Dynamic Update Settings	165
Viewing and Clearing the DNS Client Cache	168
Lesson Summary	169
Lesson Review	170
Chapter Review	171
Chapter Summary	171
Key Terms	172
Case Scenarios	173
Suggested Practices	174
Take a Practice Test	174

CHAPTER 3 Configuring a DNS Zone Infrastructure 175

Before You Begin.	175
Lesson 1: Creating and Configuring Zones.	176
Creating Zones	177
Examining Built-in Resource Records	185
Creating Resource Records	189
Enabling DNS to Use WINS Resolution	195
Aging and Scavenging	196
Using a GlobalNames Zone	199
Lesson Summary	202
Lesson Review	203

Lesson 2: Configuring Zone Replication, Transfers, and Delegations . . .	204
Configuring Zone Replication for Active Directory– Integrated Zones	205
Using Zone Transfers	210
Understanding Zone Delegations	213
Implementing Stub Zones	215
Lesson Summary	221
Lesson Review	222
Lesson 3: Implementing DNSSEC	223
Understanding Public Key Cryptography in DNSSEC	224
Understanding DNSSEC Name Resolution	227
Configuring DNSSEC	234
Lesson Summary	246
Lesson Review	247
Chapter Review	248
Chapter Summary	248
Key Terms	249
Case Scenarios	250
Suggested Practices	251
Take a Practice Test	251

CHAPTER 4 Creating a DHCP Infrastructure 253

Before You Begin.	253
Lesson 1: Installing a DHCP Server.	254
Understanding DHCP Address Assignment	254
Adding the DHCP Server Role	258
Lesson Summary	268
Lesson Review	269
Lesson 2: Configuring a DHCP Server	270
Performing Post-Installation Tasks	270
Understanding DHCP Options Classes	276
Controlling DHCP Access Through MAC Filtering	279
DHCP Delay Configuration	280

Using the DHCP Split-Scope Configuration Wizard	281
Configuring DHCP to Perform Dynamic DNS Updates for Clients	282
Installing and Configuring DHCP on a Server Core Installation	285
Lesson Summary	286
Lesson Review	287
Chapter Review	288
Chapter Summary	288
Key Terms	289
Case Scenarios	289
Suggested Practice	290
Take a Practice Test	290
CHAPTER 5 Configuring IP Routing	291
Before You Begin	291
Lesson 1: Routing	292
Routing Overview	292
Examining Network Routes	294
Routing Protocols	295
Demand-Dial Routing	299
Static Routing	300
Lesson Summary	310
Lesson Review	310
Chapter Review	311
Chapter Summary	311
Key Terms	311
Case Scenarios	312
Suggested Practices	313
Take a Practice Test	313
CHAPTER 6 Protecting Network Traffic with IPsec	315
Before You Begin	315
Lesson 1: Configuring IPsec	316
What Is IPsec?	317
Using IPsec in Tunnel Mode	322

Authentication Methods for IPsec	323
Assigning a Predefined IPsec Policy	324
Creating a New IPsec Policy	325
Creating and Configuring a Connection Security Rule	330
Lesson Summary	342
Lesson Review	343
Chapter Review	344
Chapter Summary	344
Key Terms	344
Case Scenario	345
Suggested Practices	345
Take a Practice Test	346

CHAPTER 7 Connecting to Networks 347

Before You Begin.	348
Lesson 1: Configuring Network Policy Server	349
Wireless Security Standards	350
Infrastructure and Ad Hoc Wireless Networks	352
Configuring the Public Key Infrastructure	352
Authenticating Wireless Networks by Using Windows Server 2008 R2	353
Connecting to Wireless Networks	361
Deploying Wireless Networks with WPA-EAP	362
Wired Network Security	363
Using NPS Templates	365
Lesson Summary	370
Lesson Review	371
Lesson 2: Configuring Network Address Translation	372
Network Address Translation Concepts	372
Configuring Internet Connection Sharing	374
Configuring Network Address Translation by Using Routing And Remote Access	376
Troubleshooting Network Address Translation	378
Lesson Summary	380
Lesson Review	381

Lesson 3: Connecting to Remote Networks	382
Remote Access Overview	382
Configuring Dial-Up Connections	385
Configuring VPN Connections	391
Troubleshooting VPN Connection Problems	395
Configuring Connection Restrictions	395
Testing Connectivity	397
Lesson Summary	404
Lesson Review	404
Lesson 4: Configuring DirectAccess	405
DirectAccess Connection Types	407
Using DirectAccess on IPv4 Networks	408
DirectAccess and Name Resolution	409
The Network Location Server	410
DirectAccess Requirements	410
DirectAccess Limitations	412
Firewall Configuration	412
Running the DirectAccess Setup Wizard	413
Lesson Summary	422
Lesson Review	422
Chapter Review	423
Chapter Summary	423
Key Terms	424
Case Scenarios	424
Suggested Practices	426
Take a Practice Test	427

CHAPTER 8 Configuring Windows Firewall and Network Access Protection 429

Before You Begin	429
Lesson 1: Configuring Windows Firewall	430
Why Firewalls Are Important	431
Firewall Profiles	431
Filtering Inbound Traffic	432

Filtering Outbound Traffic	434
Configuring Scope	435
Authorizing Connections	436
Configuring Firewall Settings with Group Policy	438
Enabling Logging for Windows Firewall	439
Identifying Network Communications	439
Lesson Summary	442
Lesson Review	443
Lesson 2: Configuring Network Access Protection	444
Network Access Protection Concepts	445
Planning a NAP Deployment	450
Installing and Configuring the Network Policy Server	450
Configuring NAP Enforcement	453
Configuring NAP Components	463
NAP Logging	474
Lesson Summary	480
Lesson Review	481
Chapter Review	482
Chapter Summary	482
Key Terms	482
Case Scenarios	483
Suggested Practices	484
Take a Practice Test	485

CHAPTER 9 Managing Software Updates 487

Before You Begin.	488
Lesson 1: Understanding Windows Server Update Services	489
WSUS Overview	489
Windows Update Client	490
WSUS Architecture	492
WSUS Requirements	494
Planning the WSUS Installation	495
Auditing Updates	496
Lesson Summary	497
Lesson Review	498

Lesson 2: Using Windows Server Update Services	499
Installing Windows Server Update Services	499
Configuring Windows Server Update Services	500
Troubleshooting Problems Installing Updates	510
Removing Updates	513
Lesson Summary	516
Lesson Review	516
Chapter Review	517
Chapter Summary	518
Key Terms	518
Case Scenarios	518
Suggested Practice	520
Take a Practice Test	520

CHAPTER 10 Monitoring Computers 521

Before You Begin	521
Lesson 1: Monitoring Events	523
Using Event Viewer	523
Automatically Responding to Events	525
Configuring Event Forwarding	526
Lesson Summary	537
Lesson Review	538
Lesson 2: Monitoring Performance and Reliability	539
Using Performance Monitor	539
Using Reliability Monitor	542
Using Data Collector Sets	543
Configuring Virtual Memory	549
Lesson Summary	552
Lesson Review	553
Lesson 3: Using Network Monitor and Simple Network Management Protocol	554
Installing Network Monitor	554
Capturing and Analyzing Network Communications	555

Configuring SNMP	561
Lesson Summary	564
Lesson Review	564
Chapter Review	565
Chapter Summary	566
Key Terms	566
Case Scenarios	566
Suggested Practices	567
Take a Practice Test	568

CHAPTER 11 Managing Files 569

Before You Begin	569
Lesson 1: Managing File Security	570
NTFS File Permissions	571
Encrypting File System	573
BitLocker	578
Lesson Summary	584
Lesson Review	584
Lesson 2: Sharing Folders	585
Installing the File Services Server Role	586
Quotas	587
Folder Sharing	592
Classification Management	596
Distributed File System	599
Offline Files	604
BranchCache	606
Lesson Summary	613
Lesson Review	614
Lesson 3: Backing Up and Restoring Files	615
Shadow Copies	616
Windows Server Backup	617
Lesson Summary	626
Lesson Review	627

Chapter Review	627
Chapter Summary	628
Key Terms	628
Case Scenarios	628
Suggested Practices	629
Take a Practice Test	630

CHAPTER 12 Managing Printers 631

Before You Begin	631
Lesson 1: Managing Printers	632
Installing the Print And Document Services Server Role	633
Installing Printers	634
Sharing Printers	638
Configuring Print Server and Printer Permissions	640
Adding Printer Drivers	640
Configuring Printer Pooling	642
Configuring Printer Priorities	643
Managing Internet Printing	643
Generating Notifications	644
Deploying Printers with Group Policy	646
Migrating Printers	647
Managing Printers from a Command Prompt or Script	648
Monitoring Printers	650
Lesson Summary	653
Lesson Review	654
Chapter Review	655
Chapter Summary	655
Key Terms	655
Case Scenario	655
Suggested Practices	656
Take a Practice Test	656

Answer	657
Glossary	693
Index	697
About the Authors	725

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

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

www.microsoft.com/learning/booksurvey/

Introduction

This training kit is designed for information technology (IT) professionals who work in the complex computing environment of medium-sized to large companies and who also plan to take Exam 70-642. We assume that before you begin using this training kit, you have a basic understanding of Windows server operating systems and common Internet technologies.

NOTE WINDOWS SERVER 2008 CERTIFICATION

Exam 70-642 is one of three required exams for MCSA: Windows Server 2008 certification. For a limited time, it is also valid for the MCTS certification, which will be retired. Please visit the Microsoft Learning website for the most current information about Microsoft certifications: <http://www.microsoft.com/learning/>

The material covered in this training kit and on the 70-642 exam relates to fundamental networking features such as addressing, name resolution, remote access, and printing. The topics in this training kit cover what you need to know for the exam as described in the Preparation Guide for the 70-642 exam, which is available at <http://www.microsoft.com/learning/en/us/exam.aspx?ID=70-642>.

By using this training kit, you will learn how to do the following:

- Configure IP addressing, routing, and IPsec.
- Configure name resolution by using Domain Name System (DNS).
- Configure remote and wireless network access.
- Configure Network Access Protection (NAP).
- Configure file and print services.
- Monitor and manage a network infrastructure.

Refer to the objective mapping page in the front of this book to see where in the book each exam objective is covered.

System Requirements

The following are the minimum system requirements your computer needs to meet to complete the practice exercises in this book and to run the companion CD.

Hardware Requirements

We recommend that you use a single physical computer and virtualization software to perform the exercises in this training kit. The physical computer should meet the following requirements:

- x64 processor.

- If you are using Hyper-V for virtualization software, the processor must support hardware-assisted virtualization, No eXecute (NX) bit technology, and data execution prevention (DEP).
- 2 GB RAM (8 GB is recommended).
- 100 GB of hard disk space (25 GB for each of three virtual machines plus 25 GB for the base system).

Software Requirements

The following software is required to complete the practice exercises:

- Windows Server 2008 R2. You can download an evaluation edition of Windows Server 2008 R2 at the Microsoft Download Center at <http://www.microsoft.com/downloads>.
- A web browser such as Windows Internet Explorer 7, Internet Explorer 8, or Internet Explorer 9.
- An application that can display PDF files, such as Adobe Acrobat Reader, which can be downloaded from <http://www.adobe.com/reader>.

Lab Setup Instructions

Most of the exercises in this training kit require two computers or virtual machines running Windows Server 2008 R2. (The exercises in Chapter 6, "Protecting Network Traffic with IPsec," several Lesson 4 exercises in Chapter 7, "Connecting to Networks," and Lesson 2 of Chapter 8, "Configuring Windows Firewall and Network Access Protection," require a third such computer or virtual machine.) All lab computers must be physically connected to the same network for most lessons. However, some lessons will describe different network configurations. We recommend that you use an isolated network that is not part of your production network to perform the practice exercises in this book.

To minimize the time and expense of configuring physical computers and networks, we recommend that you use virtual machines for the lab computers. To run computers as virtual machines within Windows, you can use Hyper-V or third-party virtual machine software such as the free VirtualBox. Both of these options allow you to run 64-bit guest operating systems in a virtual environment, and this feature is required to support Windows Server 2008 R2, which is 64-bit only. (Note that neither Virtual PC nor Virtual Server support 64-bit guests.) For more information about Hyper-V, visit <http://www.microsoft.com/hyperv>. To download VirtualBox, visit <http://www.virtualbox.org>.

Using a virtual environment is the simplest way to prepare the computers for this training kit. To isolate the lab computers within a single network, configure the settings in each virtual machine so that the network adapter is assigned to a private or an internal network. (Note that virtual network adapters are not assigned to such private or internal networks by default in either Hyper-V or VirtualBox.) In addition, some exercises need Internet access, which will require you to connect the network adapter to an external network. You can perform these exercises by temporarily connecting the network adapter to an external network, or you can perform them on another computer with Internet access.

Preparing the Windows Server 2008 R2 Computers

Perform the following steps to prepare the first Windows Server 2008 computer for the exercises in this training kit.

On the three lab computers, perform a default installation of Windows Server 2008 R2. Do not add any roles or adjust the networking settings. In Control Panel, use System to specify the computer name of the first computer as **Dcsrv1**, the second computer as **Boston**, and the third computer as **Binghamton**.

If you are using virtual machines, you should save a snapshot of the virtual machine after setup is complete so that you can quickly return the computer to that state.

NOTE TAKE SNAPSHOTS AFTER EACH EXERCISE

Virtual machine software allows you to take a *snapshot* of a virtual machine, which is the complete state of a virtual machine at any point in time. After each exercise, you should take a snapshot of any computers on which changes have been made. After Dcsrv1 is promoted to a domain controller, be sure to always take a snapshot of this virtual machine even when exercises are performed on another computer. (Changes made to member servers often modify settings on the domain controller.)

Using the CD

The companion CD included with this training kit contains the following:

- **Practice tests** You can reinforce your understanding of how to configure Windows Server 2008 R2 network infrastructure by using electronic practice tests you customize to meet your needs from the pool of Lesson Review questions in this book. Or you can practice for the 70-642 certification exam by using tests created from a pool of about 200 realistic exam questions, which give you many practice exams to ensure that you are prepared.
- **Webcast** To supplement your learning, the CD includes a webcast about IPsec.
- **eBook** An electronic version (eBook) of this book is included for when you do not want to carry the printed book with you. The eBook is in Portable Document Format (PDF), and you can view it by using Adobe Acrobat or Adobe Reader.

Companion Content for Digital Book Readers: If you bought a digital edition of this book, you can enjoy select content from the print edition's companion CD. Visit <http://go.microsoft.com/fwlink/?Linkid=215050> to get your downloadable content. This content is always up-to-date and available to all readers.

How to Install the Practice Tests

To install the practice test software from the companion CD to your hard disk, do the following:

1. Insert the companion CD into your CD drive, and accept the license agreement. A CD menu appears.

NOTE IF THE CD MENU DOES NOT APPEAR

If the CD menu or the license agreement does not appear, AutoRun might be disabled on your computer. Refer to the *Readme.txt* file on the CD-ROM for alternate installation instructions.

2. Click Practice Tests, and follow the instructions on the screen.

How to Use the Practice Tests

To start the practice test software, follow these steps:

1. Click Start\All Programs\Microsoft Press Training Kit Exam Prep. A window appears that shows all the Microsoft Press training kit exam prep suites installed on your computer.
2. Double-click the lesson review or practice test you want to use.

NOTE LESSON REVIEWS VS. PRACTICE TESTS

Select the (70-642) Configuring Windows Server 2008 Network Infrastructure (2nd Edition) lesson review to use the questions from the "Lesson Review" sections of this book. Select the (70-642) Configuring Windows Server 2008 Network Infrastructure (2nd Edition) practice test to use a pool of about 200 questions similar to those that appear on the 70-642 certification exam.

Lesson Review Options

When you start a lesson review, the Custom Mode dialog box appears so that you can configure your test. You can click OK to accept the defaults, or you can customize the number of questions you want, how the practice test software works, which exam objectives you want the questions to relate to, and whether you want your lesson review to be timed. If you are retaking a test, you can select whether you want to see all the questions again or only the questions you missed or did not answer.

After you click OK, your lesson review starts. The following list explains the main options you have for taking the test:

- To take the test, answer the questions and use the Next, Previous, and Go To buttons to move from question to question.
- After you answer an individual question, if you want to see which answers are correct—along with an explanation of each correct answer—click Explanation.
- If you prefer to wait until the end of the test to see how you did, answer all the questions and then click Score Test. You will see a summary of the exam objectives you chose and the percentage of questions you got right overall and per objective. You can print a copy of your test, review your answers, or retake the test.

Practice Test Options

When you start a practice test, you choose whether to take the test in Certification Mode, Study Mode, or Custom Mode:

- **Certification Mode** Closely resembles the experience of taking a certification exam. The test has a set number of questions. It is timed, and you cannot pause and restart the timer.
- **Study Mode** Creates an untimed test in which you can review the correct answers and the explanations after you answer each question.
- **Custom Mode** Gives you full control over the test options so that you can customize them as you like.

In all modes, the user interface you see when you are taking the test is basically the same but with different options enabled or disabled depending on the mode. The main options are discussed in the previous section, “Lesson Review Options.”

When you review your answer to an individual practice test question, a “References” section is provided that lists where in the training kit you can find the information that relates to that question and provides links to other sources of information. After you click Test Results to score your entire practice test, you can click the Learning Plan tab to see a list of references for every objective.

How to Uninstall the Practice Tests

To uninstall the practice test software for a training kit, use Add Or Remove Programs option (Windows XP) or the Program And Features option (Windows 7 and Windows Server 2008 R2) in Windows Control Panel.

Acknowledgments

This book was put together by a team of respected professionals, and we, the authors, would like to thank them each for the great job they did.

At Microsoft, Jeff Koch worked out our contracts as the acquisitions editor, Karen Szall was our developmental editor, and Carol Dillingham was our project editor.

Kathy Krause of Online Training Solutions, Inc., managed the editorial and production teams. Victoria Thulman, our copy editor, was responsible for making sure the book was readable and consistent, and Jaime Odell provided additional proofreading.

Bob Dean provided a technical review to help make the book as accurate as possible. Jan Bednarczuk created the index that you'll find at the back of the book.

Many people helped with this book, even though they weren't formally part of the team.

Tony Northrup would like to thank his friends, especially Brian and Melissa Rheaume, Jose and Kristin Gonzales, Chelsea and Madelyn Knowles, Eddie and Christine Mercado, Papa Jose, and Nana Lucy.

J.C. Mackin would like to thank his friends and family for always being so supportive.

It makes a huge difference when you consider the people you work with to be friends. Having a great team not only improves the quality of the book, it makes it a more pleasant experience. Writing this book was most enjoyable, and we hope we get the chance to work with everyone in the future.

Support & Feedback

The following sections provide information on errata, book support, feedback, and contact information.

Preparing for the Exam

Microsoft certification exams are a great way to build your resume and let the world know about your level of expertise. Certification exams validate your on-the-job experience and product knowledge. Although there is no substitute for on-the-job experience, preparation through study and hands-on practice can help you prepare for the exam. We recommend that you augment your exam preparation plan by using a combination of available study materials and courses. For example, you might use the Training Kit and another study guide for your “at home” preparation, and take a Microsoft Official Curriculum course for the classroom experience. Choose the combination that you think works best for you.

Configuring a DNS Zone Infrastructure

Deploying a DNS server is a fairly simple procedure, especially on a domain controller. But to manage and troubleshoot DNS, you need to understand zones in more detail.

Zones are the databases in which DNS data is stored. A DNS zone infrastructure essentially consists of the various servers and hosted zones that communicate with one another in a way that ensures consistent name resolution. This chapter introduces you to the types of zones that make up a DNS infrastructure, the options for zone replications and transfers among them, and the configurable settings within zones that you need to understand in order to manage DNS effectively on your network.

Exam objectives in this chapter:

- Configure DNS zones.
- Configure DNS records.
- Configure DNS replication.

Lessons in this chapter:

- Lesson 1: Creating and Configuring Zones **176**
- Lesson 2: Configuring Zone Replication, Transfers, and Delegations **204**
- Lesson 3: Implementing DNSSEC **223**

Before You Begin

To complete the lessons in this chapter, you must have the following:

- Two networked computers running Windows Server 2008 R2.
- The first computer must be a domain controller named Dcsrv1 in a domain named nwtraders.msft. Dcsrv1 must be assigned the static address 192.168.0.1/24 with the DNS server specified as the same address. Dcsrv1 includes the server roles Active Directory Domain Services and DNS Server.
- The second computer must be named Boston.nwtraders.msft and must be assigned the address 192.168.0.2/24. Its DNS server must be specified as 192.168.0.1. Finally, Boston must be joined to the Nwtraders.msft domain.



REAL WORLD

J.C. Mackin

DNS Manager is the main administration tool for DNS servers, but if you need to manage DNS for your job, it's a good idea to become familiar with some other DNS tools as well. Of all the alternate tools available, the `Dnscmd` command-line tool is the most important and the most powerful. By typing `dnscmd` at a command prompt, you can see all 40 or so of its subcommands. Some of the most important of these include `dnscmd /clearcache`, which clears the server cache; `dnscmd /enumdirectorypartitions`, which shows the application directory partitions available on the local server; and `dnscmd /info` (which provides a basic overview of the DNS server configuration).

If your network includes Active Directory–integrated zones, you should also review tools for managing Active Directory replication. If you want to test replication on a domain controller, type `dcdiag /test:replications`. If you want to show replication partners, type `repadmin /showrepl`. Finally, if you want to force replication with another domain controller, use the Active Directory Sites and Services console to browse to the NTDS settings beneath your server, right-click the connection object in the details pane, and click Replicate Now.

Lesson 1: Creating and Configuring Zones

A *zone* is a database that contains authoritative information about a portion of the DNS namespace. When you install a DNS server with a domain controller, the DNS zone used to support the Active Directory domain is created automatically. However, if you install a DNS server at any other time, either on a domain controller, domain member server, or stand-alone server, you have to create and configure zones manually.

This lesson describes the steps required to create and configure a zone, as well as the underlying concepts you need to understand to configure a zone properly.

After this lesson, you will be able to:

- Create and configure DNS zones.
- Create and configure resource records.

Estimated lesson time: 120 minutes

Creating Zones

A *DNS zone* is a database containing records that associate names with addresses for a defined portion of a DNS namespace. Although a DNS server can use cached information from other servers to answer queries for names, it is only through a locally hosted zone that a DNS server can answer queries authoritatively. For any portion of a DNS namespace represented by a domain name such as “proseware.com,” there can be only one authoritative source of zone data.

To create a new zone on a DNS server, you can use the New Zone Wizard in DNS Manager. To launch this wizard, right-click the server icon in the DNS Manager console tree, and then choose New Zone, as shown in Figure 3-1.

The New Zone Wizard includes the following configuration pages:

- Zone Type
- Active Directory Zone Replication Scope
- Forward or Reverse Lookup Zone
- Zone Name
- Dynamic Update

The sections that follow describe the configuration concepts related to these five wizard pages.

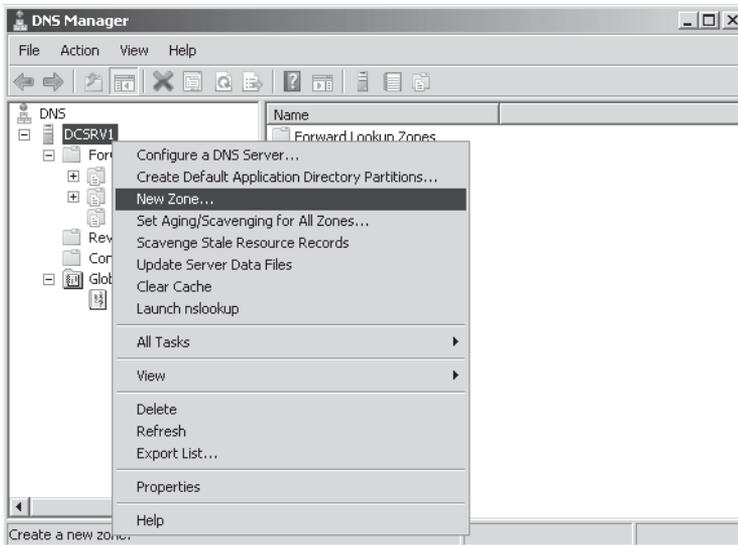


FIGURE 3-1 Creating a new zone

Choosing a Zone Type

The Zone Type page of the New Zone Wizard, shown in Figure 3-2, enables you to create your choice of a primary zone, a secondary zone, or a stub zone. If you are creating a primary or stub zone on a domain controller, you also have the option to store zone data in Active Directory.

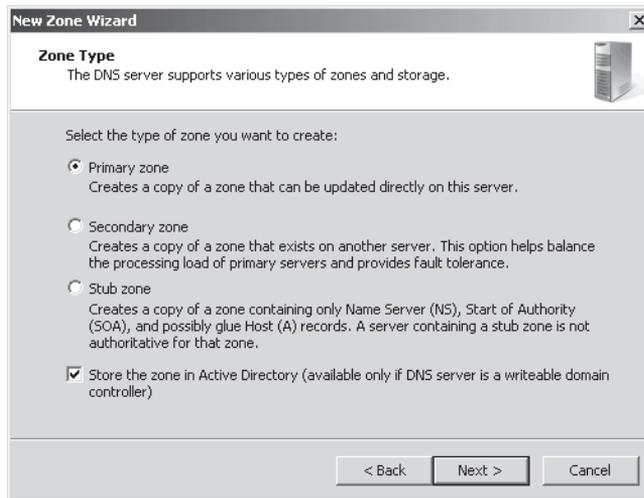


FIGURE 3-2 Choosing a zone type

PRIMARY ZONES

A *primary zone* is the main type of DNS zone. A primary zone provides original read-write source data that allows the local DNS server to answer DNS queries authoritatively about a portion of a DNS namespace.

When the local DNS server hosts a primary zone, the DNS server is the primary source for information about this zone, and the server stores the master copy of zone data in a local file or in Active Directory Domain Services (AD DS). When the zone is stored in a file instead of Active Directory, by default the primary zone file is named *zone_name.dns*, and this file is located in the %systemroot%\System32\Dns folder on the server.

SECONDARY ZONES

A *secondary zone* provides an authoritative, read-only copy of a primary zone or another secondary zone. Secondary zones provide a means to offload DNS query traffic in areas of the network where a zone is heavily queried and used. Additionally, if the zone server hosting a primary zone is unavailable, a secondary zone can provide name resolution for the namespace until the primary server becomes available again.

The source zones from which secondary zones acquire their information are called *masters*, and the data copy procedures through which this information is regularly updated are called *zone transfers*. A master can be a primary zone or other secondary zone. You can specify the master of a secondary zone when the secondary zone is created through the New Zone Wizard. Because a secondary zone is merely a copy of a primary zone that is hosted on another server, it cannot be stored in AD DS.

STUB ZONES

A *stub zone* is similar to a secondary zone, but it contains only those resource records necessary to identify the authoritative DNS servers for the master zone. Stub zones are often used to enable a parent zone like `proseware.com` to keep an updated list of the name servers available in a delegated child zone, such as `east.proseware.com`. They can also be used to improve name resolution and simplify DNS administration.

STORING THE ZONE IN ACTIVE DIRECTORY

When you create a new primary or stub zone on a domain controller, the Zone Type page gives you the option to store the zone in Active Directory. In Active Directory–integrated zones, zone data is automatically replicated through Active Directory in a manner determined by the settings you choose on the Active Directory Zone Replication Scope page. In most cases, this option eliminates the need to configure zone transfers to secondary servers.

Integrating your DNS zone with Active Directory has several advantages. First, because Active Directory performs zone replication, you do not need to configure a separate mechanism for DNS zone transfers between primary and secondary servers. Fault tolerance, along with improved performance from the availability of multiple read/write primary servers, is automatically supplied by the presence of multimaster replication on your network. Second, Active Directory allows for single properties of resource records to be updated and replicated among DNS servers. Avoiding the transfer of many and complete resource records decreases the load on network resources during zone transfers. Finally, Active Directory–integrated zones also provide the optional benefit of requiring security for dynamic updates, an option you can configure on the Dynamic Update page.

NOTE READ-ONLY DOMAIN CONTROLLERS AND ACTIVE DIRECTORY–INTEGRATED ZONES

For traditional domain controllers, the copy of the zone is a read-write copy. For read-only domain controllers (RODCs), the copy of the zone will be read-only.

STANDARD ZONES

By default, on the Zone Type page, the option to store the zone in Active Directory is selected when you are creating the zone on a domain controller. However, you can clear this check box and instead create what is called a standard zone. A standard zone is also the only option for a new zone when you are creating the zone on a server that is not a domain controller; in this case the check box on this page cannot be selected.

As opposed to an Active Directory–integrated zone, a *standard zone* stores its data in a text file on the local DNS server. Also unlike Active Directory–integrated zones, with standard zones, you can configure only a single read-write (primary) copy of zone data. All other copies of the zone (secondary zones) are read-only.

The standard zone model implies a single point of failure for the writable version of the zone. If the primary zone is unavailable to the network, no changes to the zone can be made. However, queries for names in the zone can continue uninterrupted as long as secondary zones are available.

Choosing an Active Directory Zone Replication Scope

On the Active Directory Zone Replication Scope page of the New Scope Wizard, you can choose which domain controllers in your network will store the zone. This page, shown in Figure 3-3, appears only when you have configured the zone to be stored in Active Directory. Note that the choice of where you store the zone determines the domain controllers among which the zone data will be replicated.

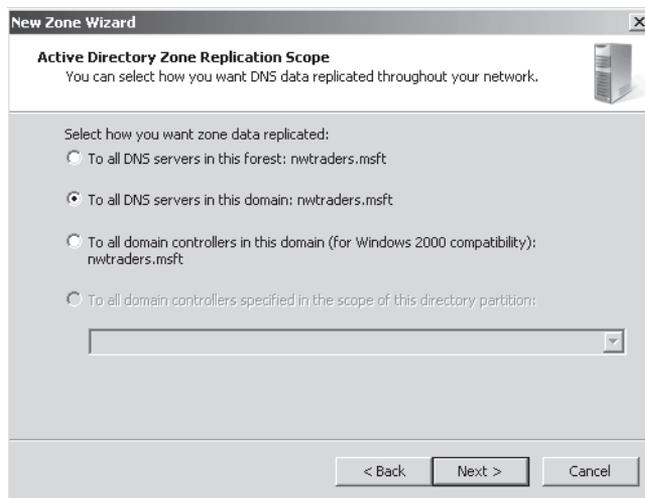


FIGURE 3-3 Choosing the domain controllers to store the zone

You have four choices:

- Store the zone in all domain controllers that are also DNS servers in the entire Active Directory forest.
- Store the zone in all domain controllers that are also DNS servers in the local Active Directory domain.
- Store the zone in all domain controllers in the local Active Directory domain (used for compatibility with Windows 2000).
- Store the zone in all domain controllers specified in the scope of a custom Active Directory directory partition.

These options are described in more detail in Lesson 2, “Configuring Zone Replication, Transfers, and Delegations.”

Creating a Forward or Reverse Lookup Zone

On the Forward Or Reverse Lookup Zone page of the New Zone Wizard, you determine whether the new zone you are creating should act as a forward or reverse lookup zone. This page is shown in Figure 3-4.

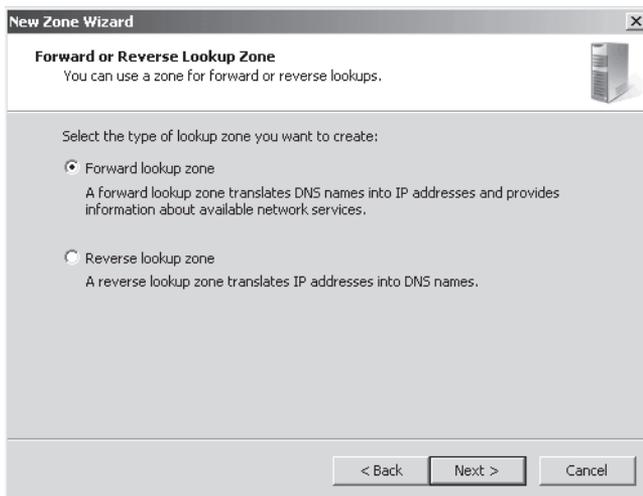


FIGURE 3-4 Choosing a forward or reverse lookup zone

In *forward lookup zones*, DNS servers map fully qualified domain names (FQDNs) to IP addresses. In *reverse lookup zones*, DNS servers map IP addresses to FQDNs. Forward lookup zones thus answer queries to resolve FQDNs to IP addresses, and reverse lookup zones answer queries to resolve IP addresses to FQDNs. Note that forward lookup zones adopt the name of the DNS domain name for whose names you want to provide resolution service, such as “proseware.com.” Reverse lookup zones are named by a reverse order of the

network ID octets in the address space for which you want to provide reverse name resolution service *plus* the final tag "in-addr.arpa." For example, if you want to provide reverse name resolution service for the subnet 192.168.1.0/24, the name of the reverse lookup zone will be "1.168.192.in-addr.arpa." Within a forward lookup zone, a single database entry or record that maps a host name to an IPv4 address is known as a *host* or an *A* record. (For an IPv6 address, the host record is also known as *AAAA*, or a "quad A" record.) In a reverse lookup zone, a single database entry that maps an address host ID to a host name is known as *pointer* or *PTR* record.

A forward lookup zone is illustrated in Figure 3-5, and a reverse lookup zone is illustrated in Figure 3-6.

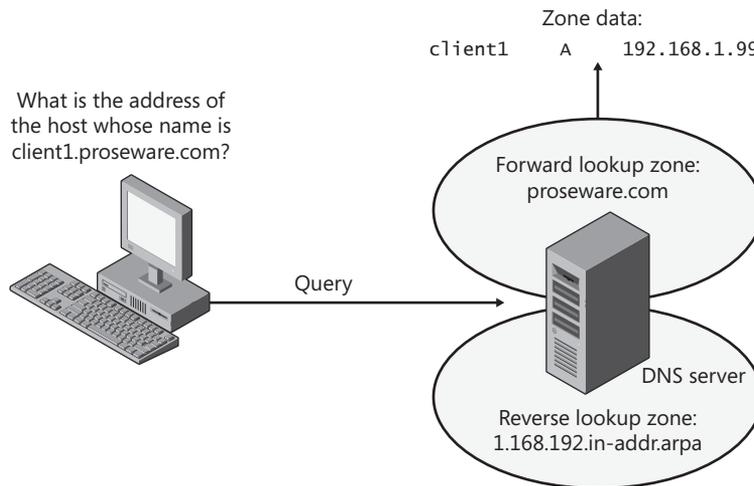


FIGURE 3-5 A forward lookup zone

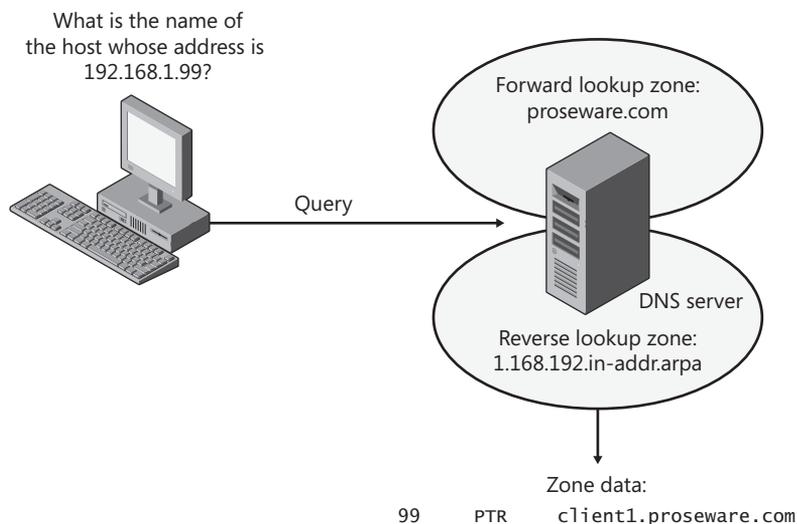


FIGURE 3-6 A reverse lookup zone

NOTE THE CONFIGURE A DNS SERVER WIZARD

To create forward and reverse lookup zones at one time, you can use the Configure A DNS Server Wizard. To open this wizard, right-click the server icon in the DNS Manager console tree, and then choose Configure A DNS Server.

Choosing a Zone Name

The Zone Name page of the New Zone Wizard enables you to choose a name for the forward lookup zone you are creating. (Reverse lookup zones have specific names corresponding to the IP address range for which they are authoritative.) The Zone Name page is shown in Figure 3-7.

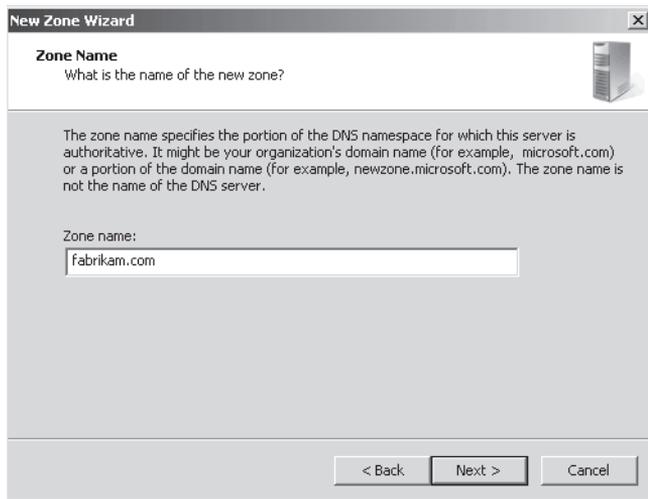


FIGURE 3-7 Choosing a zone name

In general, if the zone you are creating is going to be providing name resolution for an Active Directory domain, you want the zone to match the name of that Active Directory domain. For example, if your organization includes two Active Directory domains named `proseware.com` and `east.proseware.com`, your name resolution infrastructure should include two zones with names that match those Active Directory domains.

If you are creating a zone for a DNS namespace outside of an Active Directory environment, you should supply the name of your organization's Internet domain name, such as `fabrikam.com`.

NOTE ADDING A DNS SERVER TO A DOMAIN CONTROLLER

If you want to add a DNS server to an existing domain controller, you normally want to add a copy of the primary zone providing name resolution for the local Active Directory domain. To achieve this, create a zone whose name corresponds to the name of the existing zone in the local Active Directory domain, and the new zone will be populated with data from other DNS servers in the domain.

Configuring Dynamic Update Settings

DNS client computers can register and dynamically update their resource records with a DNS server. By default, DNS clients that are configured with static IP addresses attempt to update host (A or AAAA) and pointer (PTR) records, and DNS clients that are DHCP clients attempt to update only host records. In a workgroup environment, the DHCP server updates the pointer record on behalf of the DHCP client whenever the IP configuration is renewed.

For dynamic DNS updates to succeed, the zone in which the client attempts to register or update a record must be configured to accept dynamic updates. Two types of dynamic updates are allowed:

- **Secure updates** Allow registrations only from Active Directory domain member computers and updates only from the same computer that originally performed the registration
- **Nonsecure updates** Allow updates from any computer

The Dynamic Update page of the New Zone Wizard enables you to specify whether the zone you are creating should accept secure, nonsecure, or no dynamic updates. The Dynamic Update page is shown in Figure 3-8.

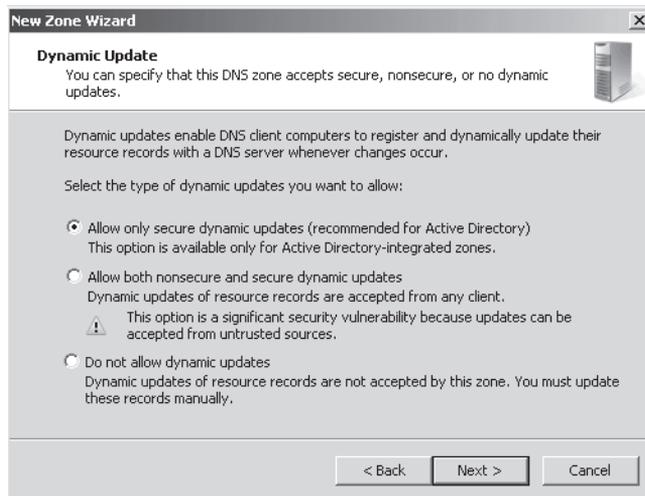


FIGURE 3-8 Configuring dynamic updates on a zone

DYNAMIC UPDATES AND DNS RECORD SECURITY

If you look at the security properties of a resource record, you can see that various users and groups are assigned permissions to the record, just as with any resource in Windows. These security permissions are used for secure dynamic updates. When only secure dynamic updates are allowed in a zone, the user listed as the owner of the resource record (in the advanced security settings) is the only user that can update that record.

The owner of a resource in Windows by default is the user who created that resource. For this reason, when a computer first registers in DNS by creating an A record, that computer becomes the owner of the record.

NOTE USER ACCOUNTS FOR COMPUTERS IN AD DS

Every computer in AD DS gets a user account corresponding to its computer name plus the "\$" symbol, such as Client1\$ or Server1\$.



EXAM TIP

To manually force a DNS client to perform a dynamic update, use the `Ipconfig/registerdns` command.



Quick Check

- What are the server requirements for storing a zone in Active Directory?

Quick Check Answer

- The DNS server needs to be a domain controller.

Examining Built-in Resource Records

When you create a new zone, two types of records required for the zone are automatically created. First, a new zone always includes a Start of Authority (SOA) record that defines basic properties for the zone. All new zones also include at least one NS record signifying the name of the server or servers authoritative for the zone. Figure 3-9 shows a new zone populated by these two records.

The following section describes the functions and features of these two resource records.

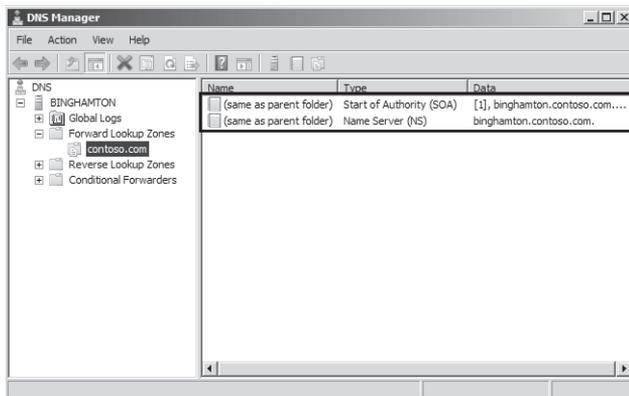


FIGURE 3-9 A new zone always includes at least an SOA and an NS record.

Start of Authority (SOA) Records

When a DNS server loads a zone, it uses the SOA resource record to determine basic and authoritative properties for the zone. These settings also determine how often zone transfers are performed between primary and secondary servers.

If you double-click the SOA record, you open the Start Of Authority (SOA) tab of the zone properties dialog box, shown in Figure 3-10.

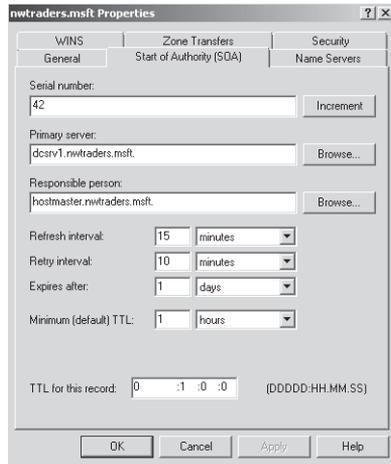


FIGURE 3-10 SOA record settings

On this tab, you can modify the following settings:

- **Serial Number** The Serial Number text box on the Start Of Authority (SOA) tab contains the revision number of the zone file. This number increases each time a resource record changes in the zone or when you manually increment the value in this tab by clicking Increment.

When zones are configured to perform zone transfers to one or more secondary servers, the secondary servers query the master server intermittently for the serial number of the zone. This query is called the *SOA query*. If, through the SOA query, the serial number of the master zone is determined to be equivalent to the serial number stored on the secondary, no transfer is made. However, if the serial number for the zone at the master server is greater than that at the requesting secondary server, the secondary server initiates a transfer.

NOTE FORCING A ZONE TRANSFER ON THE MASTER

When you click the Increment button, you force a zone transfer.

- **Primary Server** The Primary Server text box on the Start Of Authority (SOA) tab contains the full computer name for the primary DNS server of the zone. This name must end with a period.

- **Responsible Person** When this text box is configured, it contains the name of a responsible person (RP) resource record that specifies a domain mailbox name for a zone administrator. The name of the record entered into this field should always end with a period. The name “hostmaster” is used in this field by default.
- **Refresh Interval** The value you configure in the Refresh Interval text box determines how long a secondary DNS server waits before querying the master server for a zone renewal. When the refresh interval expires, the secondary DNS server requests a copy of the current SOA resource record for the zone from its master server source, which then answers this SOA query. The secondary DNS server then compares the serial number of the source server’s current SOA resource record (as indicated in the master’s response) with the serial number of its own local SOA resource record. If they are different, the secondary DNS server requests a zone transfer from the primary DNS server. The default value for this setting is 15 minutes.



EXAM TIP

Increasing the refresh interval decreases zone transfer traffic.

- **Retry Interval** The value you configure in the Retry Interval text box determines how long a secondary server waits before retrying a failed zone transfer. Normally, this time is less than the refresh interval. The default value is 10 minutes.
- **Expires After** The value you configure in the Expires After text box determines the length of time that a secondary server, without any contact with its master server, continues to answer queries from DNS clients. After this time elapses, the data is considered unreliable. The default value is one day.
- **Minimum (Default) TTL** The value you configure in the Minimum (Default) TTL text box determines the default Time to Live (TTL) that is applied to all resource records in the zone. The default value is one hour.

TTL values are not relevant for resource records within their authoritative zones. Instead, the TTL refers to the cache life of a resource record in nonauthoritative servers. A DNS server that has cached a resource record from a previous query discards the record when that record’s TTL has expired.

- **TTL For This Record** The value you configure in this text box determines the TTL of the present SOA resource record. This value overrides the default value setting in the preceding field.

After you create it, an SOA resource record is represented textually in a standard zone file in the manner shown in this example:

```
@ IN SOA computer1.domain1.local. hostmaster.domain1.local. (
  5099      ; serial number
  3600     ; refresh (1 hour)
  600      ; retry (10 mins)
  86400    ; expire (1 day)
  60      ) ; minimum TTL (1 min)
```

Name Server Records

A name server (NS) record specifies a server that is authoritative for a given zone. When you create a zone in Windows Server 2008 or Windows Server 2008 R2, every server hosting a primary copy of an Active Directory–integrated zone will have its own NS record appear in the new zone by default. If you are creating a standard primary zone, an NS record for the local server appears in the zone by default. However, you need to manually add NS records for servers hosting secondary zones on a primary copy of the zone.

Creating an NS record requires a procedure that is different from the one used for creating other resource record types. To add an NS record, double-click any existing NS record in DNS Manager. This step opens the Name Servers tab of the zone properties dialog box, shown in Figure 3-11. On the Name Servers tab, click the Add button to add the FQDN and IP address of the server hosting the secondary zone of the local primary zone. When you click OK after adding the new server, a new NS record pointing to that server appears in DNS Manager.

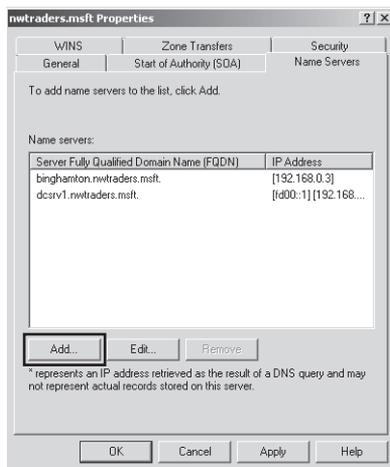


FIGURE 3-11 Adding an NS record to specify a server hosting a secondary zone

NOTE ENABLING TRANSFERS TO SECONDARY ZONES

A secondary zone will not be recognized as a valid name server until it contains a valid copy of zone data. For the secondary zone to obtain this data, you must first enable zone transfers to that server by using the Zone Transfers tab in the Zone Properties dialog box. This tab is discussed in more detail in Lesson 2.

After you create the record, a line such as the following appears in the standard zone file:

```
@ NS dns1.lucernepublishing.com.
```

In this record, the “@” symbol represents the zone defined by the SOA record in the same zone file. The complete entry, then, effectively maps the lucernepublishing.com domain to a DNS server named dns1.lucernepublishing.com.

Creating Resource Records

Beyond the SOA and NS records, some other resource records are also created automatically. For example, if you choose to install a new DNS server when promoting a server to a domain controller, many SRV records for AD DS services are automatically created in the locally hosted zone. In addition, through dynamic updates, many DNS clients automatically register host (A or AAAA) and pointer (PTR) records in a zone by default.

Even though many resource records are created automatically, in a production environment, you usually need to create some resource records manually as well. Such records might include mail exchanger (MX) records for mail servers, alias (CNAME) records for web servers or application servers, and host records for servers or clients that cannot perform their own updates.

To add a resource record for a zone manually, right-click the zone icon in the DNS Manager console, and then choose the type of resource record you want to create from the shortcut menu. Figure 3-12 demonstrates the creation of a new MX record.

After you make your selection from the shortcut menu, a new dialog box appears in which you can specify the name of the record and the computer associated with it. Figure 3-13 shows the New Resource Record dialog box that appears for the creation of a new MX record. Note that only host records associate the name of a computer with the actual IP address of the computer. Most record types associate the name of a service or alias with the original host record. As a result, the MX record shown in Figure 3-13 relies on the presence in the zone of a host record named SRV12.nwtraders.msft.

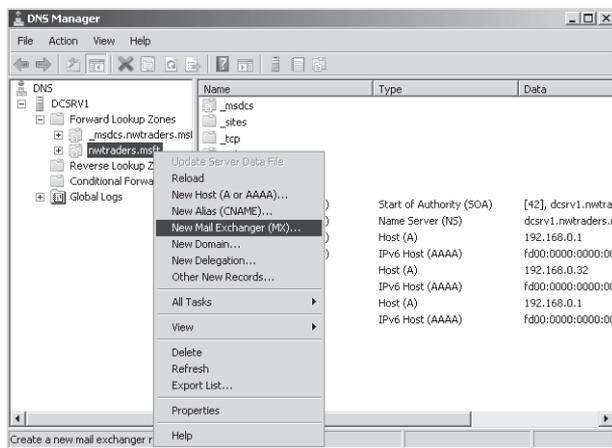


FIGURE 3-12 Creating a new resource record

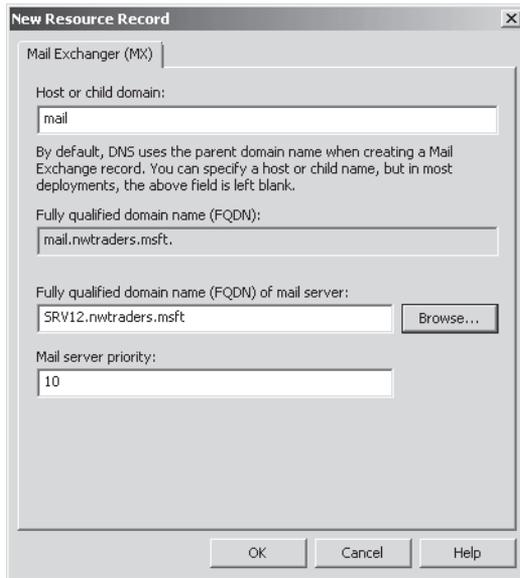


FIGURE 3-13 Defining a new MX record

Record Types

The most common resource records you need to create manually include the following:

- Host (A or AAAA)
- Alias (CNAME)
- Mail exchanger (MX)
- Pointer (PTR)
- Service location (SRV)

HOST RESOURCE RECORDS

For most networks, host (A or AAAA) resource records make up the majority of resource records in a zone database. These records are used in a zone to associate computer names (host names) to IP addresses.

After you create them in the DNS Manager console, an A resource record that maps the host name `server1.lucernepublishing.com` to the IPv4 address `192.168.0.99` and an AAAA resource record that maps the same name to the IPv6 address `fd00:0:0:5::8` would be represented textually within the standard zone file `lucernepublishing.com.dns` in the following way:

```

;
; Zone records
;

server1                A                192.168.0.99
                      AAAA             fd00:0:0:5::8

```

Even when dynamic updates are enabled for a particular zone, in some scenarios it might be necessary to add host records manually to that zone. For example, in Figure 3-14, a company named Contoso, Ltd., uses the domain name `contoso.com` for both its public namespace and its internal Active Directory domain. In this case, the public web server named `www.contoso.com` is located outside the Active Directory domain and performs updates only on the public DNS server authoritative for `contoso.com`. Internal clients, however, point their DNS requests toward internal DNS servers. Because the A record for `www.contoso.com` is not updated dynamically on these internal DNS servers, the record must be added manually for internal clients to resolve the name and connect to the public web server.

Another case in which you might need to add host records manually is when you have a UNIX server on your network. For example, in Figure 3-15, a company named Fabrikam, Inc., uses a single Active Directory domain named `fabrikam.com` for its private network. The network also includes a UNIX server named `App1.fabrikam.com` that runs an application critical to the company's daily operations. Because UNIX servers often do not perform dynamic updates (or especially secure dynamic updates) with Microsoft DNS servers, you might need to add a host record manually for `App1` on the DNS server hosting the `fabrikam.com` zone. Otherwise, users will not be able to connect to the application server when they specify it by FQDN.

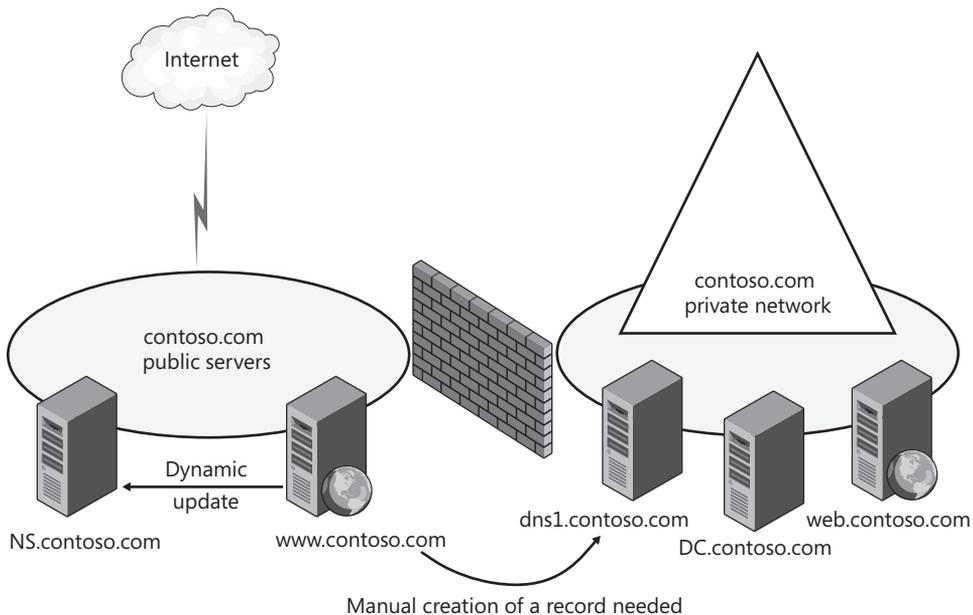


FIGURE 3-14 Adding a host record for a public web server

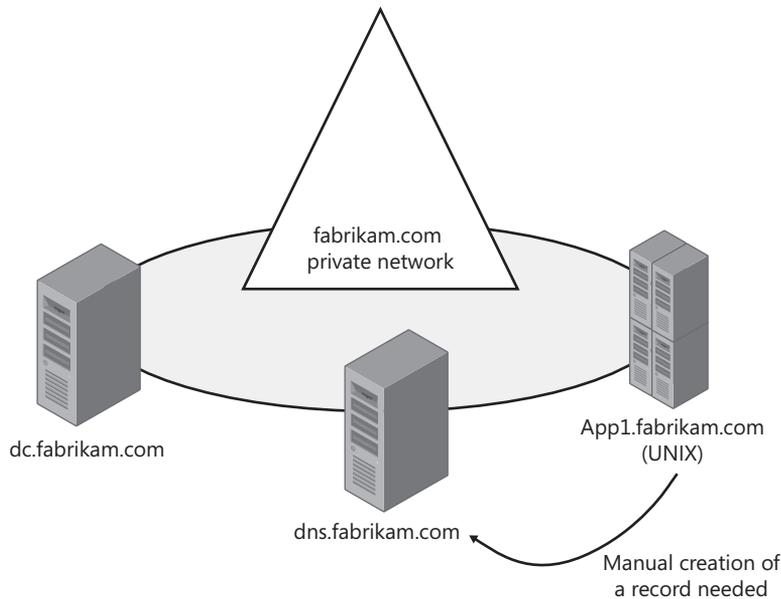


FIGURE 3-15 Adding a host record for a private UNIX server



EXAM TIP

If you can ping a computer by IP address but not by name, the computer might be missing a host record in DNS. You can attempt to remedy this situation by executing the `Ipconfig /registerdns` command at that computer—but only if the client computer is running Windows 2000 or later.

ALIAS RESOURCE RECORDS

Alias (CNAME) resource records are sometimes called *canonical names*. These records allow you to use more than one name to point to a single host. For example, the well-known server names (ftp, www) are typically registered using CNAME resource records. These records map the host name specific to a given service (such as ftp.lucernepublishing.com) to the actual A resource record of the computer hosting the service (such as server-boston.lucernepublishing.com).

CNAME resource records are also recommended for use in the following scenarios:

- When a host specified in an A resource record in the same zone needs to be renamed
- When a generic name for a well-known server such as www needs to resolve to a group of individual computers (each with individual A resource records) that provide the same service (for example, a group of redundant web servers)

After you create it in the DNS Manager console, a CNAME resource record that maps the alias ftp.lucernepublishing.com to the host name ftp1.lucernepublishing.com would be represented textually within the lucernepublishing.com.dns standard zone file as follows:

```
ftp           CNAME      ftp1.lucernepublishing.com.
```

MAIL EXCHANGER RESOURCE RECORDS

The mail exchanger (MX) resource record is used by SMTP (mail) agents to locate other SMTP servers in a remote domain, typically for the purpose of routing mail to that domain. An MX record maps the domain name found in an email address (such as joe@lucernepublishing.com) to a particular server hosting the mail server in that domain.

Multiple MX records are also often used to specify a preferred SMTP server and a backup SMTP server. Each MX record is assigned a Mail Server Priority value, with the lower values representing higher preference. The DNS server responds to the original query by returning all the MX records matching the domain name. Finally, the SMTP agent that has queried the DNS server looks at the MX records it has received and then contacts the server whose record is assigned the lowest Mail Server Priority value. If the server assigned the lowest value is unavailable, the server assigned the next lowest value is contacted.

When two or more MX records are assigned the lowest preference value, DNS round robin can be used to balance the workload evenly among the SMTP servers corresponding to those MX records. For example, if you create three MX records in DNS Manager for mailserver1, mailserver2, and mailserver3, and then assign these records preference values of 10, 10, and 20, respectively, the workload would be split evenly between mailserver1 and mailserver2. Mailserver3 would be used as a backup. These resource records would be represented textually within the lucernepublishing.com.dns zone file as follows:

```
@           MX      10    mailserver1.lucernepublishing.com.  
@           MX      10    mailserver2.lucernepublishing.com.  
@           MX      20    mailserver3.lucernepublishing.com.
```

NOTE WHAT DOES THE @ SYMBOL MEAN?

In this example, the @ symbol represents the local domain name contained in an email address.



REAL WORLD

J.C. Mackin

In theory as well as in Microsoft exams, the Mail Server Priority value you set for MX records takes precedence over round-robin distribution in DNS. It doesn't always work that way in reality, however.

At first, everything works according to plan. An SMTP agent queries a DNS server for an MX record corresponding to a particular domain name, and the DNS server responds with a list of all matching MX records. The order of that list rotates from response to response if round robin is left enabled on the DNS server. So far, so good, but what happens after that point is inconsistent. The SMTP agent is then supposed to scan through the response list and contact the server whose MX record is weighted with the lowest preference value. In reality, this happens only sometimes. Reports are common of SMTP agents ignoring the preference values in MX records and contacting merely the first server in the DNS response list.

The take-away? If you want to use preference values to set load balancing and specify a backup mail server, go ahead. Know, however, that this configuration is only approximate, and that your mail server workload will be distributed in a way that is hard to predict accurately.

POINTER RESOURCE RECORDS

The pointer (PTR) resource record is used in reverse lookup zones only to support reverse lookups, which perform queries to resolve IP addresses to host names or FQDNs. Reverse lookups are performed in zones rooted in the in-addr.arpa domain. PTR resource records can be added to zones manually or automatically.

After you create it in the DNS Manager console, a PTR resource record that maps the IP address 192.168.0.99 to the host name server1.lucernepublishing.com would be represented textually within a zone file as follows:

```
99          PTR      server1.lucernepublishing.com.
```

NOTE WHY IS THE PTR RECORD NAMED 99?

In a reverse lookup zone, the host ID portion of an IPv4 address is equivalent to a host name. The 99 therefore represents the name assigned to the host within the 0.168.192.in-addr.arpa zone. This zone corresponds to the 192.168.0.0 subnet.

SERVICE RESOURCE RECORDS

Service location (SRV) resource records are used to specify the location of specific services in a domain. Client applications that are SRV-aware can use DNS to retrieve the SRV resource records for given application servers.

Active Directory Directory Service (AD DS) is an example of an SRV-aware application. The Netlogon service uses SRV records to locate domain controllers in a domain by searching the domain for the Lightweight Directory Access Protocol (LDAP) service. (LDAP is the protocol used to query AD DS.)

If a computer needs to locate a domain controller in the `lucernepublishing.com` domain, the DNS client sends an SRV query for the name:

```
_ldap._tcp.lucernepublishing.com.
```

The DNS server then responds to the client with all records matching the query.

Although SRV resource records for AD DS are created automatically, you might need to create SRV records manually for other services or if some records have been accidentally deleted. The following example shows the textual representation of two SRV records that have been configured manually in the DNS Manager console:

```
_ldap._tcp SRV 0 0 389 dc1.lucernepublishing.com.  
SRV 10 0 389 dc2.lucernepublishing.com.
```

In the example, an LDAP server (domain controller) with a priority of 0 (highest) is mapped to port 389 at the host `dc1.lucernepublishing.com`. A second domain controller with a lower priority of 10 is mapped to port 389 at the host `dc2.lucernepublishing.com`. Both entries have a 0 value in the weight field, which means that no load balancing has been configured among servers with equal priority.

Enabling DNS to Use WINS Resolution

You can use the WINS tab in the properties of a zone to specify a WINS server that the DNS Server service can contact to look up names not found through DNS queries. When you specify a WINS server on the WINS tab in the properties of a forward lookup zone, a special WINS resource record pointing to that WINS server is added to the zone. When you specify a WINS server on the WINS tab in a reverse lookup zone, a special WINS-R resource record pointing to that WINS server is added to the zone.

For example, if a DNS client queries for the name `ClientZ.contoso.com` and the preferred DNS server cannot find the answer through any of its usual sources (local zone data, cache, queries to other servers), the server then queries the WINS server specified in the WINS record for the name "CLIENTZ." If the WINS server responds with an answer to the query, the DNS server returns this response to the original client.



EXAM TIP

For the 70-642 exam, you need to understand the function of the WINS and WINS-R records in a DNS zone.

Aging and Scavenging

Aging in DNS refers to the process of using time stamps to track the age of dynamically registered resource records. *Scavenging* refers to the process of deleting outdated resource records on which time stamps have been placed. Scavenging can occur only when aging is enabled. Together, aging and scavenging provide a mechanism to remove stale resource records, which can accumulate in zone data over time. Both aging and scavenging are disabled by default.

Enabling Aging

To enable aging for a zone, you have to enable this feature either at the server level or at the zone level.

To enable aging at the server level, first open the Server Aging/Scavenging Properties dialog box by right-clicking the server icon in the DNS Manager console tree and then choosing Set Aging/Scavenging For All Zones, as shown in Figure 3-16. Next, in the Server Aging/Scavenging Properties dialog box that opens, select the Scavenge Stale Resource Records check box. Although this setting enables aging and scavenging for all new zones at the server level, it does not automatically enable aging or scavenging on existing Active Directory–integrated zones at the server level. To do that, click OK, and then, in the Server Aging/Scavenging Confirmation dialog box that appears, enable the option to apply these settings to existing Active Directory–integrated zones, as shown in Figure 3-17.

To enable aging and scavenging at the zone level, open the properties of the zone and then, in the General tab, click Aging, as shown in Figure 3-18. Then, in the Zone Aging/Scavenging Properties dialog box that opens, select the Scavenge Stale Resource Records check box, as shown in Figure 3-19.

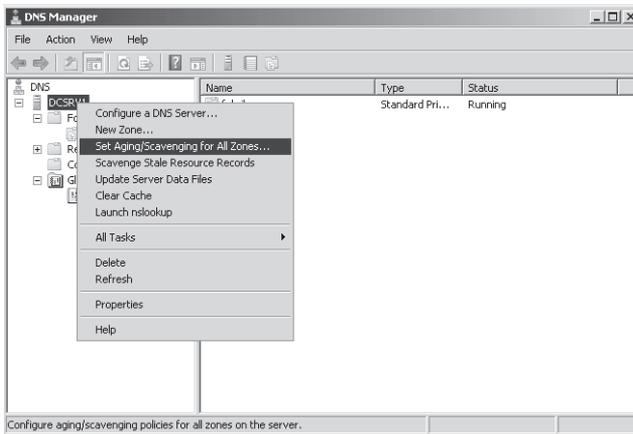


FIGURE 3-16 Enabling aging at the server level

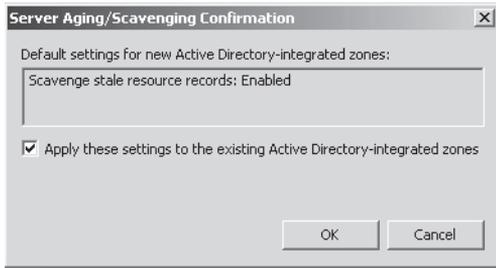


FIGURE 3-17 Enabling aging on Active Directory-integrated zones

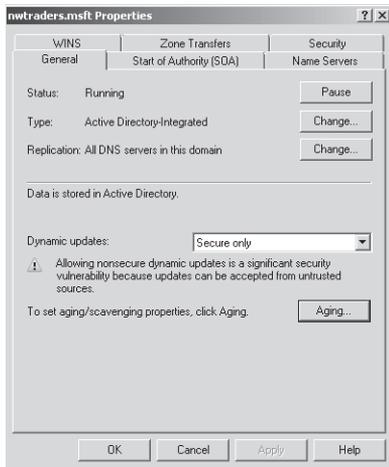


FIGURE 3-18 Accessing aging properties for a zone

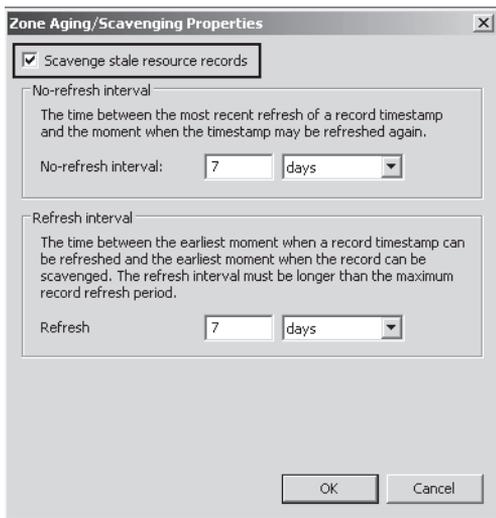


FIGURE 3-19 Enabling aging and scavenging at the zone level

Time Stamping

The DNS server performs aging and scavenging by using time stamp values set on resource records in a zone. Active Directory–integrated zones perform time stamping for dynamically registered records by default, even before aging and scavenging are enabled. However, primary standard zones place time stamps on dynamically registered records in the zone only after aging is enabled. Manually created resource records for all zone types are assigned a time stamp of 0; this value indicates that they will not be aged.

Modifying Zone Aging/Scavenging Properties

The Zone Aging/Scavenging Properties dialog box enables you to modify two key settings related to aging and scavenging: the no-refresh interval and the refresh interval.

- **Modifying the no-refresh interval** The *no-refresh interval* is the period after a time stamp during which a zone or server rejects a time stamp refresh. The no-refresh feature prevents the server from processing unnecessary refreshes and reduces unnecessary zone transfer traffic. The default no-refresh interval is 7 days.
- **Modifying the refresh interval** The *refresh interval* is the time after the no-refresh interval during which time stamp refreshes are accepted and resource records are not scavenged. After the no-refresh and refresh intervals expire, records can be scavenged from the zone. The default refresh interval is 7 days. Consequently, when aging is enabled, dynamically registered resource records can be scavenged after 14 days by default.

Performing Scavenging

Scavenging in a zone is performed either automatically or manually. For scavenging to be performed automatically, you must enable automatic scavenging of stale resource records on the Advanced tab of the DNS server properties dialog box, as shown in Figure 3-20.

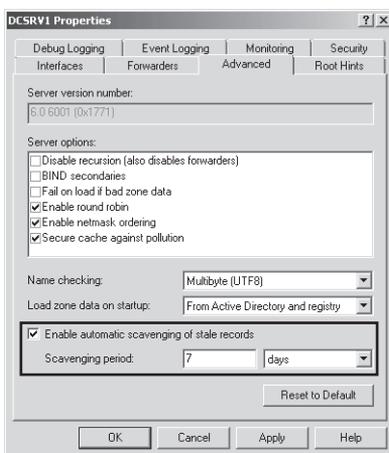


FIGURE 3-20 Enabling automatic scavenging on a DNS server

When this feature is not enabled, you can perform manual scavenging in zones by right-clicking the server icon in the DNS Manager console tree and then choosing Scavenge Stale Resource Records, as shown in Figure 3-21.

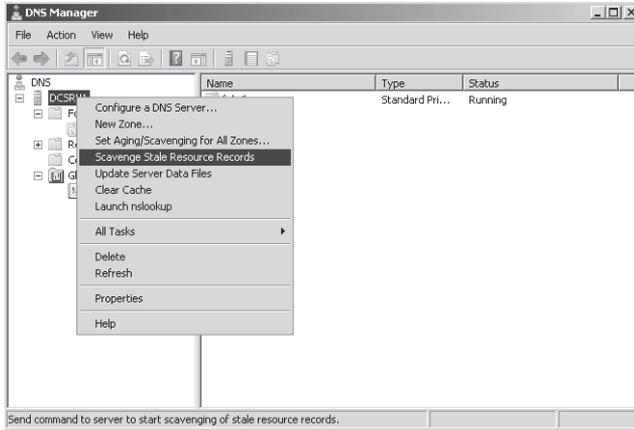


FIGURE 3-21 Performing manual scavenging for zones



Quick Check

- What kind of zones do not automatically perform time stamping on dynamically created resource records?

Quick Check Answer

- Standard zones

Using a GlobalNames Zone

Windows Server 2008 and Windows Server 2008 R2 include a new feature that enables all DNS clients in an Active Directory forest to use single-label name tags such as “Mail” to connect to specific server resources located anywhere in the forest. This feature can be useful when the default DNS suffix search list for DNS clients would not enable users to connect quickly (or connect at all) to a resource by using a single-label name.

To support this functionality, the DNS Server role in Windows Server 2008 and Windows Server 2008 R2 includes capability for a GlobalNames zone. The GlobalNames zone does not exist by default, but by deploying a zone with this name you can provide access to selected resources through single-label names without relying on WINS. These single-label names typically refer to records for important, well-known, and widely used servers—servers that are already assigned static IP addresses.

Figure 3-22 shows a GlobalNames zone with a record for a server with a single-label name of “Mail.”

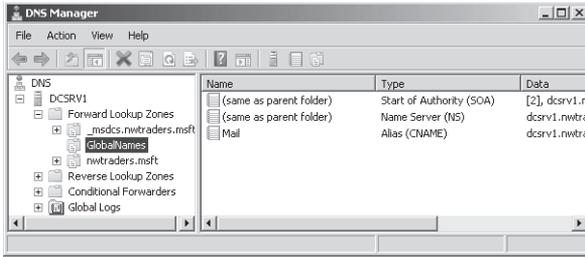


FIGURE 3-22 The GlobalNames zone

Deploying a GlobalNames Zone

The GlobalNames zone is compatible only with DNS servers running Windows Server 2008 and Windows Server 2008 R2. Therefore, it cannot replicate to servers running earlier versions of Windows Server.

There are three basic steps in deploying a GlobalNames zone.

1. Enable GlobalNames zone support. You can perform this step before or after you create the zone, but you must perform it on every DNS server to which the GlobalNames zone will be replicated.

At an elevated command prompt, type the following:

```
dnscommand . /config /enableglobalnamesupport 1
```

In this case the "." is used to represent the local server. If you want to enable GlobalNames zone support on a remote server, substitute the "." for the DNS server name.

2. Create the GlobalNames zone. The next step in deploying a GlobalNames zone is to create the zone on a DNS server that is a domain controller running Windows Server 2008 or Windows Server 2008 R2. The GlobalNames zone is not a special zone type; rather, it is simply an Active Directory–integrated forward lookup zone that is called GlobalNames. When you create the zone, make sure to select the option to replicate zone data to all DNS servers in the forest. (This option appears on the Active Directory Zone Replication Scope page of the New Zone Wizard.)
3. Populate the GlobalNames zone. For each server for which you want to be able to provide single-label name resolution, create an alias (CNAME) resource record in the GlobalNames zone. The name you give each CNAME record represents the single-label name that users will use to connect to the resource. Note that each CNAME record points to a host record in another zone.



EXAM TIP

Expect to see a question about the GlobalNames zone on the 70-642 exam.



Quick Check

- Why would you use a GlobalNames zone?

Quick Check Answer

- To facilitate the resolution of single-label computer names in a large network.

PRACTICE Deploying a GlobalNames Zone

In this practice, you will create the GlobalNames Zone to enable connectivity to a specific single-label name throughout an Active Directory forest.

EXERCISE 1 Enabling the GlobalNames Zone

In this exercise, you will enable the GlobalNames zone on Dcsrv1. In a production environment, you would need to perform this step on every DNS server in the forest.

1. Log on to Nwtraders from Dcsrv1 as a domain administrator.
2. Open an elevated command prompt.
3. At the command prompt, type **dnscmd . /config /enableglobalnamesupport 1**. Note the space in this command after the "." You receive an output message indicating that the Registry property was successfully reset.

EXERCISE 2 Creating the GlobalNames Zone

In this exercise, you will create a new DNS forward lookup zone named GlobalNames on Dcsrv1.

1. While you are logged on to Nwtraders from Dcsrv1 as a domain administrator, open DNS Manager.
2. In the DNS Manager console tree, right-click the Forward Lookup Zones container, and then choose New Zone.
3. On the Welcome page of the New Zone Wizard, read the text, and then click Next.
4. On the Zone Type page, read all the text on the page. Leave the default selections of Primary and Store The Zone In Active Directory, and then click Next.
5. On the Active Directory Zone Replication Scope page, select To All DNS Servers In This Forest, and then click Next.
6. On the Zone Name page, type **GlobalNames**, and then click Next.
7. On the Dynamic Update page, select the Do Not Allow Dynamic Updates option, and then click Next. You should choose the option because dynamic updates are not supported with the GlobalNames zone.
8. After the Completing The New Zone Wizard page, read the text, and then click Finish. In the DNS Manager console tree, the new GlobalNames zone appears.

EXERCISE 3 Adding Records to the GlobalNames Zone

In this exercise, you will add records to the GlobalNames zone so that you can later test its functionality.

1. While you are still logged on to Nwtraders from Dcsrv1 as a domain administrator, in the DNS Manager console tree, select and then right-click the GlobalNames zone, and then choose New Alias (CNAME).
2. In the New Resource Record dialog box, in the Alias Name text box, type **mail**.
3. In the Fully Qualified Domain Name (FQDN) For Target Host text box, type **dcsrv1.nwtraders.msft**, and then click OK. A new alias (CNAME) record with the name "mail" now appears in the GlobalNames zone.

EXERCISE 4 Testing the GlobalNames Zone

In this exercise, you will attempt to resolve the name of the new record you have created. The GlobalNames zone is used to resolve single-name tags anywhere in an Active Directory forest.

1. Log on to Nwtraders from Boston as a domain administrator.
2. Open an elevated command prompt.
3. At the command prompt, type **ping mail**. Boston translates the name "mail" to dcsrv1.nwtraders.msft and then pings the address of that server. You know that this name has been resolved from the GlobalNames zone because there is no record in the Nwtraders.msft zone for a host or an alias named "mail."
4. Log off both Dcsrv1 and Boston.

Lesson Summary

- A DNS zone is a database containing records that associate names with addresses for a defined portion of a DNS namespace. To create a new zone on a DNS server, you can use the New Zone Wizard in DNS Manager. The New Zone Wizard enables you to choose a zone type, specify a forward or reverse lookup zone, set the zone replication scope, name the zone, and configure options for dynamic updates.
- A primary zone provides original read-write source data that allows the local DNS server to answer DNS queries authoritatively about a portion of a DNS namespace. A secondary zone provides an authoritative, read-only copy of a primary zone or another secondary zone. A stub zone is similar to a secondary zone, but it contains only those resource records necessary to identify the authoritative DNS servers for the master zone.
- When you create a new primary or stub zone on a domain controller, the Zone Type page gives you the option to store the zone in Active Directory. There are several advantages to integrating your DNS zone with Active Directory, including ease of management, the availability of multiple primary zones, and improved security.

- When you do not store a zone in Active Directory, the zone is called a standard zone, and zone data is stored in text files on the DNS server.
- When you create a new zone, two types of records required for the zone are automatically created: an SOA record and at least one NS record. The SOA record defines basic properties for the zone. NS records determine which servers hold authoritative information for the zone.
- Aging in DNS refers to the process of using time stamps to track the age of dynamically registered resource records. Scavenging refers to the process of deleting outdated resource records on which time stamps have been placed.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 1, "Creating and Configuring Zones." The questions are also available on the companion CD if you prefer to review them in electronic form.

NOTE ANSWERS

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the "Answers" section at the end of the book.

1. You want to prevent a certain host (A) record from being scavenged. The record belongs to a portable computer named LaptopA that connects to the network only infrequently. LaptopA obtains its address from a DHCP server on the network. Which of the following steps would best enable you to achieve this goal?
 - A. Disable scavenging on the zone in which the record has been created.
 - B. Disable scavenging on the server with which the computer registers its record.
 - C. Assign the computer a static address.
 - D. Create a record for LaptopA manually.

2. You are a network administrator for a company named Fabrikam, Inc. A DNS server for the network is located on a member server named Dns1 in the Fabrikam.com Active Directory domain. Dns1 provides name resolution for the Fabrikam.com domain only. Occasionally, you see DNS records for unauthorized computers in the Fabrikam.com zone. These computers do not have accounts in the Fabrikam.com Active Directory domain. What steps should you take to prevent unauthorized computers from registering host records with the DNS server? (Choose three. Each answer represents part of the solution.)
 - A. Promote DNS1 to a domain controller.
 - B. Choose the option to store the zone in Active Directory.
 - C. Clear the option to store the zone in Active Directory.
 - D. Configure the zone not to accept dynamic updates.
 - E. Configure the zone to accept secure and nonsecure dynamic updates.
 - F. Configure the zone to accept secure updates only.

3. You work for Contoso.com as a network administrator. In your network, you use three servers to distribute the email workload. The contoso.com zone file shows the following data for the three mail servers.

```
@      MX      10      mailserver1.contoso.com.  
@      MX      10      mailserver2.contoso.com.  
@      MX      10      mailserver3.contoso.com.
```

You want to configure mail servers so that mailserver1 and mailserver2 share the email workload and mailserver3 is used as a backup. What should you do?

- A. Create a second, identical MX record for both mailserver1 and mailserver2. Enable round robin on the DNS server.
- B. Create a second, identical MX record for both mailserver1 and mailserver2. Disable round robin on the DNS server.
- C. Replace the value *10* with *1* in the mailserver3 MX record.
- D. Replace the value *10* with *20* in the mailserver3 MX record.

Lesson 2: Configuring Zone Replication, Transfers, and Delegations

To deploy DNS in a medium-to-large sized organization, you need to do more than configure DNS on an individual server. You also have to design DNS in a way that keeps the processing workload and administration workload distributed in a sensible way. For all but the smallest organizations, achieving these goals requires you to deploy more than one DNS server.

To manage your DNS data well and preserve data consistency among multiple servers, you need to understand zone replication, transfers, and delegations.

After this lesson, you will be able to:

- Configure a zone replication scope appropriate to your network.
- Create a new directory partition and enlist a server in that partition.
- Understand the benefits of a secondary zone.
- Implement a secondary zone.
- Understand the benefits of zone delegations.
- Understand the benefits of stub zones.
- Implement a stub zone.
- Enable zone transfers to secondary and stub zones.

Estimated lesson time: 90 minutes

Configuring Zone Replication for Active Directory–Integrated Zones

You can install Active Directory–integrated zones only on domain controllers on which the DNS Server role is installed. Active Directory–integrated zones are generally preferable to standard zones because they offer multimaster data replication, simpler configuration, and improved security and efficiency. With Active Directory–integrated storage, DNS clients can send updates to any Active Directory–integrated DNS server. These updates are then copied to other Active Directory–integrated DNS servers by means of Active Directory replication.

Replication and Application Directory Partitions

DNS data for any particular zone can be replicated among domain controllers in a number of ways, depending on the application directory partition on which the DNS zone data is stored.

DOMAINDNSZONES AND FORESTDNSZONES

A *partition* is a data structure in Active Directory that distinguishes data for different replication purposes. By default, domain controllers include two application directory partitions reserved for DNS data: `DomainDnsZones` and `ForestDnsZones`. The `DomainDnsZones` partition is replicated among all domain controllers that are also DNS servers in a particular domain, and the `ForestDnsZones` partition is replicated among all domain controllers that are also DNS servers in every domain in an Active Directory forest.

Each of these application directory partitions is designated by a DNS subdomain and an FQDN. For example, in an Active Directory domain named `east.nwtraders.msft` and whose root domain in the Active Directory forest is `nwtraders.msft`, the built-in DNS application partition directories are specified by these FQDNs: `DomainDnsZones.east.nwtraders.msft` and `ForestDnsZones.nwtraders.msft`.

You can see evidence of these partitions when you browse DNS Manager, as shown in Figure 3-23. Note that the `ForestDnsZones` name is located in the `nwtraders.msft` zone. Note also that each zone includes a `DomainDnsZones` name that points to the partition that is replicated only within each local domain.

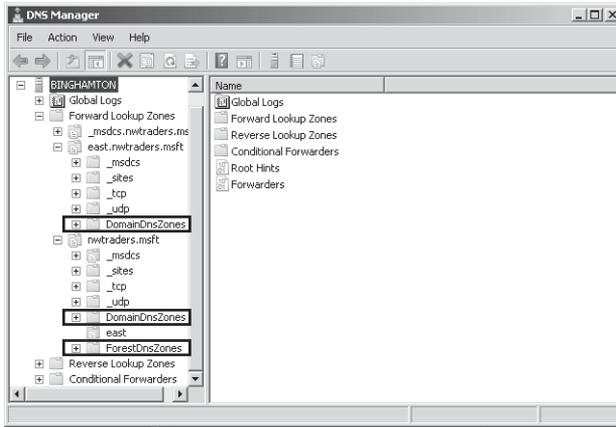


FIGURE 3-23 You can see evidence of the built-in directory partitions for DNS within an Active Directory–integrated zone.

CUSTOM APPLICATION DIRECTORY PARTITIONS

In addition to the two application directory partition types `DomainDnsZones` and `ForestDnsZones`, you can create a custom (user-defined) application directory partition with a name of your own choosing. You can then configure a zone to be stored in this new structure that you created. By default, the new application directory partition exists only on the server on which you created the partition, but you can enlist other servers in the partition so that replication of its data content are copied to those particular servers you choose.

The replication pattern displayed by these three application data partition types—`DomainDnsZones`, `ForestDnsZones`, and custom—is illustrated in Figure 3-24.

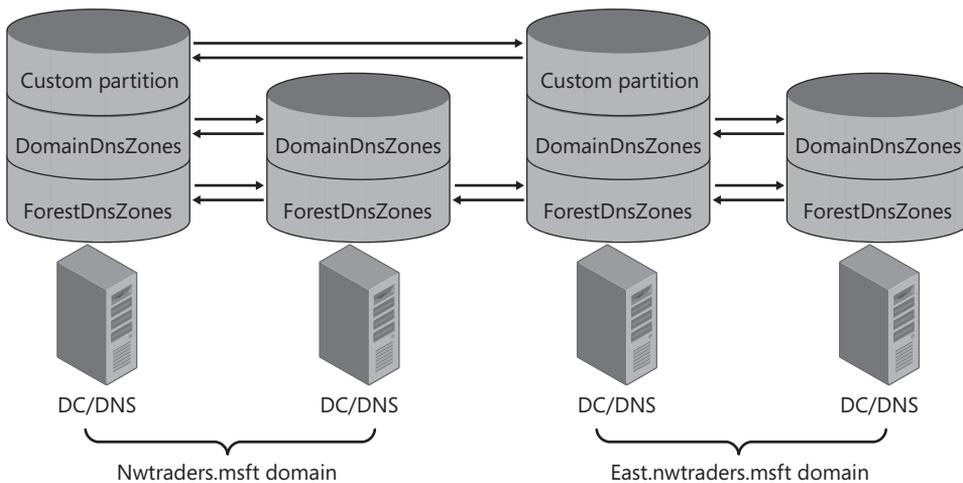


FIGURE 3-24 Replication patterns among application directory partitions

STORING DNS DATA IN THE DOMAIN PARTITION

The final storage option for an Active Directory–integrated zone is to store the zone in the domain partition along with all remaining data for a domain. In this configuration, the DNS data does not replicate merely to domain controllers that are also DNS servers; it replicates to all domain controllers in general in the local domain. This option is not ideal because it generates unnecessary replication traffic. However, you need to use it if you want your DNS data to be replicated to computers running Windows 2000 Server.

Choosing Zone Replication Scope

The partition in which a zone is stored effectively determines the replication scope for that zone. Replication scope is set when an Active Directory–integrated zone is first created. When you use Dcpromo to promote a server to a domain controller in a new domain, the new Active Directory–integrated zone created for the domain is stored automatically in the DomainDnsZones partition. However, when you create a new zone by using the New Zone Wizard instead, you are given an opportunity on the Active Directory Zone Replication Scope page to choose the partition in which to store the zone, as shown in Figure 3-25.

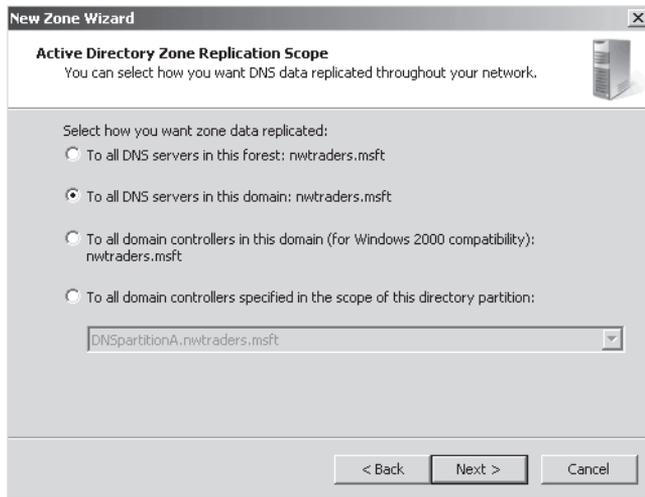


FIGURE 3-25 Choosing zone replication scope for a new zone

The four options presented on the Active Directory Zone Replication Scope page are the following:

- **To All DNS Servers In This Forest** This option stores the new zone in the ForestDnsZones partition. Every domain controller in the entire forest and on which the DNS Server role is installed will receive a copy of the zone.
- **To All DNS Servers In This Domain** This option stores the new zone in the DomainDnsZones partition. Every domain controller in the local domain and on which the DNS Server role is installed will receive a copy of the zone.

- **To All Domain Controllers In This Domain (For Windows 2000 Compatibility)** This option stores the zone in the domain partition. Every domain controller in the local domain will receive a copy of the zone, regardless of whether the DNS Server role is installed on that domain controller. This setting is required for compatibility with Windows 2000 Server domain controllers because Windows 2000 Server does not support directory partitions.
- **To All Domain Controllers Specified In The Scope Of This Directory Partition** This option stores the zone in the user-created application directory partition specified in the associated drop-down list box. For a domain controller to fall within the scope of such a directory partition, you must manually enlist that domain controller in the partition.

After a new zone is created, you can choose to change the replication scope for the zone at any time. To do so, on the General tab of the properties of the zone, click the Change button associated with replication, as shown in Figure 3-26.

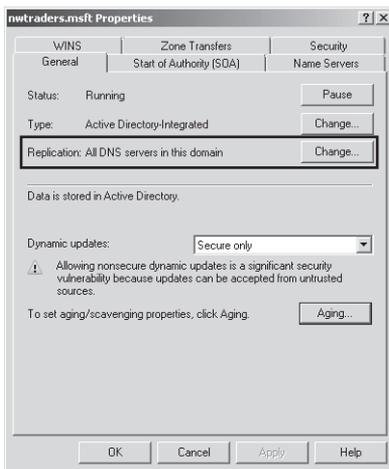


FIGURE 3-26 Changing the replication scope of an existing zone

This step opens the Change Zone Replication Scope dialog box, which, as shown in Figure 3-27, provides the same zone replication scope options that the New Zone Wizard does.

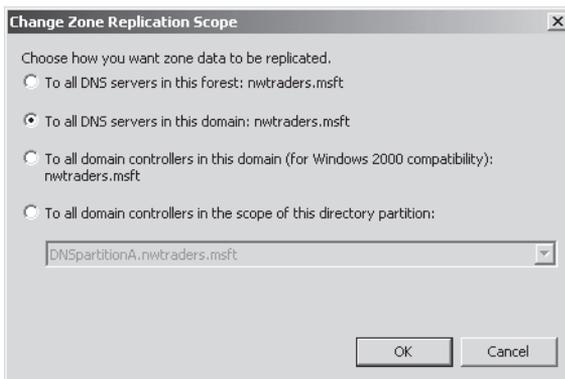


FIGURE 3-27 Modifying the replication scope for an existing zone

When deciding which replication scope to choose, consider that the broader the replication scope, the greater the network traffic caused by replication. For example, if you choose to have Active Directory–integrated DNS zone data replicated to all DNS servers in the forest, this setting produces greater network traffic than does replicating the DNS zone data to all DNS servers in the local domain only. On the other hand, replicating zone data to all DNS servers in a forest can improve forest-wide name resolution performance and increase fault tolerance.

NOTE RE-CREATING DOMAINDNSZONES AND FORESTDNSZONES

If either of the default application directory partitions is deleted or damaged, you can re-create them in DNS Manager by right-clicking the server node and choosing Create Default Application Directory Partitions.

Creating Custom Application Directory Partitions

You can create your own custom application directory partitions for use with DNS and then enlist selected domain controllers in your network to host replicas of this partition.

To accomplish this task, first create the partition by typing the following command:

```
dnscmd servername /createdirectorypartition FQDN
```

Then enlist other DNS servers in the partition by typing the following command:

```
dnscmd servername /enlistdirectorypartition FQDN
```

For example, to create an application directory partition named DNSpartitionA on a computer named Server1 in the Active Directory domain contoso.com, type the following command:

```
dnscmd server1 /createdirectorypartition DNSpartitionA.contoso.com
```

NOTE USE A DOT (".") FOR THE LOCAL SERVER NAME

You can substitute a "." for the server name if you are executing the command on the same server on which you want to create the partition.

To enlist a computer named Server2 in the application directory partition, type the following command:

```
dnscmd server2 /enlistdirectorypartition DNSpartitionA.contoso.com
```

NOTE WHO CAN CREATE A CUSTOM APPLICATION DIRECTORY PARTITION?

You must be a member of the Enterprise Admins group to create an application directory partition.

After you create a new application directory partition, that partition will appear as an option in the drop-down list box both on the Active Directory Zone Replication Scope page of the New Zone Wizard and in the Change Zone Replication Scope dialog box. To store a zone in the new partition, choose To All Domain Controllers Specified In The Scope Of This Directory Partition and then select the partition in the drop-down list box.



EXAM TIP

Expect to be tested on application directory partition concepts, as well as on the options in the Change Zone Replication Scope dialog box.

Using Zone Transfers

When all your DNS servers are located on domain controllers, you will normally want to use Active Directory replication to keep zone data consistent among all DNS servers. However, this option is not available when you install a DNS server on a computer that is not a domain controller. In such cases, you cannot store the zone in Active Directory and instead must use a standard zone that stores data in a local text file on each DNS server. If your organization requires multiple DNS servers, the source data can be copied to read-only secondary zones hosted on other servers. To keep data consistent and up to date between a primary and any secondary zones, you need to configure zone transfers.

Zone transfers are essentially pull operations initiated on secondary zones that copy zone data from a master zone, which itself can be a primary or another secondary. In fact, the master zone for a secondary zone need not even be another standard zone—you can configure a secondary zone for an Active Directory–integrated primary zone. This arrangement might be suitable, for example, if you have two sites, one in New York and one in Los Angeles, each with its own Active Directory domain. In each domain, you might want to provide name resolution for the opposite domain without installing a new domain controller and managing replication traffic between the two sites. Such an infrastructure is illustrated in Figure 3-28.

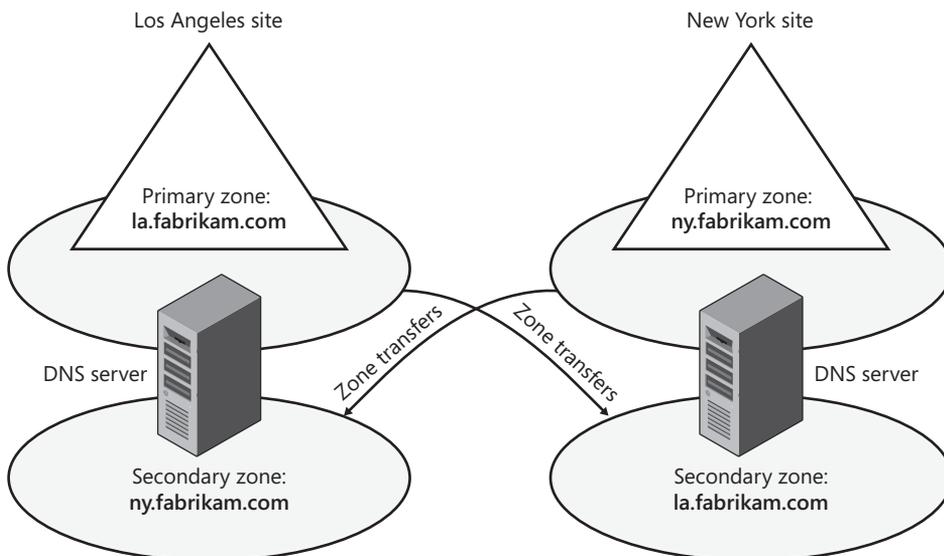


FIGURE 3-28 A DNS infrastructure with zone transfers between sites

Zone Transfer Initiation

Any of three events can trigger zone transfers on secondary zones:

- When the refresh interval of the primary zone's SOA resource record expires
- When a server hosting a secondary zone boots up
- When a change occurs in the configuration of the primary zone and this primary zone is configured to notify a secondary zone of zone updates

In the first two cases, the secondary server initiates a query to find out whether any updates in the zone have occurred. This information is determined by comparing the serial number (specified in the SOA record) of the secondary zone to the serial number of the master zone. If the master zone has a higher serial number, the secondary zone initiates a transfer from the master.

Enabling Zone Transfers

By default, zone transfers are disabled from any zone, and you must enable them on the Zone Transfers tab of the zone properties dialog box, as shown in Figure 3-29. After you have selected the option to allow zone transfers from the zone, you have a choice of three suboptions:

- **To Any Server** This option is the least secure. Because a zone transfer is essentially a copy of zone data, this setting allows anyone with network access to the DNS server to discover the complete contents of the zone, including all server and computer names along with their IP addresses. This option should therefore be used only in private networks with a high degree of security.
- **Only To Servers Listed On The Name Servers Tab** This option restricts zone transfers only to secondary DNS servers that have an NS record in the zone and are therefore already authoritative for zone data.
- **Only To The Following Servers** This option allows you to specify a list of secondary servers to which you will allow zone transfers. The secondary servers do not need to be identified by an NS record in the zone.

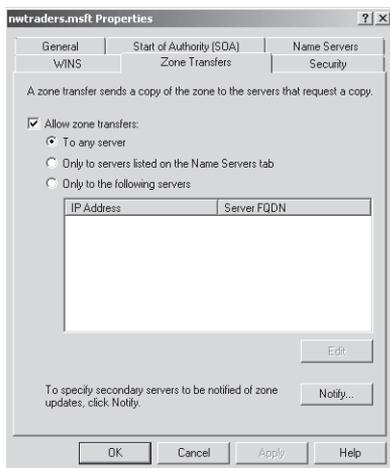


FIGURE 3-29 A zone on which transfers have been enabled

Configuring Notifications

The Zone Transfers tab also allows you to configure notification to secondary servers whenever a change occurs at the primary zone. Zone transfers are pull operations and cannot be initiated from the master to push new data to secondary zones. However, when a modification occurs in zone data, you can configure a primary zone to send a notification to any specified servers hosting secondary zones. When the secondary zone receives this notification, it initiates a zone transfer.

To configure notifications, click Notify on the Zone Transfers tab when zone transfers are enabled. This action opens the Notify dialog box, shown in Figure 3-30, in which you can specify secondary servers that should be notified whenever a zone update occurs at the local master server. By default, when zone transfers are enabled, all servers listed on the Name Servers tab are automatically notified of zone changes.

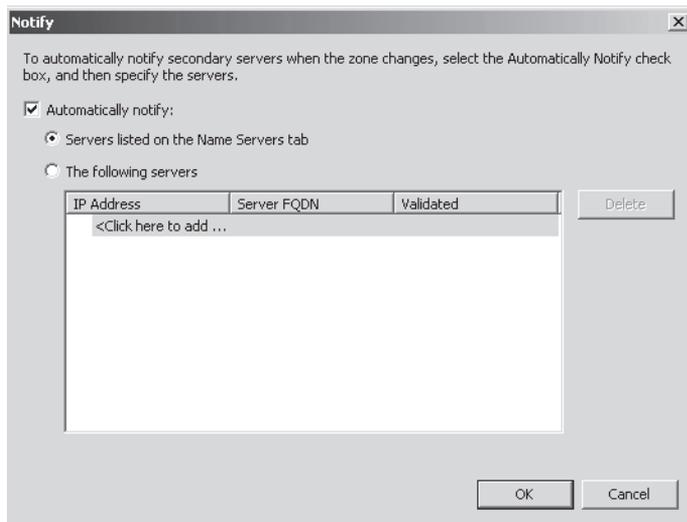


FIGURE 3-30 Notify dialog box

Manually Updating a Secondary Zone

By right-clicking a secondary zone in the DNS Manager console tree, you can use the shortcut menu to perform the following secondary zone update operations:

- **Reload** This operation reloads the secondary zone from the local storage.
- **Transfer From Master** The server hosting the local secondary zone determines whether the serial number in the secondary zone's SOA resource record has expired and then pulls a zone transfer from the master server.
- **Transfer New Copy Of Zone From Master** This operation performs a zone transfer from the secondary zone's master server regardless of the serial number in the secondary zone's SOA resource record.

Understanding Zone Delegations

To delegate a zone is to create a new zone for a subdomain within a DNS namespace and relinquish authority of that new zone. For example, the organization VeriSign manages the top-level domain “com” and creates new zones for subdomains such as microsoft.com. These child zones are then managed by the private organizations that have purchased the associated portion of the public DNS namespace. Delegations also appear within a single organization. For example, a company owning the domain contoso.com can delegate subdomains such as asia.contoso.com and eu.contoso.com to its various regional offices.

When to Delegate Zones

DNS delegations are automatically used to separate parent and child AD DS domains in a single forest. For example, if your organization originally includes a single AD DS domain northwindtraders.com and then creates a second child AD DS domain named ny.northwindtraders.com, the DNS namespace of the new child AD DS domain will automatically be configured as a new DNS zone and delegated subdomain of the parent zone. The authoritative DNS data for all computers in the child domain will be stored on DNS servers in that new AD DS domain.

Outside of an AD DS infrastructure, you should consider delegating a zone within your network whenever any of the following conditions are present:

- You need to delegate management of a DNS domain to a branch or department within your organization.
- You need to distribute the load of maintaining one large DNS database among multiple name servers to improve name resolution performance and fault tolerance.
- You need hosts and host names to be structured according to branch or departmental affiliation within your organization.

Above all, when choosing how to structure zones, you should use a plan that reflects the structure of your organization.

How Delegations Work

For a delegation to be implemented, the parent zone must contain an NS record and an associated A record (called a *glue record*) pointing to each authoritative server of the delegated domain. These records are necessary both to transfer authority to the new name servers and to provide referrals to other servers querying for names in the delegated namespace.

Figure 3-31 illustrates how DNS queries are handled with delegated subdomains. In the figure, a local DNS server named ns.contoso.com is authoritative for the domain contoso.com and has a configured delegation for the subdomain asia.contoso.com. If a client queries this local DNS server for the FQDN “hk4.asia.contoso.com”, the server consults the locally stored NS and A records that are configured for the delegation to determine that the authoritative name server for the asia.contoso.com domain is ns1.asia.contoso.com, and that this server’s IP address is 192.168.3.5. The local DNS server then queries ns1.asia.contoso.com for the name

“hk4.asia.contoso.com”. After the remote DNS server receives the query, it consults its locally stored database and responds to the querying DNS server with the IP address of the host hk4.asia.contoso.com, which is 192.168.3.10. The local DNS server then responds to the original querying client with the information requested.

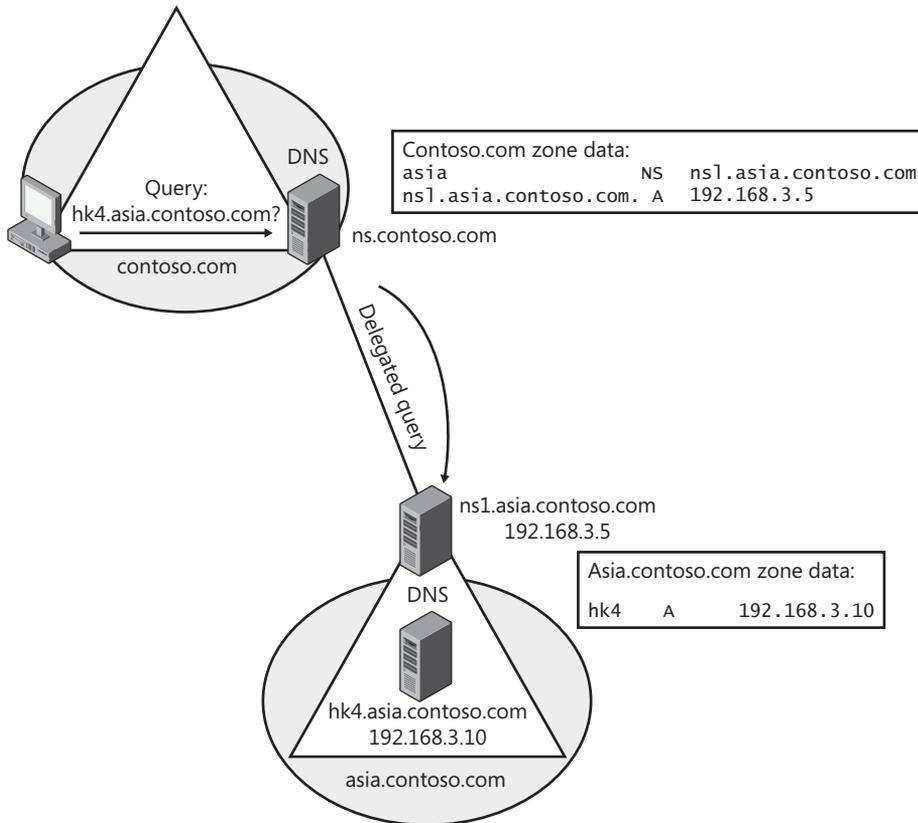


FIGURE 3-31 A delegated DNS query

Creating a Zone Delegation

To create a zone delegation, the domain to be delegated must already be created on a remote server that is authoritative for the DNS subdomain. Then, run the New Delegation Wizard on the server hosting the parent zone by right-clicking the parent zone folder in the DNS console and selecting New Delegation, as shown in Figure 3-32.

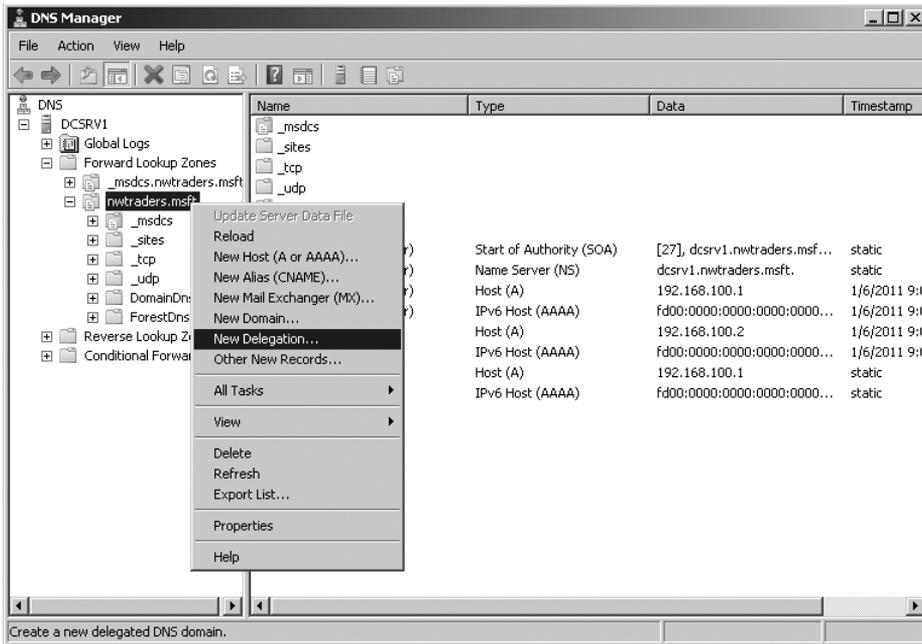


FIGURE 3-32 Creating a new delegation

To complete the New Delegation Wizard, you will need to specify the name of the delegated subdomain and the name of at least one name server that will be authoritative for the new zone. After you run the wizard, a new folder will appear in the DNS console tree representing the newly delegated subdomain.

Implementing Stub Zones

As you learned earlier, a stub zone is a copy of a zone that contains only the most basic records in the master zone. The purpose of a stub zone is to enable the local DNS server to forward queries to the name servers authoritative for the master zone. In this way, a stub zone is functionally similar to a zone delegation. However, because stub zones can initiate and receive zone transfers from the master (delegated) zone, stub zones provide the added benefit of informing parent zones of updates in the NS records of child zones.

An example of a stub zone is shown in Figure 3-33.

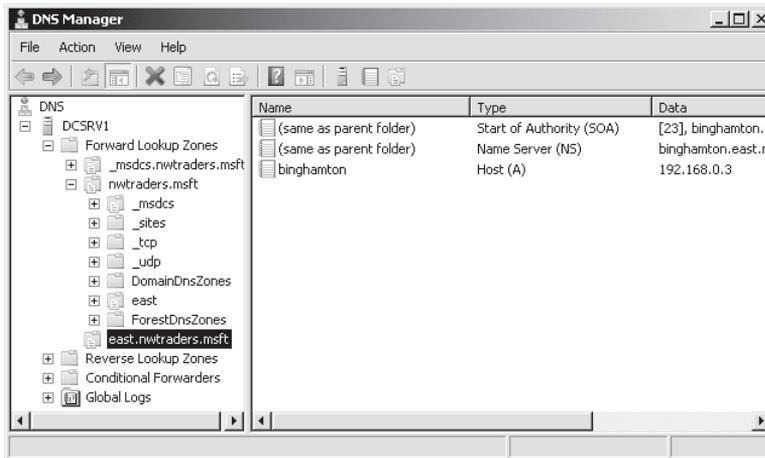


FIGURE 3-33 East.nwtraders.msft is a stub zone of a child zone hosted on remote server.

You can use stub zones to do the following:

- **Keep delegated zone information current** In a delegation scenario, a stub zone helps a parent zone stay up-to-date about the authoritative name servers for a delegated (child) subdomain.
- **Improve name resolution** Stub zones enable a DNS server to perform recursion using the stub zone's list of name servers without having to query the Internet or an internal server within the local DNS namespace. When stub zones are deployed for this reason, they are deployed not between parent and child zones but rather across domains in a large Active Directory forest or DNS namespace.

Stub Zone Example

Suppose that you are an administrator for the DNS server named Dns1.contoso.com, which is authoritative for the zone Contoso.com. Your company includes a child Active Directory domain, India.contoso.com, for which a delegation has been performed. When the delegation is originally performed, the child zone (which is Active Directory–integrated) contains only two authoritative DNS servers: 192.168.2.1 and 192.168.2.2. Later, administrators of the India.contoso.com domain deploy additional domain controllers and install the DNS Server role on these new domain controllers. However, these same administrators do not notify you of the addition of more authoritative DNS servers in their domain. As a result, Dns1.contoso.com is not configured with the records of the new DNS servers authoritative for India.contoso.com and continues to query only the two DNS servers that were defined in the original delegation.

You can remedy this problem by configuring Dns1.contoso.com to host a stub zone for India.contoso.com. As a result of this new stub zone, Dns1 learns through zone transfers about the new name servers authoritative for the India.contoso.com child zone. Dns1 is thus able to direct queries for names within the India.contoso.com namespace to all of that child zone's authoritative DNS servers.

This example is illustrated in Figure 3-34.

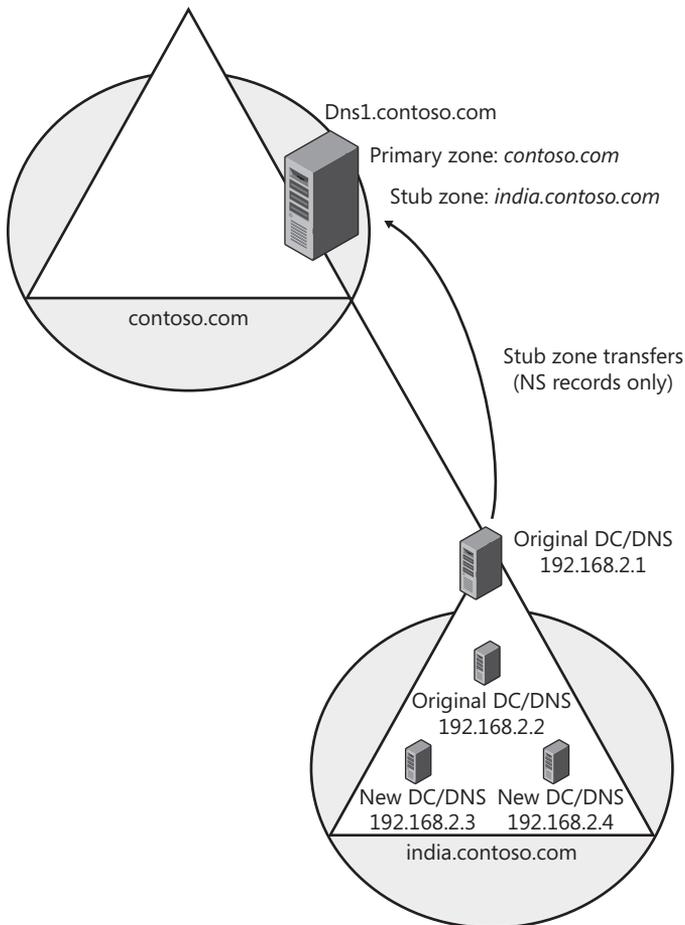


FIGURE 3-34 Stub zones enable a parent domain to keep an updated list of name servers in a child domain.

Other Uses for Stub Zones

Another use for stub zones is to facilitate name resolution across domains in a manner that avoids searching the DNS namespace for a common parent server. Stub zones can thus replace secondary zones in situations where achieving DNS connectivity across domains is important but providing data redundancy for the master zone is not. Also note that stub zones improve name resolution and eliminate the burden on network resources that would otherwise result from large zone transfers.



EXAM TIP

Expect to see a question about stub zones on the 70-642 exam. Understand that you can use them to keep an updated list of name servers in a remote zone and to improve name resolution across domains.

✓ Quick Check

1. True or False: You can perform a delegation only from a parent zone to a child zone.
2. Why does a stub zone improve name resolution when it is implemented across domains in a large forest or other DNS namespace?

Quick Check Answers

1. True.
2. A stub zone provides a DNS server with the names of servers that are authoritative for a given zone. When this information is stored locally, the DNS server does not need to query other servers to find the authoritative servers for that zone. The process of resolving a name in that zone is therefore more efficient.

PRACTICE Creating an Application Directory Partition for DNS

In this practice, you will create a custom application directory partition and then modify the Nwtraders.msft zone to store data in that partition. (Note that zone data can be stored only in directory partitions for Active Directory–integrated zones.)

EXERCISE 1 Creating the New Application Directory Partition

In this exercise, you will create an application directory partition on Dcsrv1.

1. Log on to Nwtraders from Dcsrv1 as a domain administrator.
2. At an elevated command prompt, type the following:

```
dnscmd . /createdirectorypartition DNSpartitionA.nwtraders.msft
```

This command creates an application directory partition that will replicate in Active Directory only to domain controllers that you enlist in the partition. You do not need to enlist the local server in the partition.

EXERCISE 2 Storing Zone Data in the New Application Directory Partition

In this exercise, you will modify the properties of the Nwtraders.msft zone so that its data is stored in the new application directory partition you have just created.

1. While you are logged on to Nwtraders from Dcsrv1 as a domain administrator, open DNS Manager.
2. In the DNS Manager console tree, expand the Forward Lookup Zones folder, select and then right-click the Nwtraders.msft zone, and then choose Properties.

3. On the General tab of the Nwtraders.msft Properties dialog box, click the Change button for replication. This button is directly to the right of the text Replication: All DNS Servers In This Domain.
4. In the Change Zone Replication Scope dialog box that opens, select To All Domain Controllers In The Scope Of This Directory Partition.
5. In the associated drop-down list box, select DNSpartitionA.nwtraders.msft, and then click OK.
6. In the Nwtraders.msft Properties dialog box, click OK. The Nwtraders.msft zone data is now stored in the new application directory partition you created on Dcsrv1. Other domain controllers that are DNS servers in the Nwtraders.msft forest will receive a copy of the Nwtraders.msft primary zone only if you later enlist those servers in the new partition by using the following command:

```
dnscmd <server name> /enlistdirectorypartition DNSpartitionA.nwtraders.msft
```

PRACTICE Deploying a Secondary Zone

In this practice, you will create a secondary DNS zone for Nwtraders.msft on the Boston server. Because the Boston server is not a domain controller, it cannot host an Active Directory–integrated copy of the Nwtraders.msft primary zone. In a production environment, you might choose to install a secondary zone when you want to install a DNS server without installing a domain controller.

EXERCISE 1 Adding the DNS Server Role

In this exercise, you will install the DNS server role on the Boston server.

1. Log on to Nwtraders from Boston as a domain administrator.
2. If the Initial Configuration Tasks window appears, click Add Roles. Otherwise, open Server Manager and click Add Roles in the details pane.
3. On the Before You Begin page of the Add Roles Wizard, click Next.
4. On the Select Server Roles page, select the DNS Server check box, and then click Next.
5. On the DNS Server page, read all the text, and then click Next.
6. On the Confirm Installation Selections page, click Install.
7. After the installation completes, on the Installation Results page, click Close.

EXERCISE 2 Creating the Secondary Zone

In this exercise, you will create a secondary zone named Nwtraders.msft on Boston.nwtraders.msft.

1. While you are still logged on to Nwtraders from Boston as a domain administrator, open DNS Manager.
2. Expand the DNS Manager console tree.

3. In the DNS Manager console tree, select and then right-click the Forward Lookup Zones folder, and then choose New Zone. The Welcome page of the New Zone Wizard appears. Click Next.
4. On the Zone Type page, read all the text, and then select Secondary Zone. Note that the option to store the zone in Active Directory is dimmed. This choice is unavailable because the local computer is not a domain controller. Click Next.
5. On the Zone Name page, in the Zone Name text box, type **nwtraders.msft**. Click Next.
6. On the Master DNS Servers page, read the text on the page.
7. In the Master Servers area, type **192.168.0.1**, and then press Enter.
8. Wait about 30 seconds for the name DCSRVR1 to appear beneath the Server FQDN heading in the Master Servers area. Click Next.
9. On the Completing The New Zone Wizard page, click Finish. The new zone now appears in DNS Manager.
10. In the DNS Manager console tree, select the Nwtraders.msft forward lookup zone. An error message that appears in the details pane indicates that the zone is not loaded by the DNS server. The problem is that you have not enabled zone transfers in the properties of the primary zone on Dcsrv1.

EXERCISE 3 Enabling Zone Transfers to the Secondary Zone

In this exercise, you will enable zone transfers to the Boston computer from Dcsrv1.

1. Log on to Nwtraders from Dcsrv1 as a domain administrator.
2. Open DNS Manager.
3. Expand the DNS Manager console tree.
4. Right-click the Nwtraders.msft forward lookup zone, and then choose Properties.
5. In the Nwtraders.msft Properties dialog box, click the Zone Transfers tab.
6. On the Zone Transfers tab, select the Allow Zone Transfers check box.
7. Verify that To Any Server is selected, and then click OK.

EXERCISE 4 Transferring the Zone Data

In this exercise, you will load the zone data from the primary zone to the secondary zone. You will perform this exercise while logged on to Nwtraders from the Boston computer as a domain administrator.

1. On Boston, in the DNS Manager console tree, right-click the Nwtraders.msft forward lookup zone, and then choose Transfer From Master. If you see an error, wait 15 seconds, and then press F5 or select Refresh from the Action menu.
2. The Nwtraders.msft zone data eventually appears in the details pane of DNS Manager. Note that the application directory partition DNSpartitionA appears above DomainDNSZones and ForestDNSZones.

EXERCISE 5 Creating an NS Record for the Server Hosting the Secondary Zone

In this exercise, you will create an NS record for the Boston DNS server in the primary zone. Note that you cannot create an NS record for a secondary zone server from within the secondary zone itself because a secondary zone is a read-only copy of the zone.

You perform this exercise while logged on to Nwtraders from Dcsv1 as a domain administrator.

1. On Dcsv1, in the DNS Manager console tree, select the Nwtraders.msft zone. In the details pane, note that the only name server (NS) record included in the zone points to dcsv1.nwtraders.msft. The fact that there is only one such NS record means that even if the DNS domain were connected to a larger DNS namespace, information about names in the Nwtraders.msft domain will always originate from Dcsv1.
2. In the detail pane, double-click the NS record. The Nwtraders.msft Properties dialog box opens, and the Name Servers tab is selected.
3. Click the Add button.
4. In the New Name Server Record dialog box, in the Server Fully Qualified Domain Name (FQDN) text box, type **boston.nwtraders.msft**, and then click Resolve. The name is resolved to an IPv6 address and an IPv4 address.
5. In the New Name Server Record dialog box, click OK.
6. In the Nwtraders.msft Properties dialog box, click the Zone Transfers tab.
7. Select Only To Servers Listed On The Name Servers Tab. This setting provides security for the zone by restricting copies (transfers) of the zone data to only authorized servers.
8. In the Nwtraders.msft Properties dialog box, click OK. In the details pane of DNS Manager, a new NS record appears that points to boston.nwtraders.msft.
9. Close all windows and log off both servers.

Lesson Summary

- Zone replication refers to the synchronization of zone data for Active Directory–integrated zones. Zone transfers refer to the synchronization of zone data between any master and a secondary standard zone.
- A partition is a data structure in Active Directory that distinguishes data for different replication purposes. By default, domain controllers include two application directory partitions reserved for DNS data: DomainDnsZones and ForestDnsZones. The DomainDnsZones partition is replicated among all domain controllers that are also DNS servers in a particular domain, and the ForestDnsZones partition is replicated among all domain controllers that are also DNS servers in every domain in an Active Directory forest.
- You can also create a user-defined directory partition with a name of your choice. You can then configure a zone to be stored in this new structure that you have created.

- The partition in which a zone is stored effectively determines the replication scope for that zone.
- Zone transfers are essentially pull operations initiated on secondary zones that copy zone data from a master zone, which itself can be a primary zone or another secondary zone. By default, zone transfers are disabled from any zone and you must enable them on the Zone Transfers tab of the zone properties dialog box.
- A delegation is a subdomain of a parent zone that has been created as a new, separately administered zone and for which minimal records still appear in the parent zone. A delegation enables a parent zone to direct queries intended for the subdomain to the authoritative servers of that subdomain.
- You can use stub zones to keep delegated zone information current or to improve name resolution across domains in a large DNS namespace.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 2. The questions are also available on the companion CD if you prefer to review them in electronic form.

NOTE ANSWERS

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the “Answers” section at the end of the book.

1. You are a network administrator for a large company named Northwind Traders that has many branch offices worldwide. You work at the New York office, which has its own Active Directory domain, `ny.us.nwtraders.msft`.

Recently you have noticed that when users in the New York office want to connect to resources located in the `uk.eu.nwtraders.msft` domain, name resolution for computer names in the remote domain is very slow. You want to improve name resolution response times for names within `uk.eu.nwtraders.msft` domain by keeping an updated list of remote name servers authoritative for that domain name. You also want to minimize zone transfer traffic. What should you do?
 - A. Create a stub zone of the `uk.eu.nwtraders.msft` domain on the DNS servers at the New York office.
 - B. Configure conditional forwarding so that queries for names within the `uk.eu.nwtraders.msft` domain are automatically forwarded to the name servers in that domain.
 - C. Create a secondary zone of the `uk.eu.nwtraders.msft` domain on the DNS servers at the New York office.
 - D. Perform a delegation of the `uk.eu.nwtraders.msft` domain on the DNS servers at the New York office.

2. You recently migrated a DNS zone named Contoso.com to a domain controller running Windows Server 2008 R2. You have selected the option to store the zone in Active Directory, but you find that the zone does not appear on a domain controller named DC2000 that is running Windows 2000 Server in the same domain. DC2000 is already configured with the DNS server component.

You want the zone to appear on the DC2000 domain controller in the Contoso.com domain. What should you do?

- A. Choose the option to store the zone in all DNS servers in the forest.
 - B. Choose the option to store the zone in all DNS servers in the domain.
 - C. Choose the option to store the zone in all domain controllers in the domain.
 - D. Create a new directory partition, and then choose the option to store the zone in the new partition.
3. The server ns1.contoso.com acts as the DNS server in the contoso.com domain. You create a new secondary zone for the domain on a server named ns2.contoso.com and specify ns1.contoso.com as the master. However, whenever you attempt to populate the zone in DNS Manager by selecting Transfer From Master from the Action menu, you see only a red "x" mark on the zone folder, and no data appears.

You want to populate the secondary zone with data from the master zone. What should you do?

- A. Select Reload from the Action menu in DNS Manager.
- B. Select Transfer New Copy Of Zone From Master from the Action menu in DNS Manager.
- C. Configure ns1.contoso.com to allow zone transfers.
- D. Configure ns1.contoso.com to perform notifications, and add ns2.contoso.com to the list of servers to notify.

Lesson 3: Implementing DNSSEC

DNS does not strongly validate the source of the information received through its queries. As a result, attackers can use methods such as DNS cache poisoning to provide forged data to DNS clients and trick these clients into visiting spoofed sites or addresses.

DNSSEC was created to stop the threat of forged DNS data. DNSSEC is an optional DNS server feature that provides digital signatures for its records and validates the signatures received from other DNSSEC-enabled servers. In Windows networks, DNSSEC is used as a server-to-server protocol that validates responses on behalf of Windows 7 clients.

Support for the latest version of DNSSEC is new to Windows Server 2008 R2 and Windows 7.

After this lesson, you will be able to:

- Understand the purpose of DNSSEC.
- Understand how DNSSEC validation is performed.
- Create keys for DNSSEC and sign a zone.
- Configure trust anchors.
- Configure DNS clients to request DNSSEC validation.

Estimated lesson time: 75 minutes

Understanding Public Key Cryptography in DNSSEC

DNSSEC validates information by using public key cryptography. To understand DNSSEC, therefore, you need to review some basic concepts related to this technology, such as key pairs, digital signatures, and trust relationships.

Understanding Key Pairs

Public key cryptography provides *asymmetric encryption*, which means that separate keys are used to encrypt and decrypt data. These separate keys amount to a unique *key pair* generated by each organization that wants to send or receive encrypted data with the public. One of these two keys, the *public key*, is validated and shared freely with the world, but the *private key* is kept secret. Either key may be used to encrypt or decrypt data. However, as illustrated in Figure 3-35, data encrypted by the public key can be decrypted only by the corresponding private key, and data encrypted by the private key can be decrypted only by the corresponding public key.

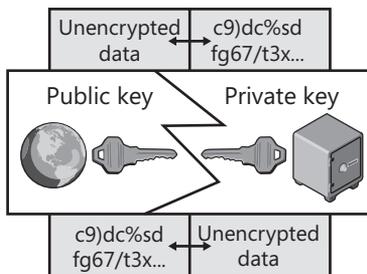


FIGURE 3-35 A key pair is used to encrypt and decrypt data.

In Windows Server 2008 R2, key pairs are generated by using the `Dnscmd` command in a manner described in the section “Generating Key Pairs” later in this lesson. After you run the command, public keys needed for DNSSEC are made available to the world in DNS zones and are stored in a resource record type called a DNSKEY. The private keys are stored in the local certificate store on the server on which the command is run, but you can export these private keys to a disk or another safe location.

Understanding Digital Signatures

A *digital signature* is a version of some unit of data, such as a specific file, that an organization has encrypted with its private key. This signature is then delivered to another party along with the original unencrypted version of the same data. The receiving party decrypts the encrypted data by using the sender's validated public key and compares the result to the original unencrypted data. If the two data sets match, the data is authenticated as truly originating from that organization. The signature check also effectively ensures that the data is unmodified from the original version that was signed. In this way, digital signatures use public key cryptography to prove that information is unspoofed and unchanged.

The procedure of creating and validating a digital signature is illustrated by an example shown in Figure 3-36. In the figure, Contoso.com creates a signature of a file by encrypting the file with its private key. The original file and the signature are then sent to another party, Nwtraders.com, who needs to verify that the file is authentic and unchanged. Nwtraders.com uses the validated public key for Contoso.com to decrypt the signature. If the decrypted signature matches the original file exactly, the file is effectively authenticated.

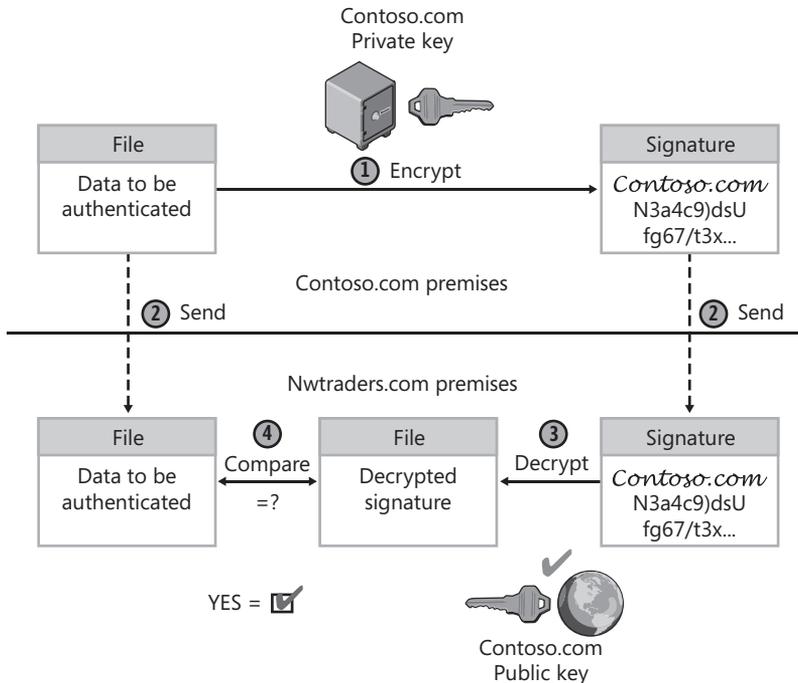


FIGURE 3-36 Creation and verification of a digital signature

Understanding DNSSEC Trust Relationships

DNSSEC in Windows Server 2008 R2 allows a DNS server to validate DNS data on behalf of its Windows 7 clients. When a DNSSEC-enabled DNS server receives a query that it cannot answer immediately, it first contacts other DNS servers to obtain the queried-for data and then verifies the digital signatures associated with that data before responding to the client.

However, a digital signature alone is not sufficient to authenticate data to other DNS servers. A spoofing DNS server could use its own private key to sign false records and present public keys claiming to belong to another source. In order for the data to be truly validated, therefore, the public key itself needs to be validated by other DNS servers through a trust relationship, either directly or indirectly.

Trust anchors are used with DNSSEC to establish trust relationships between DNS servers. A *trust anchor* is a public key for a remote DNS server that is trusted and able to provide DNSSEC responses.

When you configure a local DNS server with a trust anchor for a remote DNS server, the local DNS server is able to validate the DNS information for the zone(s) for which that remote server is authoritative. In addition, through a chain of trust relationships, the local server is also able to validate the data from any delegated DNS subdomains that also can provide DNSSEC responses.

Figure 3-37 illustrates how trust anchors extend DNSSEC validation. The DNS server authoritative for `nwtraders.com` is configured with a trust anchor for the DNS server authoritative for `contoso.com`. This latter DNS server has two delegated subdomains, one of which, `asia.contoso.com`, is also configured with DNSSEC. As a result of these configured trust relationships, the DNS server at `nwtraders.com` is able to provide DNSSEC validation to its clients for queries related to both the `contoso.com` domain and the `asia.contoso.com` domain. Validated responses are not possible, however, for the `uk.contoso.com` domain.

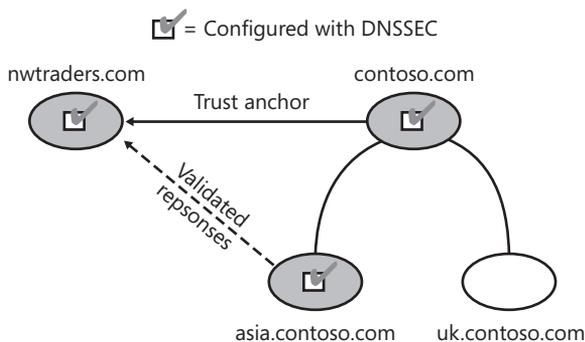


FIGURE 3-37 Trust anchors extend DNSSEC validation to remote domains and their DNSSEC-compatible subdomains.

Understanding DNSSEC Name Resolution

In Microsoft's implementation of DNSSEC, a client does not validate DNS responses received from its local server. Instead, Windows 7 clients request DNSSEC from their local Windows Server 2008 R2 server, which then uses DNSSEC to validate only the responses obtained from other servers.

Windows 7 clients can be configured to request DNSSEC through the Name Resolution Policy Table (NRPT) in Group Policy. The NRPT allows you to specify the DNS query suffixes, prefixes, FQDNs, or reverse lookup subnets for which a Windows 7 or Windows Server 2008 R2 client will request DNSSEC. (Note that you can also use the NRPT to enforce IPsec between the local DNS client and the local DNS server; this configuration essentially authenticates the local DNS server and prevents man-in-the-middle attacks.)

The next series of figures illustrates an example of a name resolution procedure with DNSSEC. In Figure 3-38, a Windows 7 client named `client1.nwtraders.com` needs to query its DNS server for the name `www.ny.contoso.com`. DNSSEC in this example validates four elements in a row: the zone delegation from `contoso.com` to `ny.contoso.com`, the Key Signing Key from `ny.contoso.com`, the Zone Signing Key from `ny.contoso.com`, and finally the host record for `www.ny.contoso.com`.

In step 1, the client first checks its NRPT to determine whether this query should be performed with a request for DNSSEC. The NRPT includes an entry for the suffix `contoso.com`, so in step 2, the client queries the local server at `ns1.nwtraders.com` for the name `www.ny.contoso.com` with a request for DNSSEC validation.

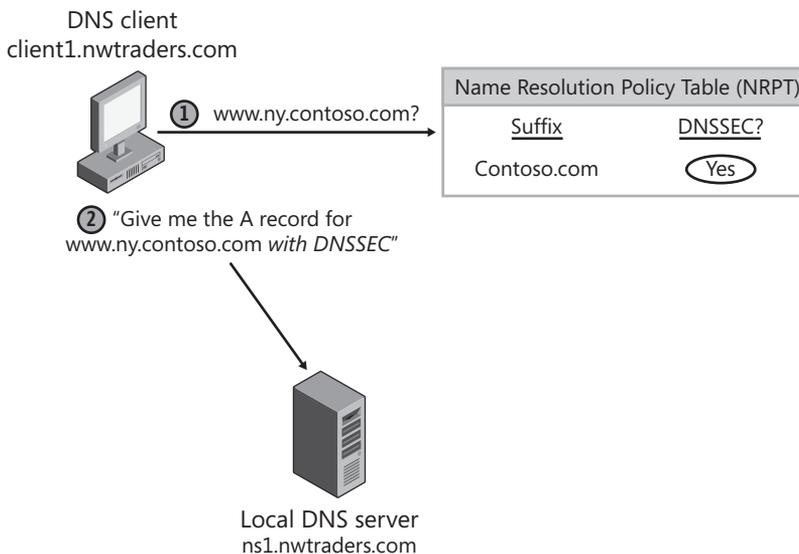


FIGURE 3-38 DNSSEC validation begins when a Windows 7 client checks its NRPT before querying the local DNS server.

In Figure 3-39, the local DNS server at ns1.nwtraders.com confirms in step 3 that it has a trust anchor for contoso.com, which is a parent DNS domain of the target domain ny.contoso.com. Then, in step 4, the local DNS server forwards the query to the server authoritative for contoso.com, the address of which has been configured as a conditional forwarder for that domain.

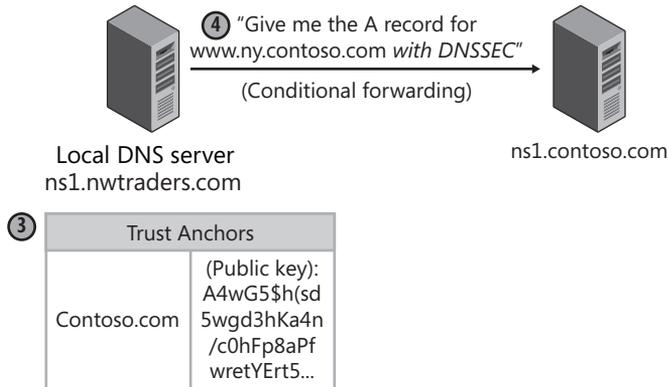


FIGURE 3-39 DNSSEC requires the local DNS server to have a trust anchor configured for the queried-for domain or parent domain.

In step 5, shown in Figure 3-40, ns1.contoso.com responds to ns1.nwtraders.com with information about the DNS server authoritative for the delegated subdomain ny.contoso.com. As with a normal DNS query, this information is contained in the NS and A records.

In addition, ns1.contoso.com provides two additional records for DNSSEC whose purpose is to validate the DNS delegation: a Delegation Signer (DS) record and a Resource Record Signature (RRSIG) record. DS records appear in parent DNS domains and include a SHA-1 or SHA-256 hash of the public key used in a delegated subdomain that is also DNSSEC-compatible. In this case, the DS record includes a hash of a particular public key called a Key Signing Key (KSK) from ny.contoso.com. (This hash will later be used to authenticate the KSK obtained directly from the ny.contoso.com domain.) The RRSIG record is a digital signature of this latter DS record. It is signed by contoso.com to enable others to verify that the DS record is authentic and unchanged.

NOTE WHAT ARE THE KSK AND ZSK?

Zones that are digitally signed typically use two separate key pairs and, therefore, two public keys. The KSK is updated rarely and is stored in other zones as a DS record or trust anchor. The second public key is called a Zone Signing Key (ZSK). The ZSK is updated frequently and is stored only in the native zone. The KSK is used to validate the ZSK, and the ZSK is used to validate the records in a zone.

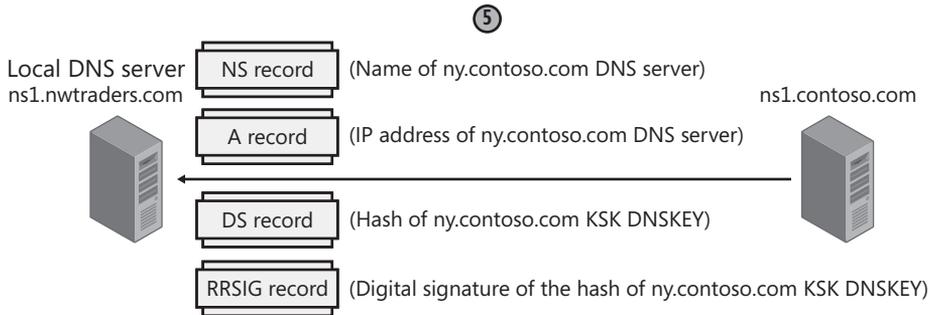


FIGURE 3-40 DNS servers configured with DNSSEC validate delegations with DS and RRSIG records.

Figure 3-41 shows the procedure for verifying the digital signature of the DS record obtained in step 5. In step 6, the public key stored as a trust anchor on ns1.nwtraders.com is used to decrypt the RRSIG record obtained with the DS record. Then, in step 7, this decrypted RRSIG record is compared to the DS record. If the two are identical, the DS record is validated.

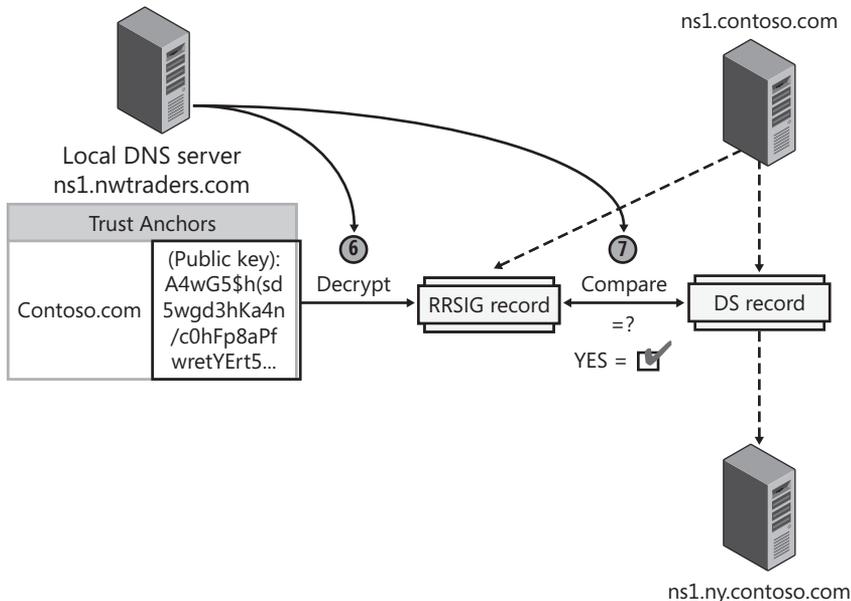


FIGURE 3-41 The trust anchor is used to validate the digital signature of the DS record.

After the DS record is validated, the local DNS server contacts the server authoritative for the ny.contoso.com domain. As shown in step 8 of Figure 3-42, the local DNS server requests the KSK for ny.contoso.com, a hash of which is contained in the just-validated DS record. In step 9, ns1.ny.contoso.com provides the requested KSK in a DNSKEY record. (DNSKEY records always contain public keys.)

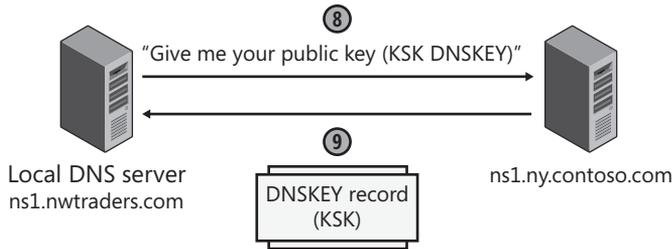


FIGURE 3-42 The local DNS server obtains a public key in a DNSKEY record from the target zone.

Figure 3-43 shows the process of validating the public key obtained in step 9. In step 10, the local DNS server applies the same SHA-1 or SHA-256 hashing algorithm to the just-obtained KSK from ns1.ny.contoso.com. In step 11, the local DNS server compares the hash of the DNSKEY generated by the procedure in step 10 to the DS record obtained in step 5 from ns1.contoso.com. If the two hashes match, the KSK from ny.contoso.com is validated.

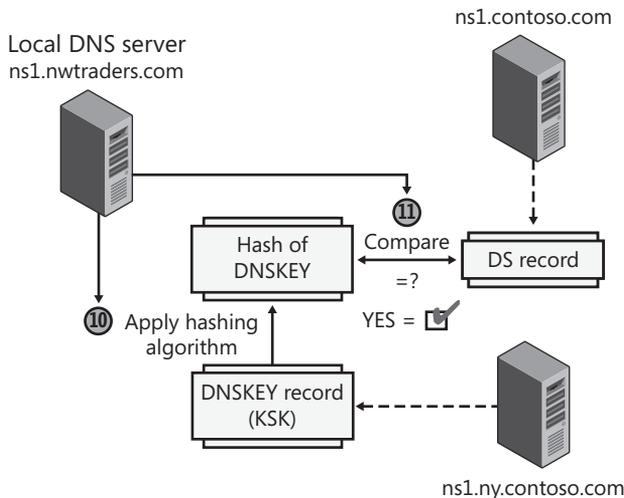


FIGURE 3-43 The KSK is validated by comparing hashes.

Figure 3-44 shows how, in step 12, the local DNS server next contacts ns1.ny.contoso.com to request its ZSK public key. In step 13, the remote DNS server responds with the ZSK and an accompanying RRSIG record, or digital signature.

Figure 3-45 illustrates the process of verifying the digital signature for the ZSK obtained in this last step. In step 14, the KSK obtained and validated in steps 9–11 is now used to unlock the digital signature of the ZSK. The decrypted digital signature in step 15 is then compared to the DNSKEY record that includes the ZSK. If the two match, the ZSK is validated.

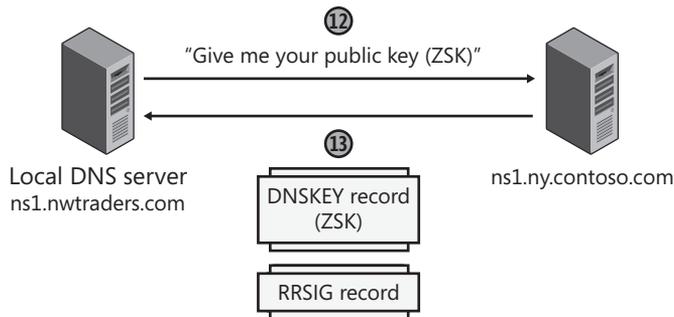


FIGURE 3-44 The local server obtains the remote zone's ZSK and a digital signature used to validate it.

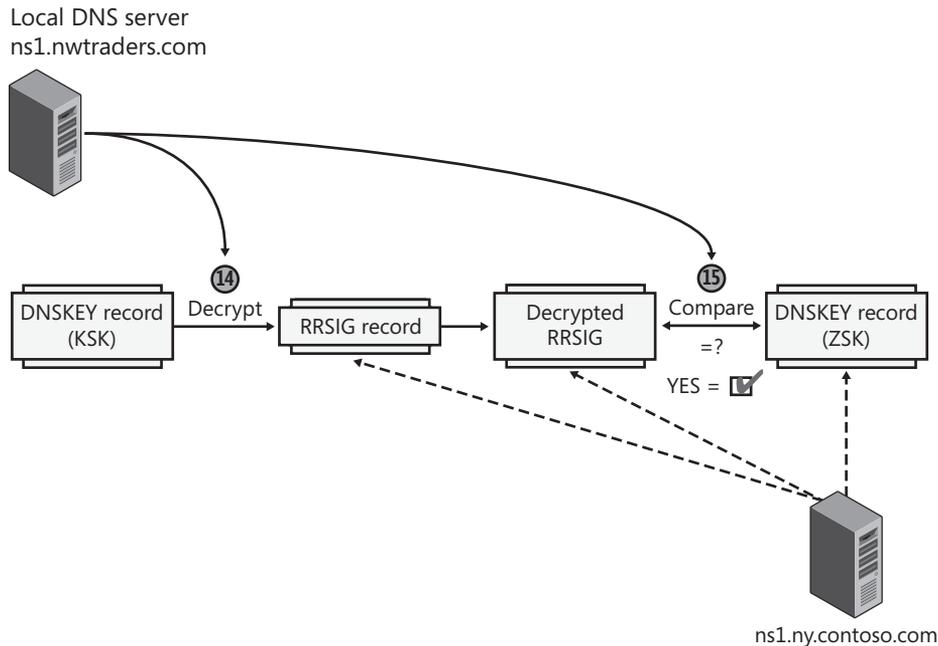


FIGURE 3-45 The KSK public key is used to validate the ZSK public key.

Now that the ZSK has been obtained, the local DNS server in step 16 can query ns1.ny.contoso.com for the name `www.ny.contoso.com` with DNSSEC, as shown in Figure 3-46. In step 17, the response is received in the form of an A record and accompanying digital signature in an RRSIG record. Note that even if the DNS responded with many matching A records, still only one RRSIG record would be included in the response because an RRSIG record is a digital signature of the *set* of DNS resource records that can match a query.

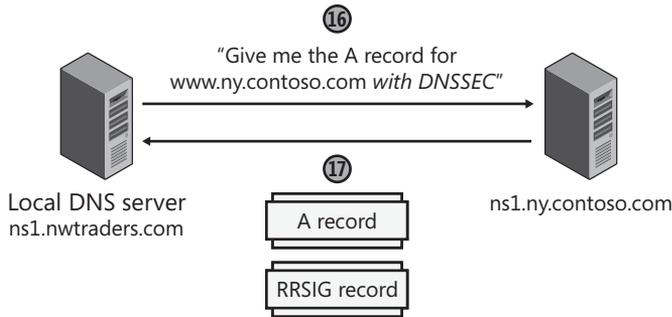


FIGURE 3-46 A DNSSEC-enabled DNS server responds to queries with a digital signature.

Next, in Figure 3-47, the local DNS server can validate the A record received in step 17, which is the response to the original query. In step 18, the now-validated ZSK public key is used to decrypt the RRSIG of the A record received in step 17. Then, in step 19, this decrypted RRSIG is compared to the A record received in step 17. If the two match, the A record is validated.

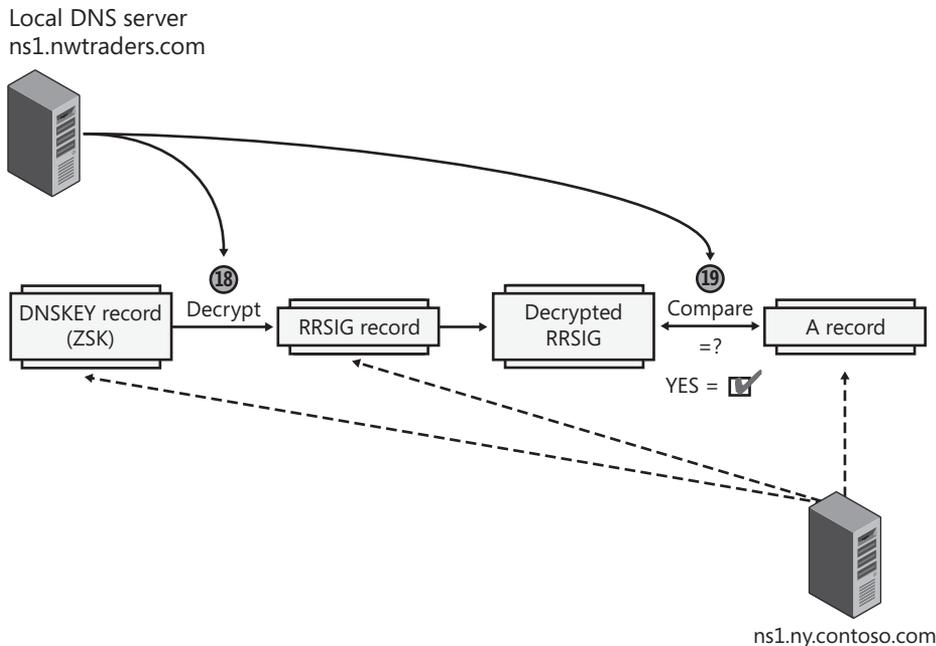
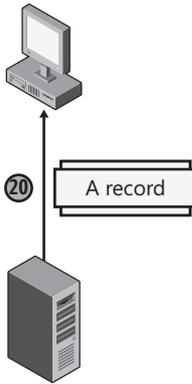


FIGURE 3-47 The digital signature of the A record is checked by using the validated ZSK from the remote zone.

Step 19 marks the end of the validation process. Now that the client's original query has been answered for the local DNS server with DNSSEC validation, that local DNS server in step 20 finally responds to the DNS client with an A record for www.ny.contoso.com. This step is shown in Figure 3-48.



Local DNS server
ns1.nwtraders.com

FIGURE 3-48 After DNSSEC validation, the DNS server responds to the client with an A record.

NOTE WHAT IS A NEXT SECURE (NSEC) RECORD?

The earlier description of DNSSEC name validation applies to queried-for names that have a matching A record in the authoritative DNS server. If the queried-for name had no matching A record, however, the authoritative server would have responded with an NSEC record and accompanying RRSIG. An NSEC record is used with DNSSEC to prove that a queried-for name does not exist in a zone.

As a zone is signed, all the records in the zone are ordered alphabetically, and the record corresponding to the name of the zone is placed at the end of the list. Then, an NSEC record is created for every gap between successive records. For example, in a zone named Contoso.com containing only three records, named Boston, Phoenix, and the domain name record "Contoso.com," three NSEC records would be created: the first to prove the absence of names between Boston and Phoenix, the second to prove the absence of names between Phoenix and "Contoso.com," and the third to prove the absence of names between "Contoso.com" and Boston.

When the server receives a query for a name that doesn't exist in the zone, it responds with the NSEC record that shows that no name exists in the corresponding alphabetical range.

✓ Quick Check

- What is the name of the record that contains a hash of the public key in a delegated subdomain?

Quick Check Answer

- A DS record

Configuring DNSSEC

For DNSSEC to work, DNSSEC has to be configured on the local DNS server that receives the client's original query and on every remote DNS server that the local DNS server contacts to resolve the query. In addition, the DNS client needs to be configured through Group Policy to request DNSSEC for the specific name queried.

NOTE DNSSEC IN WINDOWS SERVER 2008 AND WINDOWS SERVER 2008 R2

Windows Server 2003 and Windows Server 2008 implementations of DNSSEC are not interoperable with the Windows Server 2008 R2 or Windows 7 implementation. In Windows Server 2003 and Windows Server 2008, DNSSEC is implemented on secondary zones as described in RFC 2535, which is now obsolete. DNSSEC on Windows Server 2008 R2 and Windows 7 are based on RFC 4033, RFC 4034, and RFC 4035.

Configuring DNSSEC in Windows Server 2008 R2

Configuring DNSSEC on a DNS server requires you to prepare the zone to be signed; devise a rollover mechanism for your keys; create the keys; sign the zone; and, finally, configure trust anchors.

PREPARING A ZONE FOR DNSSEC

The first step in preparing a zone for DNSSEC is to back up the current zone data. To back up a zone, type the following command at an elevated command prompt:

```
dnscmd /ZoneExport <zone name> <zone file name>
```

This command exports the zone data to a file in the %windir%\System32\DNS directory.

Next, disable dynamic updates for the zone because DNSSEC works only with static zones. (If you make a change to a zone, you have to resign the zone.)

Finally, make the zone a standard zone. Although DNSSEC is compatible with AD-integrated zones, the process of signing a zone requires that the zone data be stored in a zone file. Therefore, if your zone is AD-integrated, you should temporarily configure it as a standard zone. After you have signed the zone, you can reconfigure it as an AD-integrated zone.

CHOOSING A KEY ROLLOVER MECHANISM

DNSSEC keys do not have a permanent lifetime, so you need to plan a method to regularly generate new keys, resign the zone, and distribute the new keys for trust anchors and delegations. This process of updating keys and digital signatures is also called *key rollover*. Note that to facilitate rollover, DNSSEC allows multiple ZSKs and KSKs in a zone at the same time. When other DNS servers detect multiple public keys in the zone, each public key is used to attempt the validation of signatures.

The following two methods can be used for key rollover:

- **Prepublished rollover** With this method, you introduce a new, second key to the zone before using it. After introducing the second key, you resign the zone with the old key to account for the presence of a new record (the new DNSKEY) in the zone. Then, when the old zone data has expired from remote server caches, you resign the zone with the new key. Resigning a zone with a single new key automatically deletes the old signatures in the zone. This type of rollover method is typically best suited for the ZSK.
- **Double signature rollover** With this method, you add a new, second key to the zone and resign the zone with both the old and new keys at the same time. Then, when the old zone data has expired from remote server caches and any necessary updates are made at parent domain servers or remote servers using the key as a trust anchor, you resign the zone again with only the new key. This type of rollover method is typically best suited for the KSK.

MORE INFO PERFORMING KEY ROLLOVERS

For more information about performing a key rollover, search for the “DNSSEC Deployment Guide” on the Microsoft website.

GENERATING KEY PAIRS

The purpose of using two keys with a zone is to give a long validity period to one key (KSK) that is distributed to other zones in trust anchors and delegations. If you used a long validity period with the same key to sign the zone records, however, security could be compromised. The presence of a KSK allows the key that is signing the zone (the ZSK) to be frequently updated without burdening administrators with the frequent task of distributing new keys to other zones. Note that the use of the KSK is recommended but is not required. When a single key is used both to sign the zone records and act as a trust anchor on remote servers, that key is a ZSK.

Use the following procedures to generate the KSK and ZSK key pairs. Once created, the keys are stored in a self-signed certificate in the local computer certificate store, in the MS-DNSSEC container.

To generate a KSK, follow these steps:

1. Open an elevated command prompt.
2. Type the following command:

```
DnsCmd /OfflineSign /GenKey /Alg rsasha1 /Flags KSK /Length <length> /Zone <zone name> /SSCert /FriendlyName KSK-<zone name>
```

To generate a ZSK, follow these steps:

1. Open an elevated command prompt.
2. Type the following command:

```
DnsCmd /OfflineSign /GenKey /Alg rsasha1 /Length <length> /Zone <zone name> /SSCert/FriendlyName ZSK-<zone name>
```

Table 3-1 describes the switches and options used with key generation and the Dnscmd. The values are listed in the table in the order in which you would type them in a command.

TABLE 3-1 Dnscmd Values for DNSSEC Key Generation

VALUE	DESCRIPTION
/OfflineSign	Required. Used with the GenKey, DeleteKey, ImportKey, or SignZone commands to modify certificates and keys or to sign a zone file.
/GenKey	Required. Generates a self-signed certificate with a private key.
/Alg	Required. Used with rsasha1 to specify the algorithm of the signing key. Currently, only RSA/SHA-1 is supported.
rsasha1	Required. Specifies the RSA/SHA-1 algorithm is used for the signing key.
/Flags	Used with KSK to specify the flags in DNSKEY. Currently, only KSK is supported, which indicates that the Zone Key bit and the Secure Entry Point bit are turned on. If /flags is not specified, only the Zone Key bit is turned on, which indicates a zone signing key.
KSK	Specifies that the KSK flag in DNSKEY is set.
/Length	Required. Used with <length> to specify the number of bits used in the key.
<length>	Required. Numerical value of bits used in the key. The allowed values for length are from 512 bits through 4096 bits, in 64 bit increments.
/Zone	Required. Used with <zone name> to specify the fully qualified domain name (FQDN) of the zone.
<zone name>	Required. The FQDN of the zone.
/SSCert	Required. Specifies that the key will be stored in a self-signed certificate.
/FriendlyName	Used with KSK-<zone name> or ZSK-<zone name> to specify the friendly name of the self-signed certificate.
KSK-<zone name>	Specifies the friendly name of the self-signed certificate used with a KSK.
ZSK-<zone name>	Specifies the friendly name of the self-signed certificate used with a ZSK.
/ValidFrom	Optional. Used with <validfromtime> to specify the start time for the validity period of the certificate. If not specified, the default will be Current time minus 1 hour.

VALUE	DESCRIPTION
<validfromtime>	Optional. Specifies the local start time for the validity period of the certificate. The required format is YYYYMMDDHHMMSS.
/ValidTo	Optional. Used with <validtotime> to specify the end time for the validity period of the certificate. If not specified, the certificate will be valid for 5 years.
<validtotime>	Optional. Specifies the local end time for the validity period of the certificate. The required format is YYYYMMDDHHMMSS.

After you generate the keys, you may choose to back up the associated certificates using the Certificates snap-in. In the console tree, navigate to Certificates\MS-DNSSEC\Certificates. Then, in the details pane, right-click each certificate and choose Export from the shortcut menu, as shown in Figure 3-49. When exporting, note that the certificate labeled “257” corresponds to the KSK, and the certificate labeled “256” corresponds to the ZSK.

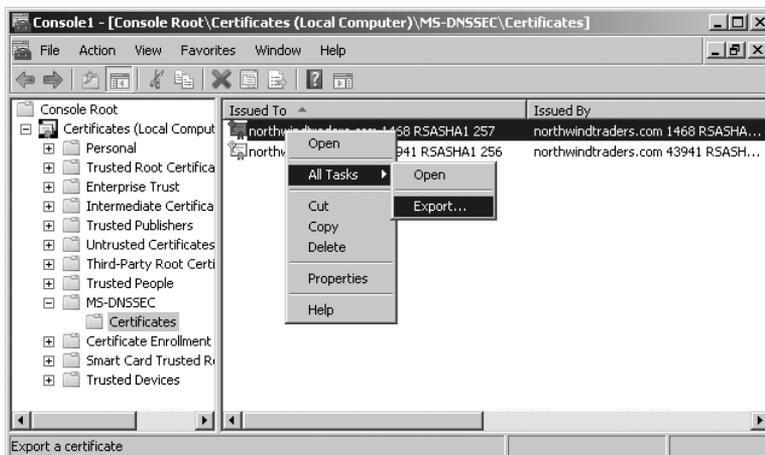


FIGURE 3-49 Backing up a KSK

SIGNING A ZONE FILE

After the keys are generated, you can sign a zone file. If the zone is Active Directory–integrated, you should first export the zone to a file. To sign a zone file, open an elevated command prompt and browse to the %windir%\System32\DNS directory.

Type the following command, and then press Enter:

```
DnsCmd /OfflineSign /SignZone /input <input zone file> /output <output zone file>
/zone <zone name> /signkey /ValidTo <validtodate> /ValidFrom <validfromdate> /cert
/friendlyname ksk-<zone name> /signkey /cert /friendlyname zsk-<zone name>
```

Table 3-2 describes the particular values used in signing a file with the Dnscmd command.

TABLE 3-2 Dnscmd Values for Zone Signing

VALUE	DESCRIPTION
/SignZone	Required. Used to sign a zone file.
/input	Required. Used with <i><input filename></i> to designate the zone file to be signed.
<i><input filename></i>	Required. The name of the zone file to be signed.
/output	Required. Used with <i><output filename></i> to designate the name of the zone file after it has been signed.
<i><output filename></i>	Required. The file name of the signed zone.
/Zone	Required. Used with <i><zone name></i> to specify the fully qualified domain name (FQDN) of the zone.
<i><zone name></i>	Required. The FQDN of the zone.
/Signkey	Required. Specifies the key that will be used to sign the zone.
/ValidFrom	Optional. Used with <i><validfromdate></i> to specify the start time of the validity period of RRSIG records created using this key. If not specified, the validity period will start 1 hour prior to the current UTC time.
<i><validfromdate></i>	Optional. Specifies the UTC start time of the validity period in YYYYMMDDHHMMSS format.
/ValidTo	Optional. Used with <i><validtodate></i> to specify the end time of the validity period of RRSIG records created using this key. If not specified, the validity period will end 30 days from the start of the validity period for zone signing keys, or 13 months from the start of the validity period for key signing keys.
<i><validtodate></i>	Optional. Specifies the UTC end time of the validity period in YYYYMMDDHHMMSS format.
/Cert	Required. Specifies that keys are stored in a certificate.

CONFIGURE TRUST ANCHORS

A validating DNS server must be configured with one or more trust anchors from a remote zone in order to perform validation. Note that it is not necessary to configure a trust anchor for a locally-hosted zone (a zone for which the DNS server is authoritative).

If the DNS server is running on a domain controller, trust anchors are stored in AD DS and are replicated to all domain controllers in the forest. On stand-alone DNS servers, trust anchors are stored in a file named TrustAnchors.dns in %windir%\System32\DNS.

To add a trust anchor, perform the following steps:

1. In the DNS console tree, right-click the name of the DNS server and then click Properties.
2. On the Trust Anchors tab, click the Add button shown in Figure 3-50.

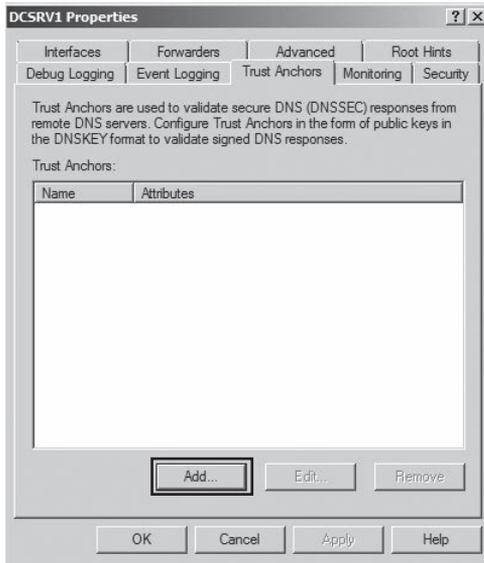


FIGURE 3-50 Adding a trust anchor to a DNS server

3. In the New Trust Anchor dialog box, shown in Figure 3-51, type the name of the signed zone in the Name text box. Do not change the settings for Protocol (DNSSEC) or Algorithm (RSA/SHA-1).

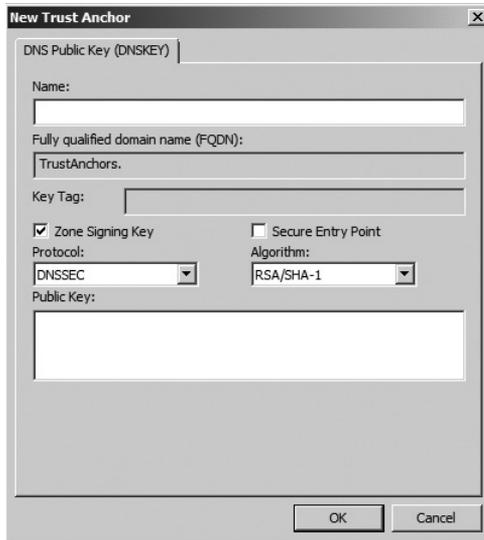


FIGURE 3-51 Configuring a new trust anchor

4. Paste the public key of the signed zone into the Public Key text box. Both the Zone Signing Key and Secure Entry Point check boxes must be selected if the remote anchor is a KSK, which is the typical configuration. However, some zones might use only a single key with the zone and not use a KSK. If the remote zone provides the ZSK as a trust anchor, you should select only the Zone Signing Key check box and leave the Secure Entry Point check box cleared.

Configuring Clients for DNSSEC

You use Group Policy to configure DNS clients to request DNSSEC validation for specific queries. To locate the DNSSEC configuration settings in a GPO, navigate to Computer Configuration \Policies\Windows Settings\Name Resolution Policy, as shown in Figure 3-52.

You can then use this Name Resolution Policy section in a GPO to create rules for the NRPT. The purpose of the NRPT is to provide special instructions for particular DNS client queries. These special instructions can relate to either DNSSEC or DirectAccess, which is a new IPv6-based VPN technology covered in Chapter 7, “Connecting to Networks.”

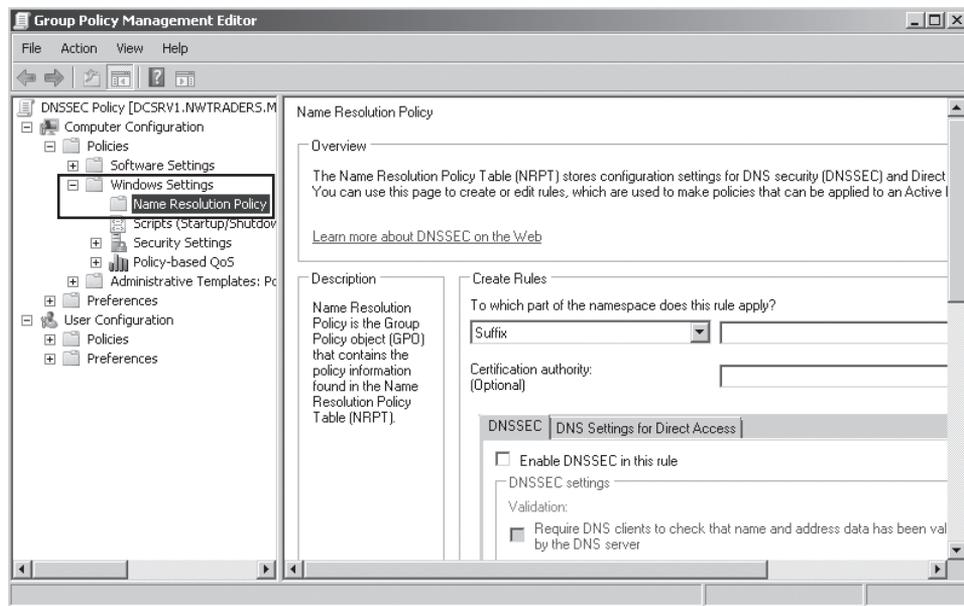


FIGURE 3-52 Configuring a name resolution policy for DNSSEC

To create a name resolution policy rule that enables DNSSEC for a given query, use the Create Rules area in the details pane to specify a name or portion of a name that will match the DNS query. First, in the drop-down list beneath To Which Part Of The Namespace Does This Rule Apply, choose Suffix, Prefix, FQDN, Subnet (IPv4), Subnet (IPv6), or Any, as shown in Figure 3-53.

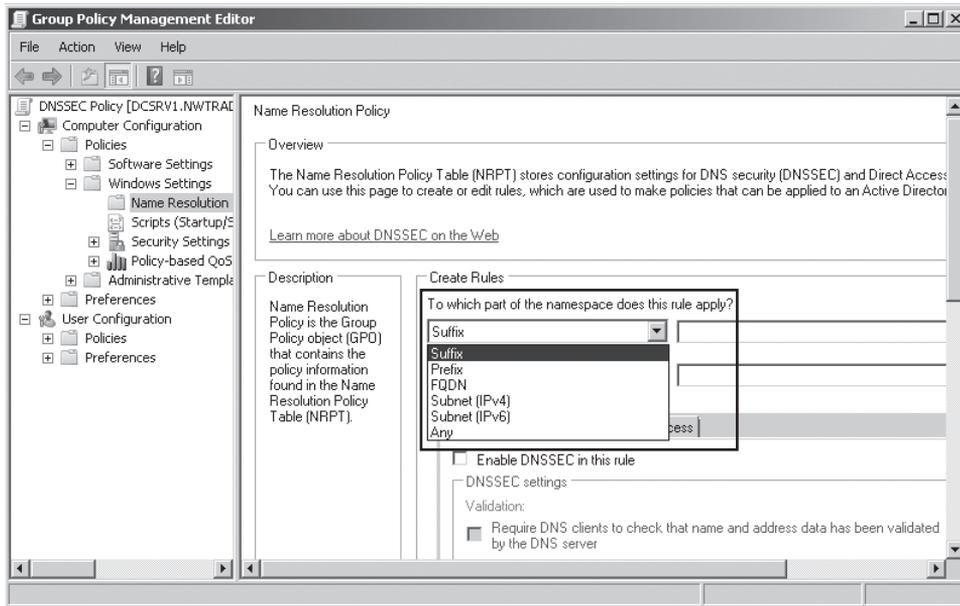


FIGURE 3-53 Choosing at least a portion of a name to match a DNS query

Then, in the text box to the right of the drop-down list, type in the name or portion of a name that represents the query you want to match, and click **Enable DNSSEC In This Rule**. By itself, this step enables DNS clients merely to receive DNSSEC data. If in addition you want to *require* DNS clients to request DNSSEC validation, click **Require DNS Clients To Check That Name And Address Data Has Been Validated**. Finally, you can also enforce IPsec authentication or encryption for the specified query by clicking **Use IPsec In Communication Between DNS Client And DNS Server**. These configuration settings are all indicated in Figure 3-54. After choosing the desired settings, click **Create** to create the rule.

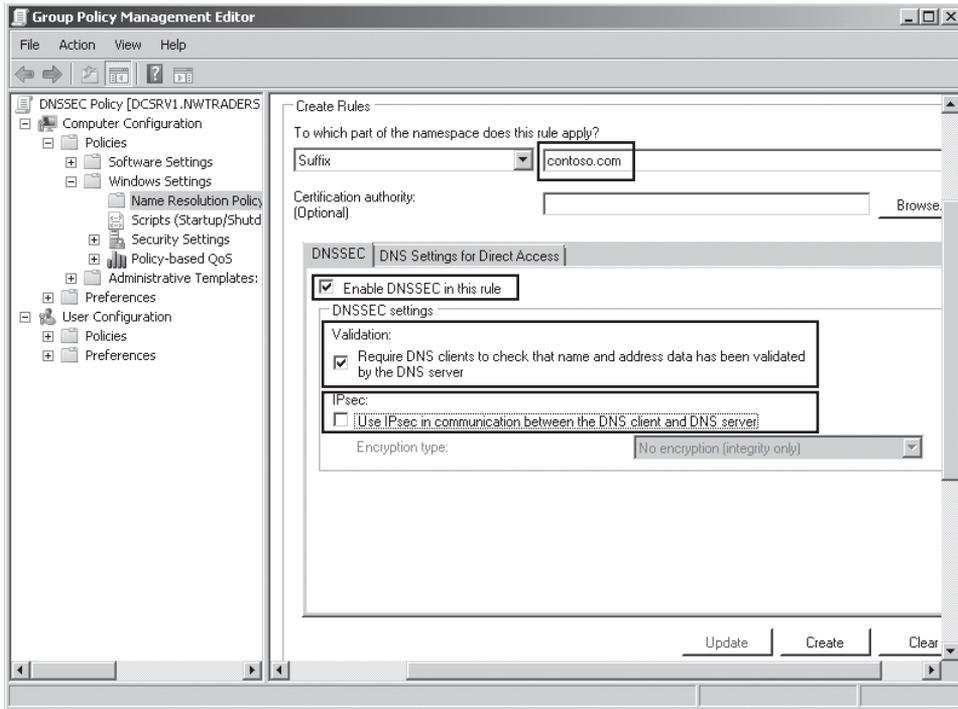


FIGURE 3-54 Adding a DNSSEC rule for the domain contoso.com

After you create the rule, a new entry is added to the NRPT, as shown in Figure 3-55. (The NRPT is visible when you scroll down the details pane in Name Resolution Policy.) Be sure to click Apply to apply the new rule to the policy.

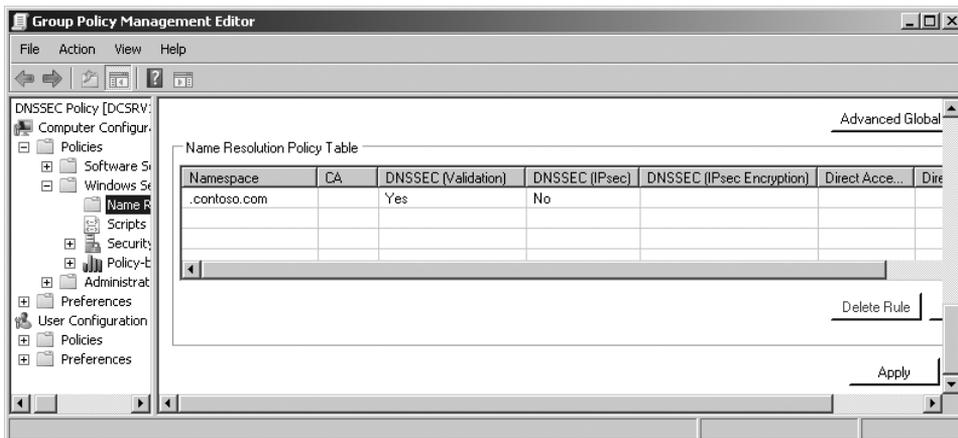


FIGURE 3-55 Use Group Policy and the Name Resolution Policy Table to configure specific DNS queries for DNSSEC.

In this practice, you create and sign a new zone and then configure clients in the domain to request DNSSEC.

EXERCISE 1 Creating a New Static Zone

In this exercise, you will create a new zone whose data is stored in a text file instead of AD DS. The zone will not allow any dynamic updates. You will then create resource records for the Dcsvr1 and Boston computers.

1. Log on to Nwtraders.msft from Dcsvr1 as a domain administrator.
2. Open the DNS console by clicking Start, Administrative Tools, DNS.
3. In the DNS console tree, right-click the Forward Lookup Zones container, and then click New Zone. The New Zone Wizard opens.
4. Click Next on the first page of the New Zone Wizard.
5. On the Zone Type page, clear the Store The Zone In Active Directory check box, and then click Next.
6. On the Zone Name page, type **northwindtraders.com**, and then click Next.
7. On the Zone File page, click Next.
8. On the Dynamic Update page, verify that Do Not Allow Dynamic Updates is selected, and then click Next.
9. On the last page of the New Zone Wizard, click Finish.
10. In the DNS console tree, navigate to and select the Northwindtraders.com container.
11. Right-click the Northwindtraders.com container and then choose New Host (A Or AAAA) from the shortcut menu.
12. In the New Host dialog box, type **dcsvr1** in the Name box and **192.168.0.1** in the IP Address text box, and then click Add Host.
13. In the DNS message box, click OK.
14. In the New Host dialog box, type **boston** in the Name box and **192.168.0.2** in the IP Address text box, and then click Add Host.
15. In the DNS message box, click OK.
16. Click Done to close the New Host dialog box.

EXERCISE 2 Signing Zone Data

In this exercise, you will create the KSK and ZSK keys used to sign the zone. You will then create a signed version of the zone file and re-create the zone using the new signed DNS file.

1. While you are logged on to Nwtraders.msft from Dcsrv1 as a domain administrator, open an elevated command prompt.
2. At the elevated command prompt, type **cd \windows\system32\dns**.
3. Generate the KSK by typing the following command:
dnscmd /OfflineSign /GenKey /Alg rsasha1 /Flags KSK /Length 1024 /Zone northwindtraders.com /SSCert /FriendlyName KSK-northwindtraders.com.
4. Generate the ZSK by typing the following command:
dnscmd /OfflineSign /GenKey /Alg rsasha1 /Length 1024 /Zone northwindtraders.com /SSCert /FriendlyName ZSK-northwindtraders.com.

The new keys are stored in the Local Computer certificate store for the computer account, in the MS-DNSSEC container. The next step is to use the keys to create a signed version of the zone file, northwindtraders.com.dns. You will call the new version signed.northwindtraders.com.dns.

5. Leave the command prompt open, and then open the DNS console. In the DNS console tree, right-click the northwindtraders.com folder and then click Update Server Data File. You should perform this step to ensure that the newest records have been added to the file.
6. Switch back to the command prompt and type the following command:
dnscmd /OfflineSign /SignZone /Input northwindtraders.com.dns /Output signed.northwindtraders.com.dns /Zone northwindtraders.com /SignKey /Cert /FriendlyName KSK-northwindtraders.com /SignKey /Cert /FriendlyName ZSK-northwindtraders.com.

You now need to recreate the zone using the new zone file.

7. At the command prompt, type **dnscmd /ZoneDelete northwindtraders.com /f**.
8. At the command prompt, type **dnscmd /ZoneAdd northwindtraders.com /Primary /File signed.northwindtraders.com.dns /Load**.
9. In the DNS console tree, right-click the northwindtraders.com folder, and then click Refresh. In the details pane, you can see new RRSIG, DNSKEY, and NSEC records. The RRSIG records are the signatures of hashes of all other records in the zone: A, NS, SOA, DNSKEY, and NSEC.

EXERCISE 3 Configuring DNS Clients to Request DNS

DNS clients do not receive DNSSEC information unless they request it. You now need to create and apply a Group Policy object that configures DNS clients to request DNSSEC information from the DNS server.

1. While you are logged on to Nwtraders.msft from Dcsv1 as a domain administrator, open the Group Policy Management console by clicking Start, Administrative Tools, and Group Policy Management.
2. In the Group Policy Management console tree, open the Domains container, right-click the Nwtraders.msft icon, and then click Create A GPO In This Domain, And Link It Here.
3. In the New GPO dialog box, type **DNSSEC Policy** in the Name box, and then click OK.
4. In the Group Policy Management console tree, right-click DNSSEC Policy, and then click Edit. The Group Policy Management Editor opens.
5. In the Group Policy Management Editor console tree, navigate to Computer Configuration \Policies\Windows Settings\Name Resolution Policy.
6. In the details pane, in the Create Rules area, ensure that Suffix is selected beneath To Which Part Of The Namespace Does This Rule Apply? Then, in the text box to the right of Suffix, type **northwindtraders.com**.

To properly configure DNSSEC, you would normally type in this text box the suffix of a remote domain. We are specifying the local suffix here only as a simplified example for our limited test environment.

7. On the DNSSEC tab, click Enable DNSSEC In This Rule.
8. In the DNSSEC Settings area, beneath Validation, click the option to require DNS clients to check that name and address data has been validated by the DNS server.
9. Click Create, and then click Apply. A new entry appears in the Name Resolution Policy Table lower in the pane.
10. Close the Group Policy Management Editor.
11. Switch to the Boston computer. At a command prompt on Boston, type **gpupdate**.
12. On a command prompt on Boston, type **ping dcsv1.northwindtraders.com**.

The name is successfully resolved and authenticated, and Dcsv1 responds to the ping. You can use a protocol analyzer such as Network Monitor to verify that Boston has received RRSIG and NSEC records along with the A record that resolves the name Dcsv1.northwindtraders.com. The frame containing an RRSIG from such a capture is shown in Figure 3-56.

```

Frame Details
  Dns: QueryId = 0xCA81, QUERY (Standard query), Response - Success, 49, 0
    QueryIdentifier: 51841 (0xCA81)
    Flags: Response, Opcode - QUERY (Standard query), AA, RD, RA, Rcode - Success
    QuestionCount: 1 (0x1)
    AnswerCount: 2 (0x2)
    NameServerCount: 0 (0x0)
    AdditionalCount: 1 (0x1)
    QRecord: dcsrv1.northwindtraders.com of type Host Addr on class Internet
    ARecord: dcsrv1.northwindtraders.com of type Host Addr on class Internet: 192.168.100.1
    ARecord: dcsrv1.northwindtraders.com of type RRSIG on class Internet
      ResourceName: dcsrv1.northwindtraders.com
      ResourceType: RRSIG, 46(0x2e)
      ResourceClass: Internet, 1(0x1)
      TimeToLive: 3600 (0xE10)
      ResourceDataLength: 168 (0xA8)
      DNSSECRRSIG: Resource Type: RRSIG
        RRSIGRData:
          TypeCovered: 1 (0x1)
          Algorithm: RSA/SHA-1 [RSASHA1]
          Labels: 3 (0x3)
          OriginalTTL: 3600 (0xE10)
          SignatureExpiration: 01/18/2011, 22:58:59 .0000 UTC
          SignatureInception: 12/19/2010, 22:58:59 .0000 UTC
          KeyTag: 43941 (0xABA5)
          SignerName: northwindtraders.com
          Signature: Binary Large Object (128 Bytes)
    AdditionalRecord: Type: OPT, Sender's largest UDP payload size: 4000

```

FIGURE 3-56 Boston receives a resource record signature (RRSIG) along with the Host record.

Lesson Summary

- DNSSEC enables a DNS server to digitally sign the resource records in its zones. DNSSEC also enables DNS servers to validate the digital signatures of resource records received from other servers.
- Trust anchors are public keys from other zones that are used to validate digitally signed records originating from those zones and from delegated subdomains that are also DNSSEC-compatible. A DNS server configured with DNSSEC must have configured at least one trust anchor from a remote zone.
- You use Group Policy to configure DNS clients to request DNSSEC validation. In Group Policy, you specify queries in the Name Resolution Policy Table for which clients request DNSSEC validation.
- You use the Dnscmd utility to create the keys needed for DNSSEC and to sign the zone with these keys.
- It is recommended that you use two key pairs with your zone: the first, a frequently updated zone signing key (ZSK) to sign the zone; and a second, a rarely updated key signing key (KSK) that is stored on other servers and that validates the ZSK.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 3, "Implementing DNSSEC." The questions are also available on the companion CD if you prefer to review them in electronic form.

NOTE ANSWERS

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the "Answers" section at the end of the book.

- 1.** You work for an organization named Fabrikam.com whose network includes 10 servers running Windows Server 2008, 60 clients running Windows Vista Professional, and 40 clients running Windows 7 Professional. You need to implement the latest version of DNSSEC so that your clients can receive validated DNS responses. Which of the following steps do you need to take? (Choose all that apply.)
 - A.** Upgrade the DNS server to Windows Server 2008 R2.
 - B.** Configure a trust anchor for the public root domain.
 - C.** Upgrade all clients to Windows 7.
 - D.** Configure a trust anchor for every top-level domain.

- 2.** You are a network administrator for Contoso.com, whose network includes 25 servers running Windows Server 2008 R2 and 300 clients running Windows 7 Professional. Contoso.com is in the process of implementing DNSSEC so that its local DNS server can validate responses received from a remote DNS server in a partner organization named Northwindtraders.com.

Both the Contoso.com domain and the Northwindtraders.com domain have made a single zone-signing key (ZSK) and a single key-signing key (KSK) available in their respective DNS zones. You want to validate responses from Northwindtraders.com and minimize administrative effort in the future. What should you do?

 - A.** Inform the administrator at Northwindtraders.com to configure the ZSK from Contoso.com as a trust anchor.
 - B.** Inform the administrator at Northwindtraders.com to configure the KSK from Contoso.com as a trust anchor.
 - C.** Import the ZSK from Northwindtraders.com and configure it as a trust anchor.
 - D.** Import the KSK from Northwindtraders.com and configure it as a trust anchor.

3. You are an administrator for northwindtraders.com. You want to configure the clients in the northwindtraders.com domain to request DNSSEC validation whenever they query for a name in the fabrikam.com domain. What should you do?
 - A. Use Name Resolution Policy to create and enforce a rule that enables DNSSEC for the FQDN "fabrikam.com."
 - B. Use Name Resolution Policy to create and enforce a rule that enables DNSSEC for the suffix "fabrikam.com."
 - C. Import the trust anchor for the fabrikam.com domain to all your DNS servers.
 - D. Export the trust anchor for the northwindtraders.com domain to the fabrikam.com DNS servers, and configure conditional forwarding to the fabrikam.com domain.

Chapter Review

To further practice and reinforce the skills you learned in this chapter, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenarios. The scenarios set up real-world situations involving the topics of this chapter and ask you to create solutions.
- Complete the suggested practices.
- Take a practice test.

Chapter Summary

- A zone is a database that contains authoritative information about a portion of the DNS namespace. Zones are created on DNS servers. Primary zones provide the original read-write source data for a zone. Secondary zones are read-only copies of a zone. Stub zones contain only the names of servers containing primary or secondary zones.
- When you create a zone on a domain controller, you have the option to store the zone in Active Directory. This option offers a number of benefits, including reduced administration, improved security for dynamic updates, and multiple primary servers. If you do not store a zone in Active Directory, the zone is known as a standard zone, and the zone file is a text file. In standard zones, there is only one copy of the primary zone.
- Aging and scavenging provide a mechanism for removing stale resource records in a zone.
- The GlobalNames zone enables the resolution of single-label names in a multidomain forest.

- An application directory partition is a type of data structure used by DNS to store data for Active Directory–integrated zones. By default, every domain controller includes application directory partitions called DomainDnsZones and ForestDnsZones. These partitions are replicated among all domain controllers in the domain and the forest, respectively. You can also create custom application directory partitions and enlist chosen servers in the partition. You can choose to store a zone in any of these partitions. This decision affects what is called “the replication scope of the zone.”
- Zone transfers keep DNS data consistent between secondary zones and a master zone, which is usually a primary zone.
- DNSSEC is a means for a DNS server to generate digital signatures for records in its zone and to validate the digital signatures of records from remote zones. With DNSSEC, you use the Dnscmd utility both to create the keys needed to sign the zone and to sign the zone. The latest version of DNSSEC requires Windows Server 2008 R2 and Windows 7.

Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- aging
- application directory partition
- key pair
- key rollover
- key signing key (KSK)
- master zone
- Name Resolution Policy Table (NRPT)
- primary zone
- private key
- public key
- replication
- scavenging
- secondary zone
- stub zone
- trust anchor
- zone
- zone signing key (ZSK)
- zone transfers

Case Scenarios

In the following case scenarios, you will apply what you've learned about configuring DNS zones and zone transfers. You can find answers to these questions in the "Answers" section at the end of this book.

Case Scenario 1: Managing Outdated Zone Data

You work as a domain administrator for Fabrikam, Inc. Your responsibilities include managing the Active Directory and network infrastructure, including DNS. The DNS servers for the Fabrikam.com domain are all installed on domain controllers.

1. Recently you have noticed that some records in the Fabrikam.com zone refer to computers that were removed from the network several months ago. What is the best way to remove these stale records?
2. What is the best way to prevent such data from accumulating in the future?
3. You want to allow records to remain in the zone for 21 days without being scavenged. However, you want to prevent time stamps from being refreshed for the first seven days after each record is first created in the zone. How should you configure the No-Refresh and the Refresh intervals?

Case Scenario 2: Configuring Zone Transfers

You are a network administrator for City Power & Light, whose network is composed of a single Active Directory domain, Cpandl.com. The Cpandl.com zone is stored in Active Directory.

At the company headquarters, the Cpandl.com domain controllers host the DNS zones for the domain. The Cpandl.com network also includes several branch offices.

1. The Rochester office does not include a DNS server. You want to improve name resolution of computer names in the Cpandl.com domain, but you don't want to host a domain controller at the Rochester site. Minimizing zone transfer traffic is not a priority. What should you do?
2. You want zone transfers to the Rochester office to occur whenever a change occurs in the zone data. How can you enable this functionality?

Suggested Practices

To help you successfully master the exam objectives presented in this chapter, complete the following tasks.

Configure a DNS Infrastructure

The following practices will deepen your understanding of DNS replication within multi-domain forests. They both require three computers, but you can still perform these practices easily by using virtual machine software such as Hyper-V or VirtualBox.

- **Practice 1** Using virtual machines, create an Active Directory forest with two domain controllers in a domain named Contoso.com, and one domain controller in a child domain called East.contoso.com. Choose the option to store both DNS zones in all DNS servers in the forest. View the zone data and then add a record manually to each zone. Force replication by using Active Directory Sites and Services.
- **Practice 2** Using the same three-computer network, create a custom application directory partition on the domain controller in the East.contoso.com domain. Configure the zone to store its data in the newly created partition. Enlist only one of the domain controllers in the Contoso.com domain in the partition. Reboot each computer and then verify that the zone data is stored on only two of the three servers.

Take a Practice Test

The practice tests on this book's companion CD offer many options. For example, you can test yourself on just one exam objective, or you can test yourself on all the 70-642 certification exam content. You can set up the test so that it closely simulates the experience of taking a certification exam, or you can set it up in study mode so that you can look at the correct answers and explanations after you answer each question.

MORE INFO PRACTICE TESTS

For details about all the practice test options available, see the "How to Use the Practice Tests" section in this book's Introduction.

Configuring Windows Firewall and Network Access Protection

By their nature, networks can allow healthy computers to communicate with unhealthy computers and malicious tools to attack legitimate applications. This can result in costly security compromises, such as a worm that spreads rapidly through an internal network or a sophisticated attacker who steals confidential data across the network.

Windows Server 2008 R2 supports two technologies that are useful for improving network security: Windows Firewall and Network Access Protection (NAP). Windows Firewall can filter incoming and outgoing traffic, using complex criteria to distinguish between legitimate and potentially malicious communications. NAP requires computers to complete a health check before allowing unrestricted access to your network and facilitates resolving problems with computers that do not meet health requirements.

This lesson describes how to plan and implement Windows Firewall and NAP using Windows Server 2008 R2.

Exam objectives in this chapter:

- Configure Windows Firewall with Advanced Security.
- Configure Network Access Protection (NAP).

Lessons in this chapter:

- Lesson 1: Configuring Windows Firewall **430**
- Lesson 2: Configuring Network Access Protection **444**

Before You Begin

To complete the lessons in this chapter, you should be familiar with Windows networking and be comfortable with the following tasks:

- Adding roles to a computer running Windows Server 2008 R2
- Configuring Active Directory domain controllers and joining computers to a domain
- Configuring a basic network, including configuring IP settings

You will also need the following nonproduction hardware connected to test networks:

- A computer named Dcsvr1 that is a domain controller in the Nwtraders.msft domain. This computer must have at least one network interface that you can connect to either the Internet or a private network.

NOTE COMPUTER AND DOMAIN NAMES

The computer and domain names you use will not affect these exercises. The practices in this chapter refer to these computer names for simplicity, however.

- A computer named Hartford that is running Windows 7 Professional, Enterprise, or Ultimate, and is a member of the Nwtraders.msft domain. You must use Windows 7 because Windows Server 2008 R2 does not support the Windows Security Health Validator.



REAL WORLD

Tony Northrup

Instead of absolutes, security can be measured only in degrees of risk. Although NAP can't prevent a determined, skilled attacker from connecting to your network, NAP can improve your network security by helping keep computers up to date and ensuring that legitimate users do not accidentally connect to your internal network without meeting your security requirements.

When evaluating NAP as a way to protect against malicious attackers, remember that NAP trusts the System Health Agent (SHA) to report on the health of the client. The SHA is also running on the client computer. So it's a bit like airport security merely asking people if they are carrying any banned substances—people without any malicious intent would happily volunteer anything they accidentally brought. People with malicious intent would simply lie.

It's not *quite* as easy as simply lying, because the SHA signs the Statement of Health (SoH) to help prove that the health report is genuine. Additional security measures, such as requiring IPsec connection security, can help further reduce the opportunity for attackers. Nonetheless, with some time and effort, it's entirely possible that someone will create a malicious SHA that impersonates a legitimate SHA.

Lesson 1: Configuring Windows Firewall

Windows Firewall filters incoming traffic to help block unwanted network traffic. Optionally, Windows Firewall can also filter outgoing traffic to help limit the risk of malware. Although Windows Firewall's default settings will work well with components built into Windows, they might prevent other applications from functioning correctly. Windows Firewall's default settings can also be significantly improved to provide even stronger protection by requiring authorization or limiting the scope of allowed connections.

After this lesson, you will be able to:

- Describe the purpose of firewalls.
- List the three firewall profiles and how each is used.
- Create a firewall rule to allow inbound traffic.
- Create a firewall rule to allow outbound traffic and enable outbound filtering.
- Configure the scope of a firewall rule to limit communications to specific subnets.
- Configure firewall rules to require IPsec connection security and, optionally, limit authorization to specific users and computers.
- Use Group Policy settings to configure firewall rules in an Active Directory domain environment.
- Enable Windows Firewall logging so that you can isolate problems related to firewall rules.
- Identify network communications used by a specific application so that you can create rules for the application.

Estimated lesson time: 45 minutes

Why Firewalls Are Important

In networking, *firewalls* analyze communications and drop packets that haven't been specifically allowed. This is an important task, because connecting to the Internet means any of the millions of other Internet-connected computers can attack you. A successful compromise can crash a service or computer, compromise confidential data, or even allow the attacker to take complete control of the remote computer. In the case of *worms*, automated software attacks computers across the Internet, gains elevated privileges, copies itself to the compromised computer, and then begins attacking other computers (typically at random).

The purpose of a firewall is to drop unwanted traffic, such as traffic from worms, while allowing legitimate traffic, such as authorized file sharing. The more precisely you use firewall rules to identify legitimate traffic, the less you risk exposure to unwanted traffic from worms.

Firewall Profiles

When you create firewall rules to allow or block traffic, you can separately apply them to the Domain, Private, and Public profiles. These profiles enable mobile computers to allow incoming connections while connected to a domain network (for example, to allow incoming Remote Desktop connections) but block connection attempts on less secure networks (such as public wireless hotspots).

The firewall profiles are:

- **Domain** Applies when a computer is connected to its Active Directory domain. Specifically, any time a member computer's domain controller is accessible, this profile will be applied.

- **Private** Applies when a computer is connected to a private network location. By default, no networks are considered private—users must specifically mark a network location, such as their home office network, as private.
- **Public** The default profile applied to all networks when a domain controller is not available. For example, the Public profile is applied when users connect to Wi-Fi hotspots at airports or coffee shops. By default, the Public profile allows outgoing connections but blocks all incoming traffic that is not part of an existing connection.

Most servers are always connected to a domain environment. To ensure consistent operation even when a domain controller is not available, configure the same firewall rules for all three profiles when configuring a server.

Filtering Inbound Traffic

By default, Windows Firewall (as well as most other firewalls) blocks any inbound traffic that hasn't been specifically allowed. By default, the Public profile allows absolutely no incoming connections—this provides excellent security when connecting to public hotspots or other untrusted networks. The Domain and Private profiles allow some incoming connections, such as connections for file and printer sharing.

If you install or enable a Windows feature that requires incoming connections, Windows will automatically enable the required firewall rules. Therefore, you do not need to manually adjust the firewall rules. Figure 8-1 shows the default inbound firewall rules for a Windows Server 2008 R2 computer configured as a domain controller. As you can see, rules exist to allow each of the protocols required for a domain controller.

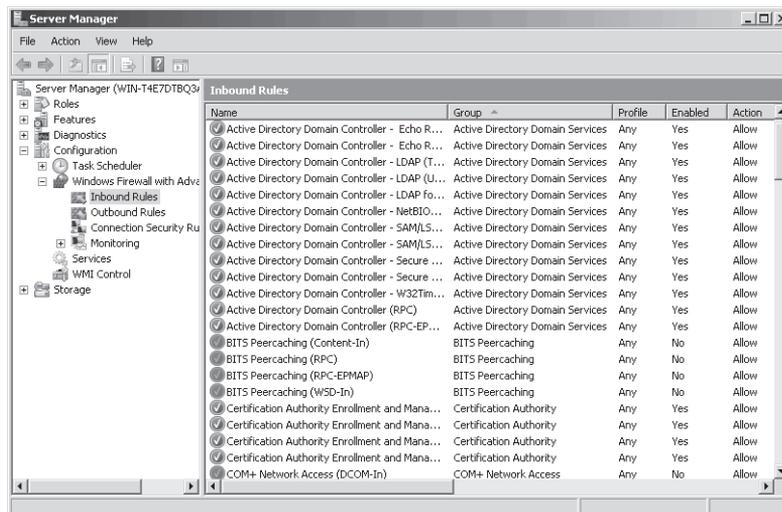


FIGURE 8-1 Default inbound firewall rules

If you install an application that does not automatically enable the required firewall rules, you will need to create the rules manually. You can create firewall rules by using the stand-alone Windows Firewall With Advanced Security console, or you can apply the rules with Group Policy by using the same interface at Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall With Advanced Security\Windows Firewall With Advanced Security.

To create an inbound filter, follow these steps:

1. In the Windows Firewall With Advanced Security snap-in, right-click Inbound Rules, and then choose New Rule. The New Inbound Rule Wizard appears.
2. On the Rule Type page, select one of the following options, and then click Next:
 - **Program** A rule that allows or blocks connections for a specific executable file, regardless of the port numbers it might use. You should use the Program rule type whenever possible. The only time it's not possible to use the Program rule type is when a service does not have its own executable.
 - **Port** A rule that allows or blocks communications for a specific TCP or UDP port number, regardless of the program generating the traffic.
 - **Predefined** A rule that controls connections for a Windows component, such as Active Directory Domain Services, File And Printer Sharing, or Remote Desktop. Typically, Windows enables these rules automatically.
 - **Custom** A rule that can combine program and port information.
3. Complete the page or pages that appear after you select one of the rule types. The page or pages you see will vary depending on the rule type you selected. Click Next.
4. On the Action page, select one of the following options, and then click Next.
 - **Allow The Connection** Allows any connection that matches the criteria you specified on the previous pages.
 - **Allow The Connection If It Is Secure** Allows connections that match the criteria you specified on the previous pages only if those connections are protected with IPsec. Optionally, you can select the Require The Connections To Be Encrypted check box, which requires encryption in addition to authentication. Selecting the Override Block Rules check box configures the rule to take precedence over other rules that might prevent a client from connecting. If you select this rule type, the wizard will also prompt you to select users and computers that are authorized to establish this type of connection.
 - **Block The Connection** Drops any connection attempt that matches the criteria you specified on the previous pages. Because inbound connections are blocked by default, you rarely need to create this rule type. However, you might use this action for an outbound rule if you specifically want to prevent an application from initiating outgoing connections.

5. On the Profile page, choose which profiles to apply the rule to. For most servers, you should apply the rule to all three profiles, because servers are usually continually connected to a single network. For mobile computers in domain environments, you typically need to apply firewall rules only to the Domain profile. If you do not have an Active Directory domain or if users need to use the firewall rule when connected to their home networks, apply the rule to the Private profile. Avoid creating firewall rules on mobile computers for the Public profile, because an attacker on an unprotected network might be able to exploit a vulnerability exposed by the firewall rule. Click Next.
6. On the Name page, type a name for the rule, and then click Finish.

The inbound rule takes effect immediately, allowing incoming connections that match the criteria you specified.

Filtering Outbound Traffic

By default, Windows Firewall allows all outbound traffic. Allowing outbound traffic is much less risky than allowing inbound traffic. However, outbound traffic still carries some risk:

- If malware infects a computer, it might send outbound traffic containing confidential data (such as content from a Microsoft SQL Server database, email messages from a Microsoft Exchange server, or a list of passwords).
- Worms and viruses seek to replicate themselves. If they successfully infect a computer, they will attempt to send outbound traffic to infect other computers. After one computer on an intranet is infected, network attacks can allow malware to rapidly infect computers on an intranet.
- Users might use unapproved applications to send data to Internet resources and either knowingly or unknowingly transmit confidential data.

By default, all versions of Windows (including Windows Server 2008 R2) do not filter outbound traffic. However, Windows Server 2008 R2 does include outbound filters for core networking services, enabling you to quickly enable outbound filtering while retaining basic network functionality. By default, outbound rules are enabled for:

- Dynamic Host Configuration Protocol (DHCP) requests
- DNS requests
- Group Policy communications
- Internet Group Management Protocol (IGMP)
- IPv6 and related protocols

Blocking outbound communications by default will prevent many built-in Windows features, and all third-party applications you might install, from communicating on the network. For example, Windows Update will no longer be able to retrieve updates, Windows will no longer be able to activate across the Internet, and the computer will be unable to send Simple Network Management Protocol (SNMP) alerts to a management host.

If you do enable outbound filtering, you must be prepared to test every application to verify that it runs correctly. Most applications are not designed to support outbound filtering and will require you to both identify the firewall rules that need to be created and then create those rules.

To create an outbound filter, follow these steps:

1. In Windows Firewall With Advanced Security (which you can access in Server Manager under Configuration), right-click Outbound Rules, and then choose New Rule. The New Outbound Rule Wizard appears.
2. On the Rule Type page, select a rule type (as described in the section “Filtering Inbound Traffic” earlier in this lesson), and then click Next.
3. On the Program page, click This Program Path. In the text box, type the path to the application’s executable file. Click Next.
4. On the Action page, select an action type (as described in the section “Filtering Inbound Traffic” earlier in this lesson), and then click Next.
5. On the Profile page, select the check boxes for the profiles that you want to apply the rule to, and then click Next.
6. On the Name page, type a name for the rule, and then click Finish.

The outbound rule takes effect immediately, allowing outgoing packets that match the criteria you specified.

To block outbound connections by default, first create and enable any outbound firewall rules so that applications do not immediately stop functioning. Then, follow these steps:

1. In Server Manager, right-click Configuration\Windows Firewall With Advanced Security, and then choose Properties.
2. Click the Domain Profile, Private Profile, or Public Profile tab.
3. From the Outbound Connections drop-down list, select Block. If necessary, return to the previous step to block outbound traffic for other profiles. Then click OK.

You will need to perform extensive testing to verify that all required applications function correctly when outbound connections are blocked by default. This testing should include background processes, such as Automatic Updates.

Configuring Scope

One of the most powerful ways to increase computer security is to configure firewall scope. Using *scope*, you can allow connections from your internal network and block connections from external networks. Scope can be used in the following ways:

- For a server that is connected to the Internet, you can allow anyone on the Internet to connect to public services (such as the web server) while allowing only users on your internal network to access private servers (such as Remote Desktop).
- For internal servers, you can allow connections only from the specific subnets that contain potential users. When planning such scope limitations, remember to include remote access subnets.

- For outgoing connections, you can allow an application to connect to servers only on specific internal subnets. For example, you might allow SNMP traps to be sent to only your SNMP management servers. Similarly, you might allow a network backup application to connect to only your backup servers.
- For mobile computers, you can allow specific communications (such as Remote Desktop) from only the subnets you use for management.

To configure the scope of a rule, follow these steps:

1. In the Windows Firewall With Advanced Security snap-in, select Inbound Rules or Outbound Rules.
2. In the details pane, right-click the rule you want to configure, and then choose Properties.
3. Click the Scope tab. In the Remote IP Address group, select These IP Addresses.
4. In the Remote IP Address group, click Add.

NOTE CONFIGURING SCOPE FOR LOCAL IP ADDRESSES

The only time you would want to configure the scope using the Local IP Address group is when the computer is configured with multiple IP addresses, and you do not want to accept connections on all IP addresses.

5. In the IP Address dialog box, select one of the following three options, and then click OK:
 - **This IP Address Or Subnet** Type an IP address (such as 192.168.1.22) or a subnet using Classless Inter Domain Routing (CIDR) notation (such as 192.168.1.0/24) that should be allowed to use the firewall rule.
 - **This IP Address Range** Using the From and To boxes, type the first and last IP address that should be allowed to use the firewall rule.
 - **Predefined Set Of Computers** Select a host from the list: Default Gateway, WINS Servers, DHCP Servers, DNS Servers, and Local Subnet.
6. Repeat steps 4 and 5 for any additional IP addresses that should be allowed to use the firewall rule, and then click OK.

Authorizing Connections

If you are using IPsec connection security in an Active Directory environment, you can also require the remote computer or user to be authorized before a connection can be established.

For example, imagine that your organization had a custom accounting application that used TCP port 1073, but the application had no access control mechanism—any user who connected to the network service could access confidential accounting data. Using Windows Firewall connection authorization, you could limit inbound connections to users who are members of the Accounting group—adding access control to the application without writing any additional code.

Most network applications do have access control built in, however. For example, you can configure Internet Information Server (a web server installed as part of the Application

Server role) to authenticate users and allow only authorized users to connect to a web application. Similarly, if you share a folder on the network, you can use file permissions and share permissions to restrict who can access the folder. Application-layer authorization should always be your first layer of security; however, connection authorization using Windows Firewall can provide an additional layer of security. Using multiple layers of security—a technique known as *defense-in-depth*—reduces risk by providing protection even when one layer has a vulnerability.

To configure connection authorization for a firewall rule, follow these steps:

1. In Server Manager, select Configuration\Windows Firewall With Advanced Security\Inbound Rules or Configuration\Windows Firewall With Advanced Security\Outbound Rules.
2. In the details pane, right-click the rule you want to configure, and then choose Properties.
3. Click the General tab. Select Allow Only Secure Connections. Because the authorization relies on IPsec, you can configure authorization only on secure connections.
4. Click the Users And Computers tab for an inbound rule or the Computers tab for an outbound rule. Select the proper options based on the rule you selected:
 - **To allow connections only from specific computers** Select the Only Allow Connections From These Computers check box for an inbound rule or the Only Allow Connections To These Computers check box for an outbound rule.
 - **To allow connections only from specific users** If you are editing an inbound rule, select the Only Allow Connections From These Users check box. You can use this option only for inbound connections.
5. Click Add and select the groups containing the users or computers you want to authorize. Figure 8-2 shows how the Users And Computers tab appears after you have configured connections for an inbound rule. Click OK.

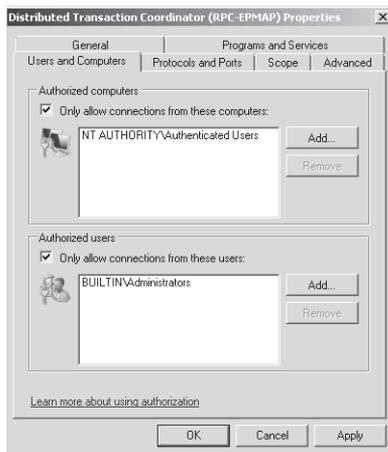


FIGURE 8-2 The Users And Computers tab

6. Click OK again.

Any future connections that match the firewall rule will require IPsec for the connection to be established. Additionally, if the authenticated computer or user is not on the list of authorized computers and users that you specified, the connection will be immediately dropped.

Configuring Firewall Settings with Group Policy

You can configure Windows Firewall locally, by using Server Manager or the Windows Firewall With Advanced Security console in the Administrative Tools folder; or globally, by using the Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall With Advanced Security\Windows Firewall With Advanced Security node of a Group Policy Object (GPO). Typically, you edit server-specific policies (such as configuring the range of IP addresses a DNS server accepts queries from) by using local tools, and you configure policies that apply to groups of computers (including IPsec connection security policies) by using GPOs.

You can use Group Policy to manage Windows Firewall settings for computers running Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2 by using two nodes:

- **Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall With Advanced Security\Windows Firewall With Advanced Security** This node applies settings only to computers running Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2 and provides exactly the same interface as the same node in Server Manager. You should always use this node when configuring computers running these recent versions of Windows because it provides for more detailed configuration of firewall rules.
- **Computer Configuration\Policies\Administrative Templates\Network\Network Connections \Windows Firewall** This node applies settings to computers running Windows XP, Windows Server 2003, Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2. This tool is less flexible than the Windows Firewall With Advanced Security console; however, settings apply to all versions of Windows that support Windows Firewall. If you are not using the new IPsec features in recent versions of Windows, you can use this node to configure all your clients.

For best results, create one GPO for Windows 7, Windows Vista, Windows Server 2008 R2, and Windows Server 2008, and create a second GPO for Windows Server 2003 and Windows XP. Then, use WMI filters to target the GPOs to computers running only the appropriate version of Windows.

MORE INFO CREATING WMI FILTERS

For more information about creating WMI filters, read Microsoft Knowledge Base article 555253, "HOWTO: Leverage Group Policies with WMI Filters," at <http://support.microsoft.com/kb/555253>.

Enabling Logging for Windows Firewall

If you are ever unsure about whether Windows Firewall is blocking or allowing traffic, you should enable logging, re-create the problem you're having, and then examine the log files. To enable logging, follow these steps:

1. In the console tree of the Windows Firewall With Advanced Security snap-in, right-click Windows Firewall With Advanced Security, and then choose Properties. The Windows Firewall With Advanced Security Properties dialog box appears.
2. Select the Domain Profile, Private Profile, or Public Profile tab.
3. In the Logging group, click the Customize button. The Customize Logging Settings dialog box appears.
4. To log packets that Windows Firewall drops, from the Log Dropped Packets drop-down list, select Yes. To log connections that Windows Firewall allows, from the Log Successful Connections drop-down list, select Yes.
5. Click OK.

By default, Windows Firewall writes log entries to %SystemRoot%\System32\LogFiles\Firewall\%Pfirewall.log and stores only the last 4 MB of data. In most production environments, this log will be almost constantly written to, which can cause a performance impact. For that reason, you should enable logging only when actively troubleshooting a problem and then immediately disable logging when you're finished.

Identifying Network Communications

The documentation included with network applications often does not clearly identify the communication protocols the application uses. Fortunately, creating Program firewall rules allows any communications required by that particular program.

If you prefer to use Port firewall rules, or if you need to configure a network firewall that can identify communications based only on port number and the application's documentation does not list the firewall requirements, you can examine the application's behavior to determine the port numbers in use.

The simplest tool to use is Netstat. On the server, run the application, and then run the following command to examine which ports are listening for active connections:

```
netstat -a -b
```

Any rows in the output with a State of LISTENING are attempting to receive incoming connections on the port number specified in the Local Address column. The executable name listed after the row is the executable that is listening for the connection. For example, the following output demonstrates that RpcSs, running under the SvcHost.exe process (which runs many services), is listening for connections on TCP port 135.

Active Connections

Proto	Local Address	Foreign Address	State
TCP	0.0.0.0:135	Dcsrv1:0	LISTENING
RpcSs [svchost.exe]			

Similarly, the following output demonstrates that the DNS service (Dns.exe) is listening for connections on TCP port 53:

Active Connections

Proto	Local Address	Foreign Address	State
TCP	0.0.0.0:53	Dcsrv1:0	LISTENING
[dns.exe]			

Although Windows Firewall has existing rules in place for these services (because they are built into Windows), the same technique would allow you to identify the port numbers used by any third-party application.

PRACTICE Configuring Windows Firewall

In this practice, you configure both inbound and outbound filtering. These are common tasks that occur when you install new applications in almost any network environment, from small businesses to large enterprises.

EXERCISE 1 Configuring Inbound Filtering

In this exercise, you will install the Telnet Server feature, which configures Windows Server 2008 R2 to accept incoming connections on TCP port 23. Then, you will examine the incoming firewall rule that applies to the Telnet Server and adjust the rule configuration.

1. In the console tree of Server Manager, select Features. In the details pane, click Add Features. The Add Features Wizard appears.
2. On the Select Features page, select the Telnet Server check box. Click Next.
3. On the Confirm Installation Selections page, click Install.
4. On the Installation Results page, click Close.
5. In Server Manager, select Configuration\Services. Then, in the details pane, right-click the Telnet Server service and choose Properties. From the Startup Type drop-down list, select Manual. Click the Apply button. Then, click the Start button to start the telnet server. Click OK.
6. On a client computer, open a command prompt and run the following command (where *ip_address* is the telnet server's IP address):

```
telnet ip_address
```

The telnet server should prompt you for a user name. This proves that the client was able to establish a TCP connection to port 23.

7. Press Ctrl+] to exit the telnet session. Type **quit** and press Enter to close telnet.
8. On the telnet server, in Server Manager, select Configuration\Windows Firewall With Advanced Security\Inbound Rules. In the details pane, right-click the Telnet Server rule, and then choose Properties.

NOTE AUTOMATICALLY ENABLING REQUIRED RULES

Notice that the Telnet Server rule is enabled; the Add Features Wizard automatically enabled the rule when it installed the Telnet Server feature.

9. Click the Programs And Services tab. Notice that the default rule is configured to allow communications for %SystemRoot%\system32\TlntSvr.exe, which is the executable file for the telnet server service. Click the Settings button and verify that Telnet is selected. Click Cancel twice.
10. In Server Manager, right-click the Telnet Server rule, and then choose Disable Rule.
11. On the telnet client computer, run the same Telnet command again. This time the command should fail because Windows Firewall is no longer allowing incoming telnet requests.
12. Use Server Manager to remove the Telnet Server feature and restart the computer if necessary.

EXERCISE 2 Configuring Outbound Filtering

In this exercise, you configure Windows Server 2008 R2 to block outbound requests by default. Then, you test it by attempting to visit a website with Internet Explorer. Next, you create an outbound rule to allow requests from Internet Explorer and verify that the outbound rule works correctly. Finally, you return your computer to its original state.

1. Open Internet Explorer and visit *http://www.microsoft.com*. If an Internet Explorer Enhanced Security Configuration dialog box appears, you can click Close to dismiss it.
2. In Server Manager, right-click Configuration\Windows Firewall With Advanced Security, and then choose Properties.
3. Click the Domain Profile tab. From the Outbound Connections drop-down list, select Block. Repeat this step for the Private Profile and Public Profile tabs. Then click OK.
4. Open Internet Explorer and attempt to visit *http://support.microsoft.com*. You should be unable to visit the website because outbound filtering is blocking Internet Explorer's outgoing HTTP queries.
5. In Server Manager, within Configuration\Windows Firewall With Advanced Security, right-click Outbound Rules, and then choose New Rule. The New Outbound Rule Wizard appears.
6. On the Rule Type page, select Program. Then, click Next.
7. On the Program page, select This Program Path. In the box, type **%ProgramFiles%\Internet Explorer\iexplore.exe** (the path to the Internet Explorer executable file). Click Next.
8. On the Action page, select Allow The Connection. Then, click Next.

9. On the Profile page, accept the default selection of applying the rule to all three profiles. Click Next.
10. On the Name page, type **Allow Internet Explorer outgoing communications**. Then click Finish.
11. In Internet Explorer, attempt to visit *http://support.microsoft.com* again. This time the connection succeeds because you created an outbound filter specifically for Internet Explorer.
12. In Server Manager, disable outbound filtering by right-clicking Configuration\Windows Firewall With Advanced Security, and then choosing Properties. On the Domain Profile tab, click the Outbound Connections list, and then click Allow (Default). Repeat this step for the Private Profile and Public Profile tabs. Click OK.

Lesson Summary

- Firewalls are designed to drop unwanted communications (such as packets generated by a worm) while still allowing legitimate communications (such as packets generated by a network management tool).
- Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2 support three firewall profiles: Domain, Private, and Public. The Domain profile applies whenever a computer can communicate with its domain controller. The Private profile must be manually applied to a network. The Public profile applies any time a domain controller is not available, and a network has not been configured as Private.
- Use the Windows Firewall With Advanced Security snap-in to create an inbound firewall rule that allows a server application to receive incoming connections.
- Use the Windows Firewall With Advanced Security snap-in to create an outbound firewall rule that allows a client application to establish outgoing connections. You need to create outbound firewall rules only when you configure outbound connections to be blocked by default.
- You can edit the properties of a firewall rule to configure the scope, which limits the subnets an application can communicate with. Configuring scope can greatly reduce the risk of attacks from untrusted networks.
- If you use IPsec in your environment, you can configure firewall rules to allow only secure connections and to allow only connections for authorized users and computers.
- Group Policy is the most effective way to configure firewall settings for all computers in a domain. Using Group Policy, you can quickly improve the security of a large number of computers and control which applications are allowed to communicate on the network.
- Windows Firewall logging identifies connections that Windows Firewall allows or blocks. This information is very useful when troubleshooting a connectivity problem that might be caused by Windows Firewall.

- If an application must accept incoming connections but the developers have not documented the communication ports that the application uses, you can use the Netstat tool to identify which ports the application listens on. With this information, you can then create Port firewall rules.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 1, “Configuring Windows Firewall.” The questions are also available on the companion CD if you prefer to review them in electronic form.

NOTE ANSWERS

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the “Answers” section at the end of the book.

1. You need to install an internally developed automation tool on a computer running Windows Server 2008 R2. The tool acts as a network client and needs to connect to a server on your intranet using TCP port 88 and to a server on the Internet using TCP port 290. Additionally, a client component you install on your workstation running Windows 7 will connect to the computer running Windows Server 2008 R2 using TCP port 39. Windows Firewall is currently configured with the default settings on both computers. Which of the following changes do you need to make to allow the application to work?
 - A. On the computer running Windows Server 2008 R2, add a firewall rule to allow outbound connections on TCP port 290.
 - B. On the computer running Windows Server 2008 R2, add a firewall rule to allow inbound connections on TCP port 39.
 - C. On the computer running Windows Server 2008 R2, add a firewall rule to allow inbound connections on TCP port 290.
 - D. On your workstation, add a firewall rule to allow outbound connections on TCP port 39.
2. You have recently installed an internal server application on a computer running Windows Server 2008 R2 that accepts incoming connections on TCP port 1036. The application does not include any access control capability. How can you configure the inbound firewall rule properties to allow connections only from authorized users in your domain? (Choose all that apply. Each answer forms part of the complete solution.)
 - A. On the General tab, click Allow Only Secure Connections.
 - B. On the Advanced tab, click These Profiles, and then select Domain.
 - C. On the Users And Computers tab, select Only Allow Connections From These Users. Then, add the Domain Users group.
 - D. On the Scope tab, in the Local IP Address group, select These IP Addresses. Then, add each of your internal networks.

3. You need to use Group Policy settings to configure firewall settings on your client computers running Windows XP and Windows 7. You would like to configure firewall rules using only the Windows Firewall node rather than the Windows Firewall With Advanced Security node. Which of the following features are *not* available when using the Windows Firewall node in Group Policy settings?
- A. Filtering UDP traffic
 - B. Allowing a specific executable to accept incoming connections on any port number
 - C. Dropping connections not originating from a specific subnet
 - D. Requiring IPsec authentication for a connection

Lesson 2: Configuring Network Access Protection

Consider this common scenario: an enterprise has thousands of computers on a private network. Perimeter firewalls protect the network from Internet threats, including network attacks from worms. Suddenly, someone creates a worm that can exploit a vulnerability in computers running Windows that do not have the latest security updates installed. The worm spreads quickly across the Internet, but the private network's perimeter firewalls protect the vulnerable computers on the internal network. A traveling salesperson then returns to the office with his mobile computer. While on his trip, he connected his computer to the wireless network at the hotel, where another guest's computer transmitted a worm across the network. When he connects to the private network, the worm immediately begins spreading to the vulnerable computers, completely bypassing the perimeter security. In a few hours, most of the computers on the internal network are infected.

This exact scenario has happened to many organizations and resulted in countless financial losses. NAP can prevent this scenario. When computers connect to your local area network (LAN), they must meet specific health requirements, such as having recent updates installed. If they can't meet those health requirements, they can be quarantined to a network where they can download updates, install antivirus software, and obtain more information about how to meet the requirements of the LAN. This lesson describes NAP and how you can deploy it on your network.

After this lesson, you will be able to:

- Describe how NAP works to protect your network.
- Plan a NAP deployment while minimizing the impact on users.
- Install and configure the Network Policy Service.
- Configure NAP enforcement.
- Configure various NAP components.
- Examine NAP log files.

Estimated lesson time: 90 minutes

Network Access Protection Concepts

As shown in Figure 8-3, NAP is designed to connect hosts to different network resources depending on their current health state. This division of network resources can be implemented using virtual LANs (VLANs, as Figure 8-3 illustrates), IP filters, IP subnet assignment, static routes, or IPsec enforcement.

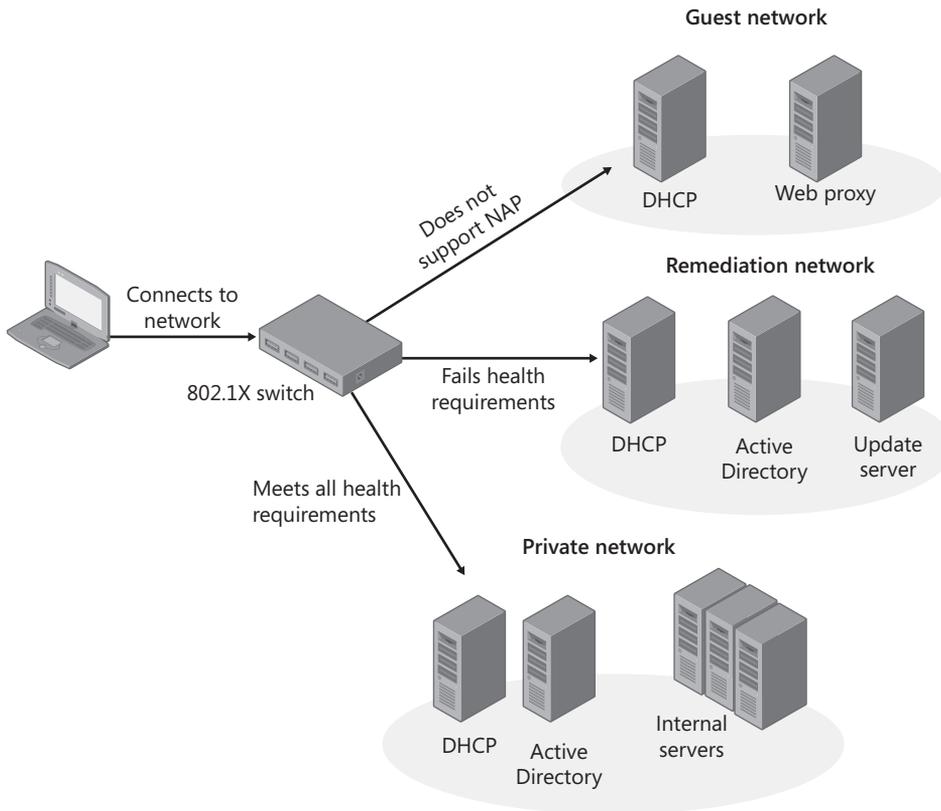


FIGURE 8-3 A typical NAP VLAN architecture

If you choose to provide a remediation network (rather than simply denying network access), you might need additional infrastructure servers for it. For example, if you configure an Active Directory domain controller on the remediation network, you should use a read-only domain controller to limit the risk if the domain controller is attacked. Similarly, you should provide separate DHCP and DNS servers from your infrastructure servers to reduce the risk that a noncompliant computer might spread malware to the production server.

Enforcement Types

For NAP to work, a network component must enforce NAP by either allowing or denying network access. The sections that follow describe the different NAP enforcement types you can use: IPsec connection security, 802.1X access points, VPN servers, DHCP servers, and Remote Desktop Gateways (RD Gateway).

NOTE TERMINAL SERVICES GATEWAY

Terminal Services Gateway enforcement is not discussed in this book because it is not covered on the exam.

IPSEC CONNECTION SECURITY

The IPsec connection security enforcement type requires clients to perform a NAP health check before they can receive a health certificate. In turn, this health certificate is required for IPsec connection security before the client can connect to IPsec-protected hosts. IPsec enforcement allows you to require health compliance on a per-IP address or a per-TCP/UDP port number basis. For example, you could allow noncompliant computers to connect to a web server but allow only compliant computers to connect to a file server—even if the two services are running on a single computer.

You can also use IPsec connection security to allow healthy computers to communicate only with other healthy computers. IPsec enforcement requires a CA running Windows Server 2008 (or Windows Server 2008 R2) Certificate Services and NAP to support health certificates. In production environments, you will need at least two CAs for redundancy. Other public key infrastructures (PKIs) will not work. IPsec enforcement provides a very high level of security, but it can protect only computers that are configured to support IPsec.

MORE INFO DEPLOYING A PKI

For more information about deploying a new Windows-based PKI in your organization, see <http://www.microsoft.com/pki> and *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008).

802.1X ACCESS POINTS

The 802.1X access points enforcement type uses Ethernet switches or wireless access points that support 802.1X authentication. Compliant computers are granted full network access, and non-compliant computers are connected to a remediation network or completely prevented from connecting to the network. If a computer falls out of compliance after connecting to the 802.1X network, the 802.1X network access device can change the computer's network access. This provides some assurance of compliance for desktop computers, which might remain connected to the network indefinitely.

802.1X enforcement uses one of two methods to control which level of access compliant, noncompliant, and unauthenticated computers receive:

- **An access control list (ACL)** A set of Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) packet filters configured on the 802.1X access point. The 802.1X access point applies the ACL to the connection and drops all packets that are not allowed by the ACL. Typically, you apply an ACL to noncompliant computer connections and allow compliant computers to connect without an ACL (thus granting them unlimited network access). ACLs allow you to prevent noncompliant computers from connecting to one another, thus limiting the ability of a worm to spread, even among noncompliant computers.
- **A virtual local area network (VLAN)** A group of ports on the switch that are grouped together to create a separate network. VLANs cannot communicate with one another unless you connect them using a router. VLANs are identified using a VLAN identifier, which must be configured on the switch itself. You can then use NAP to specify in which VLAN the compliant, noncompliant, and unauthenticated computers are placed. When you place noncompliant computers into a VLAN, they can communicate with one another. This can allow a noncompliant computer infected with a worm to attack, and possibly infect, other noncompliant computers. Another disadvantage of using VLANs is that the client's network configuration must change when transitioning from being a non-compliant NAP client to being a compliant NAP client (for example, if they are able to successfully apply updates). Changing the network configuration during system startup and user logon can cause Group Policy updates or other boot processes to fail.

Your 802.1X access points may support ACLs, VLANs, or both. If they support both and you're already using either ACLs or VLANs for other purposes, use the same technique for 802.1X enforcement. If your 802.1X access point supports both ACLs and VLANs and you are not currently using either, use ACLs for 802.1X enforcement so that you can take advantage of their ability to limit network access between noncompliant clients.

VPN SERVERS

The VPN servers type enforces NAP for remote access connections using a VPN server running Windows Server 2008 or Windows Server 2008 R2 and Routing and Remote Access (other VPN servers do not support NAP). With VPN server enforcement enabled, only compliant client computers are granted unlimited network access. The VPN server can apply a set of packet filters to connections for noncompliant computers, limiting their access to a remediation server group that you define. You can also define IPv4 and IPv6 packet filters, exactly as you would when configuring a standard VPN connection.

MORE INFO CONFIGURING VPN CONNECTIONS

For more information about configuring VPN connections, refer to Chapter 7, "Connecting to Networks."

DHCP SERVERS

The DHCP servers enforcement type uses a computer running Windows Server 2008 or Windows Server 2008 R2 and the Dynamic Host Configuration Protocol (DHCP) Server service that provides IP addresses to intranet clients. Only compliant computers receive an IP address that grants full network access; noncompliant computers are granted an IP address with a subnet mask of 255.255.255.255 and no default gateway.

Additionally, noncompliant hosts receive a list of *host routes* (routes that direct traffic to a single IP address) for network resources in a remediation server group that you can use to allow the client to apply any updates required to become compliant. This IP configuration prevents noncompliant computers from communicating with network resources other than those you configure as part of a remediation server group.

If the health state of a NAP client changes (for example, if Windows Firewall is disabled), the NAP client performs a new health evaluation using a DHCP renewal. This allows clients that become noncompliant after successfully authenticating to the network to be blocked from further network access. If you change the health policy on NAP servers, the changes will not be enforced until the client's DHCP lease is renewed.

Although 802.1X network access devices and VPN servers are capable of disconnecting computers from the network and IPsec enforcement can allow connections only from healthy computers, DHCP server enforcement points can be bypassed by an attacker who manually configures an IP address. Nonetheless, DHCP server enforcement can reduce the risk from nonmalicious users who might attempt to connect to your network with a non-compliant computer.

REMOTE DESKTOP GATEWAYS

If you use RD Gateway (called Terminal Services Gateway in Windows Server 2008) to allow users to control their desktops from remote computers across the Internet, you can use the RD Gateway enforcement type to block access using RD Gateway unless the client computer passes a health check. RD Gateway enforcement does not provide remediation.

System Health Agents and System Health Validators

NAP health validation takes place between two components:

- **System Health Agents (SHAs)** The client components that create a Statement of Health (SoH) containing a description of the health of the client computer. Windows 7, Windows Vista, Windows Server 2008, Windows Server 2008 R2, and Windows XP with Service Pack 3 include an SHA that monitors Windows Security Center settings. Microsoft and third-party developers can create custom SHAs that provide more complex reporting.
- **System Health Validators (SHVs)** The server components that analyze the SoH generated by the SHA and create an SoH Response (SoHR). The NAP health policy server uses the SoHR to determine the level of access the client computer should have and whether any remediation is necessary. Windows Server 2008 and Windows Server 2008 R2 include an SHV that corresponds to the SHA built into Windows 7, Windows Vista, and Windows XP with Service Pack 3.

The NAP connection process is as follows:

1. The NAP client connects to a network that requires NAP.
2. Each SHA on the NAP client validates its system health and generates an SoH. The NAP client combines the SoHs from multiple SHAs into a System Statement of Health (SSoH), which includes version information for the NAP client and the set of SoHs for the installed SHAs.
3. The NAP client sends the SSoH to the NAP health policy server through the NAP enforcement point.
4. The NAP health policy server uses its installed SHVs and the health requirement policies that you have configured to determine whether the NAP client meets health requirements. Each SHV produces a Statement of Health Response (SoHR), which can contain remediation instructions (such as the version number of an antivirus signature file) if the client doesn't meet that SHV's health requirements.
5. The NAP health policy server combines the SoHRs from the multiple SHVs into a System Statement of Health Response (SSoHR).
6. The NAP health policy server sends the SSoHR back to the NAP client through the NAP enforcement point. The NAP enforcement point can now connect a compliant computer to the network or connect a noncompliant computer to a remediation network.
7. Each SHA on the NAP client processes the SoHR created by the corresponding SHV. If possible, any noncompliant SHAs can attempt to come into compliance (for example, by downloading updated antivirus signatures).
8. If any noncompliant SHAs were able to meet the requirements specified by the SHV, the entire process starts over again—hopefully with a successful result.



Quick Check

1. Which NAP enforcement types do not require support from your network infrastructure?
2. Which versions of Windows can act as NAP clients?

Quick Check Answers

1. IPsec connection security, DHCP, VPN, and Remote Desktop Gateway enforcement do not require support from your network infrastructure. They can be implemented using only Windows Server 2008 R2. The 802.1X type provides very powerful enforcement, but requires a network infrastructure that supports 802.1X.
2. Windows XP with Service Pack 3, Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2.

Planning a NAP Deployment

NAP has the potential to prevent legitimate users from accessing the network. Any security mechanism that reduces productivity will be quickly removed, so you must carefully plan a NAP deployment to minimize user impact.

Typically, a NAP deployment occurs in three phases:

- **Testing** Test the NAP using examples of each different operating system, client computer configuration, and enforcement points in your environment.
- **Monitoring** Deploy NAP in a monitoring-only mode that notifies administrators if a computer fails to meet health requirements but does not prevent the user from connecting to the network. This allows you to identify computers that are not meeting health requirements and to bring them into compliance. You could bring computers into compliance manually or by using automated tools, such as Microsoft System Center Configuration Manager 2007. For more information, read the section entitled “Configuring NAP for Monitoring Only” later in this chapter.
- **Limited access** If, during the monitoring phase, you reach a point where almost all of your computers are compliant, you can enable NAP enforcement to prevent non-compliant computers from connecting to your production network. Users can then use resources on the remediation network to bring their computers into compliance, if necessary. Typically, you will need to configure exceptions for computers that are not NAP-compliant.

Installing and Configuring the Network Policy Server

NAP depends on a Windows Server 2008 or Windows Server 2008 R2 NAP health policy server, which acts as a RADIUS server, to evaluate the health of client computers. If you have existing RADIUS servers that are running Windows Server 2003 or Windows 2000 Server and Internet Authentication Service (IAS), you can upgrade them to Windows Server 2008 or Windows Server 2008 R2 and configure them as NAP health policy servers. If you have RADIUS servers running any other operating system, you will need to configure new Windows Server 2008 or Windows Server 2008 R2 NAP health policy servers, configure the health policy, and then migrate your existing RADIUS clients to the NAP health policy servers.

Typically, you will need to deploy at least two NAP health policy servers for fault tolerance. If you have only a single NAP health policy server, clients will be unable to connect to the network if it is offline. As described in Chapter 7, you can use connection request policies to allow a single RADIUS server to act as a NAP health policy server and authenticate requests from other RADIUS clients.

Installing NAP

To install NAP, follow these steps:

1. In the console tree of Server Manager, select Roles. In the details pane, click Add Roles. The Add Roles Wizard appears.
2. If the Before You Begin page appears, click Next.

3. On the Select Server Roles page, select the Network Policy And Access Services check box. Click Next.
4. On the Network Policy And Access Services page, click Next.
5. On the Select Role Services page, select the Network Policy Server check box. Click Next.
6. On the Confirmation page, click Install.
7. On the Results page, click Close.

This installs the core NPS service, which is sufficient for using the Windows Server 2008 computer as a RADIUS server for 802.1X, VPN, or DHCP enforcement.

Using the Configure NAP Wizard

After installing the Network Policy And Access Services role, follow these steps to configure NAP:

1. In Server Manager, select Roles\Network Policy And Access Services\NPS. You might need to close and reopen Server Manager if you recently installed the Network Policy And Access Services role.
2. In the details pane, select Network Access Protection, and then click Configure NAP. The Configure NAP Wizard appears.
3. On the Select Network Connection Method For Use With NAP page, choose your enforcement method. Then, click Next.
4. On the next page (whose title depends on the previously selected network connection method), you need to add any HRA servers (other than the local computer) and RADIUS clients, for example:
 - If you are using 802.1X enforcement, add the IP address of each switch.
 - If you are using VPN enforcement, add the IP address of each VPN server.
 - If you are configuring DHCP servers, add each of your NAP-capable DHCP servers.Click Add for each host and configure a friendly name, address, and shared secret. Then click OK. After you have configured any external HRA servers and RADIUS clients, click Next.
5. Depending on the network method you chose, you might be presented with additional page options, such as DHCP scopes or RD gateway redirection options. Configure these options appropriately.
6. On the Configure User Groups And Machines page, click the Add buttons to allow computers or groups to connect. Click Next.
7. The pages that follow vary depending on your NAP enforcement method:
 - For the 802.1X or VPN enforcement methods, you use the Configure An Authentication Method page (shown in Figure 8-4) to specify the NAP health policy server certificate and the EAP types to use for user or computer-level authentication.
 - For the 802.1X enforcement method, you use the Configure Traffic Controls page to configure the unlimited VLAN and the restricted network VLAN.

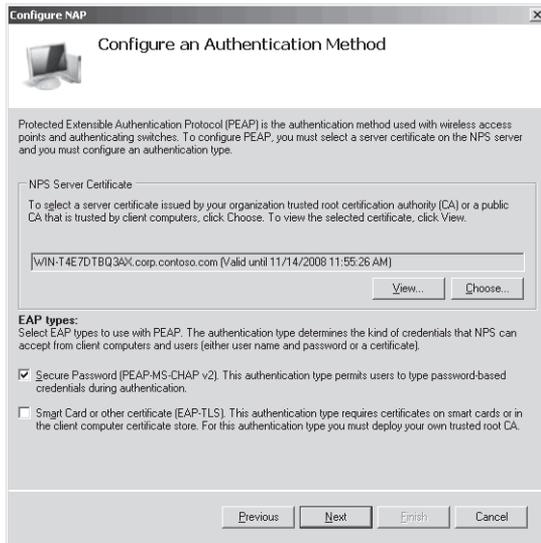


FIGURE 8-4 Configuring an 802.1X enforcement authentication method

8. On the Define NAP Health Policy page, you can select from the installed SHVs. By default, only the Windows Security Health Validator is installed. As shown in Figure 8-5, you should leave autoremediation enabled for enforcement types that support it to allow client computers to automatically change settings to meet health requirements. During initial production deployments, select Allow Full Network Access To NAP-Ineligible Client Computers to configure NAP in monitoring-only mode. Noncompliant computers will generate an event in the event log, allowing you to fix noncompliant computers before they are prevented from connecting to the network. Click Next.

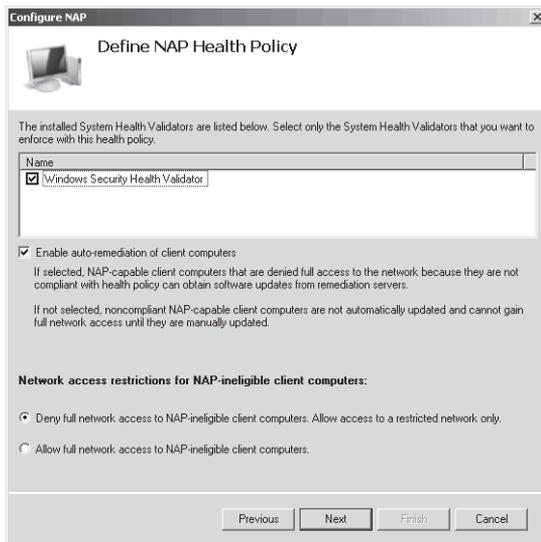


FIGURE 8-5 Defining NAP health policy

9. On the Completing NAP Enforcement Policy And RADIUS Client Configuration page, click Finish.

The Configure NAP Wizard creates the following policies:

- A connection request policy with the name specified on the Select Network Connection Method For Use With NAP page
- Compliant and noncompliant health policies, based on the name specified on the Select Network Connection Method For Use With NAP page
- Compliant and noncompliant network policies, based on the same name as the health policies

Configuring NAP Enforcement

After you have installed and configured NAP, you must perform additional steps to enable NAP enforcement. The steps you follow vary depending on whether you are using IPsec, 802.1X, DHCP, VPN, or RD Gateway enforcement. The sections that follow describe how to configure each of these enforcement types at a high level, cross-referencing other sections in this lesson that have more detailed instructions.

Configuring IPsec Enforcement

Configuring IPsec enforcement requires the following high-level steps:

1. Install the Health Registration Authority (HRA) role service and the Active Directory Certificate Services role (if it's not already present).
2. Use the Configure NAP Wizard to configure the connection request policy, network policy, and NAP health policy, as described earlier in this chapter in the section titled "Using the Configure NAP Wizard." Although you can configure these elements individually, using the wizard is much easier.
3. Configure HRA, as described in the sections that follow.
4. Enable the NAP IPsec Relying Party enforcement client and start the NAP service on NAP-capable client computers, as described later in this chapter in the sections entitled "Configuring Client Computers for IPsec Enforcement" and "Configuring NAP Clients."
5. Require IPsec connection security using health certificates for computers that should communicate only with other healthy computers, as described in the sections that follow.

The following sections describe these steps in more detail.

STEP 1: INSTALLING THE HRA ROLE SERVICE

If you plan to use IPsec enforcement, you will also need to install the HRA role service. In production environments, you should always configure at least two HRAs for fault tolerance. Large networks might require additional HRAs to meet the performance requirements.

Installing the HRA role service configures the following:

- **A certification authority (if one does not already exist)** HRA requires a certification authority running Windows Server 2008 or Windows Server 2008 R2 Certificate Services, which can be an existing CA or a new CA. For a Windows Server 2003–based CA, you must manually create a System Health Authentication certificate template so that members of the IPsec exemption group can autoenroll a long-lived health certificate.

MORE INFO CONFIGURING A CA FOR IPSEC NAP ENFORCEMENT

For more information about configuring a Windows Server 2003–based CA, read “Step By Step Guide: Demonstrate IPsec NAP Enforcement in a Test Lab” at http://download.microsoft.com/download/d/2/2/d22daf01-a6d4-486c-8239-04db487e6413/NAIPsec_StepByStep.doc.

- **A web application** The Add Role Services Wizard creates a web application named DomainHRA under the default website in Internet Information Services (IIS).

You can install the HRA role service using the Add Roles Wizard by selecting the Health Registration Authority check box on the Select Role Services page and following the prompts that appear, or you can install the role service after installing the Network Policy And Access Services role by following these steps:

1. In Server Manager, right-click Roles\Network Policy and Access Services, and then choose Add Role Services. The Add Role Services Wizard appears.
2. On the Select Role Services page, select the Health Registration Authority check box. When prompted, click Add Required Role Services. Click Next.
3. On the Choose The Certification Authority To Use With The Health Registration Authority page, choose to install a CA, use the local CA, specify a remote CA, or defer the decision until later. Then, click Next.
4. On the Choose Authentication Requirements For The Health Registration Authority page, select Yes if all client computers are members of a trusted domain. If some computers are not members of a domain, you can select No—but you must accept slightly weaker security. Click Next.
5. If the Server Authentication Certificate page appears, you can select an SSL certificate to encrypt communications with the HRA server using one of the following three options. After you select an option, click Next.
 - **Choose An Existing Certificate For SSL Encryption** If you have an SSL certificate, select this option, and then select the certificate you want to use. If your certificate does not appear in the list, click Import.
 - **Create A Self-Signed Certificate For SSL Encryption** Clients do not trust self-signed certificates by default, which means you will need to manually configure the certificate on every client computer. For this reason, it is not a practical option in most circumstances.

- **Don't Use SSL Or Choose A Certificate For SSL Encryption Later** If you are installing Certificate Services as part of this wizard, select this option so that you can manually add an SSL certificate after you have completed the Certificate Services installation.

NOTE INSTALLING AN SSL CERTIFICATE AFTER COMPLETING THE WIZARD

You can install an SSL certificate later using the Internet Information Services Manager. Right-click Sites\Default Web Site, and then choose Edit Bindings. In the Site Bindings dialog box, click Add and create an HTTPS binding with your SSL certificate.

6. If you are installing the Windows Server 2008 Certificate Services role at this time, the Active Directory Certificate Services page appears. If it does not appear, skip to step 15. On this page, click Next.
7. On the Role Services page, click Next.
8. On the Setup Type page, select whether to configure an enterprise or stand-alone CA. In Active Directory environments, configuring an Enterprise CA is much easier because you can automatically issue certificates to client computers. Click Next.
9. On the CA Type page, select Root CA if this is your first CA. If you have an existing PKI, select Subordinate CA. The remainder of these steps apply to configuring a root CA; some pages are different if you configure a subordinate CA. Click Next.
10. On the Private Key page, click Next.
11. On the Cryptography page, click Next.
12. On the CA Name page, you can type a new common name for the CA. This name must be the name clients will use to connect to the server. The default will typically work. Click Next.
13. On the Validity Period page, click Next.
14. On the Certificate Database page, click Next.
15. If you are installing IIS role services at this time, the Web Server page appears. If it does not appear, skip to step 17. Otherwise, Click Next.
16. On the Role Services page, click Next.
17. On the Confirmation page, click Install.
18. On the Results page, click Close.

STEP 2: CONFIGURING THE NAP WIZARD

Follow the steps in "Using The Configure NAP Wizard" and, on the Select Network Connection Method For Use With NAP page, select IPsec With Health Registration Authority. Completing the wizard creates the following:

- A connection request policy named NAP IPsec With HRA (at Roles\Network Policy And Access Server\NPS\Policies\Connection Request Policies in Server Manager). This connection request policy configures the local server to process NAP IPsec requests using the HRA.

- A health policy named NAP IPsec With HRA Compliant (at Roles\Network Policy And Access Server\NPS\Policies\Health Policies in Server Manager). This health policy applies to compliant computers that pass all SHV checks.
- A network policy named NAP IPsec With HRA Compliant (at Roles\Network Policy And Access Server\NPS\Policies\Network Policies in Server Manager). This network policy grants access to compliant computers.
- A health policy named NAP IPsec With HRA Noncompliant (at Roles\Network Policy And Access Server\NPS\Policies\Health Policies in Server Manager). This health policy applies to noncompliant computers that fail one or more SHV checks.
- A network policy named NAP IPsec With HRA Noncompliant (at Roles\Network Policy And Access Server\NPS\Policies\Network Policies in Server Manager). This network policy grants limited network access to noncompliant computers. Specifically, non-compliant computers will be able to access only remediation servers. You should never set the Access Permission to Deny Access, because doing so prevents the health check from being performed.

STEP 3: CONFIGURING HRA

Now you can configure HRA settings using Server Manager by selecting the Roles\Network Policy And Access Services\NPS\Health Registration Authority node. Before you can use IPsec enforcement, you must configure a CA (such as Windows Server 2008 R2 Certificate Services) that will issue health certificates. If you didn't configure the CA while installing HRA, you can install it afterward.

To configure the CA that will be used to issue health certificates for IPsec enforcements, follow these steps:

1. In Server Manager, right-click Roles\Network Policy And Access services\Health Registration Authority\Certification Authority, and then choose Add Certification Authority.
2. In the Add Certification Authority dialog box, click Browse to select an enterprise CA. Select the appropriate server, and then click OK. Alternatively, you can type the fully qualified domain name (FQDN) of your CA. Figure 8-6 shows the Add Certification Authority dialog box with an enterprise CA selected.

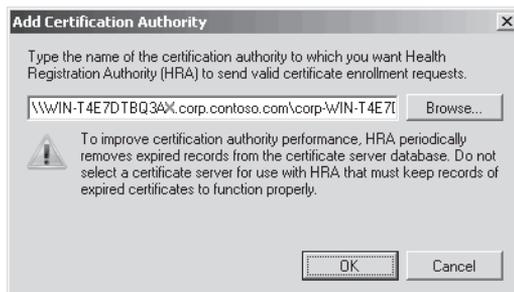


FIGURE 8-6 Selecting a CA for IPsec enforcement

3. Click OK.

4. Right-click Roles\Network Policy And Access Services\Health Registration Authority\Certification Authority, and then click Properties. The Certification Authorities Properties dialog box appears.
5. If you are using an enterprise CA, select Use Enterprise Certification Authority. Then click OK.

The CA appears in the details pane when you select the Roles\Network Policy And Access Services\Health Registration Authority\Certification Authority node in Server Manager. You can repeat the previous steps to add CAs, which allows for fault tolerance. If you have only a single CA and it goes offline, clients will be unable to undergo a NAP health check. If you have NAP enforcement enabled, this means clients will be unable to connect to the network.

You can also configure the mechanisms used for IPsec enforcement using the Roles\Network Policy And Access Services\Health Registration Authority\Certification Authority node in Server Manager. However, the default settings are typically sufficient.

STEP 4: CONFIGURING CLIENT COMPUTERS FOR IPSEC ENFORCEMENT

After configuring the NPS server for IPsec enforcement, you must configure client computers for IPsec enforcement. First, configure clients to use IPsec, as described in Chapter 6, “Protecting Network Traffic with IPsec.” Then, configure the client by following these steps:

1. Use the Group Policy Management Editor to open the GPO you want to use to apply the NAP enforcement client settings.
2. Right-click the Computer Configuration\Policies\Windows Settings\Security Settings\Network Access Protection\NAP Client Configuration\Health Registration Settings\Trusted Server Groups node, and then choose New. The New Trusted Server Group Wizard appears.
3. On the Group Name page, type a name that describes the group of HRA servers you will use for IPsec enforcement. Click Next.
4. On the Add Servers page, type the URL for each HRA. If you have an SSL certificate (that clients trust) installed on the server, type the URL as **https://<servername>**, where <servername> matches the common name on the SSL certificate. If you do not have an SSL certificate, clear the Require Server Verification check box and type the URL as **https://<servername>**. Click Add and repeat the process for any additional HRAs. NAP clients always start with the first HRA and continue through the list until an HRA can be contacted. Click Finish.

Now that you have configured clients to trust your HRAs, you should enable IPsec enforcement.

1. Select the Computer Configuration\Policies\Windows Settings\Security Settings\Network Access Protection\NAP Client Configuration\Enforcement Clients node.
2. In the details pane, double-click IPsec Relying Party.
3. In the IPsec Relying Party Properties dialog box, select the Enable This Enforcement Client check box. Then, click OK.

Additionally, follow the steps described in the section “Configuring NAP Clients” later in this chapter.

STEP 5: CONFIGURING IPSEC CONNECTION SECURITY RULES

Next, configure any servers that should be accessed only by compliant computers to require IPsec for inbound (but not outbound) connections. Note that this will prevent network communications from all computers that are not NAP-compliant or NAP-capable. In the Windows Firewall With Advanced Security snap-in (which you can access within the Configuration node of Server Manager), follow these steps:

1. Click Connection Security Rules. Then, right-click Connection Security Rules, and then choose New Rule. The New Connection Security Rule Wizard page appears.
2. On the Rule Type page, select Isolation. Then, click Next.
3. On the Requirements page, select Require Authentication For Inbound Connections And Request Authentication For Outbound Connections. Click Next.
4. On the Authentication Method page, select Advanced. Then, click Customize. In the First Authentication Group, click Add. In the Add First Authentication Method dialog box, shown in Figure 8-7, click Computer Certificate From This Certification Authority (CA). Click Browse and select the CA used to generate the certificate for your HRA. Click OK. Select the Accept Only Health Certificates and Enable Certificate To Account Mapping check boxes and then click OK. When you return to the wizard, click Next.

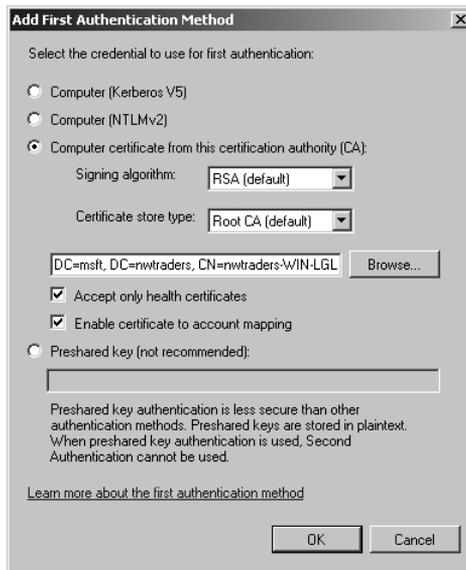


FIGURE 8-7 Requiring health certificates for a server

5. On the Profile page, click Next.
6. On the Name page, type a name, and then click Finish.

After the policy is applied to computers, only clients with a valid health certificate will be able to communicate. For this reason, you can't require health certificates for your HRA server, or clients would be unable to retrieve their health certificates.

For the HRA server, remediation servers, and any other computer that should be accessible by either noncompliant or non-NAP-capable computers, configure an IPsec connection security rule to request, but not require, security for inbound connections. For more information, read Chapter 6.

For NAP clients running Windows XP SP3, you will need to configure the equivalent policies using the IP Security Policies snap-in, available in Group Policy at Computer Configuration \Policies \Windows Settings\IP Security Policies. To configure a Windows XP SP3-based NAP client to use its health certificate for IPsec authentication, you must set the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PolicyAgent\Oakley \IKEFlags registry value to 0x1c.

Configuring 802.1X Enforcement

Configuring 802.1X enforcement requires the following high-level steps:

1. Use the Configure NAP Wizard to configure the connection request policy, network policy, and NAP health policy, as described in the section of this chapter entitled “Using the Configure NAP Wizard.” Although you can configure these elements individually, using the wizard is much easier. On the Virtual LAN (VLAN) Configuration page, you will need to specify the ACLs or VLANs for both compliant and noncompliant NAP clients, as shown in Figure 8-8. Refer to your switch documentation for information about which RADIUS attributes to use to specify the VLAN or ACL.

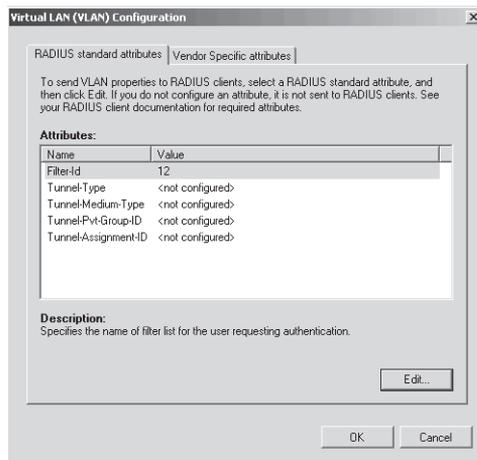


FIGURE 8-8 Configuring the VLAN for unrestricted network access

2. Configure your 802.1X authenticating switches to perform Protected Extensible Authentication Protocol (PEAP)-based authentication (either PEAP-MS-CHAP v2 or PEAP-TLS) and submit RADIUS requests to your NAP server. Additionally, configure a reauthentication interval to require authenticated client computers that remain connected to the network to be reauthenticated regularly. Microsoft suggests a reauthentication interval of four hours. Refer to your switch documentation for instructions.

3. If you plan to use certificates for authentication (using either PEAP-TLS or EAP-TLS), deploy a PKI such as the Certificate Services role and distribute certificates to client computers using a mechanism such as Active Directory autoenrollment. For more information, refer to Chapter 7. If you plan to use PEAP-MS-CHAP v2 domain authentication, use a PKI to issue server certificates to the NAP server.
4. Create NAP exemptions for computers that cannot complete a NAP health evaluation by creating a network policy that grants wireless or wired access and uses the Windows Groups condition set to the security group for the exempted computers but does not use the Health Policy condition. For more information, read “Configuring Network Policies” later in this lesson.
5. Enable the NAP EAP Quarantine Enforcement Client and start the NAP service on NAP-capable client computers. For more information, read “Configuring NAP Clients” later in this lesson.

Configuring DHCP Enforcement

Configuring DHCP enforcement requires the following high-level steps:

1. Use the Configure NAP Wizard to configure the connection request policy, network policy, and NAP health policy, as described in the section of this chapter entitled “Using the Configure NAP Wizard.” Although you can configure these elements individually, it’s much easier to use the wizard.
2. Configure remediation servers to define the computers that noncompliant clients can access. For more information, read “Configuring Remediation” later in this lesson.
3. Configure a DHCP server. For more information, refer to Chapter 4, “Creating a DHCP Infrastructure.” NPS must be installed on the DHCP server. If your DHCP and primary NPS servers are different computers, configure NPS on the remote DHCP NPS server as a RADIUS proxy to forward connection requests to the primary NPS server. For more information about configuring RADIUS proxies, refer to Chapter 7.
4. In the DHCP console, enable NAP for individual scopes or for all scopes on the DHCP server, as described in the sections that follow.
5. Enable the NAP DHCP Quarantine Enforcement Client and start the NAP service on NAP-capable client computers. For more information, read “Configuring NAP Clients” later in this chapter.

ENABLING NAP ON ALL DHCP SCOPES

To enable NAP for all DHCP scopes on a DHCP server, follow these steps:

1. In Server Manager, right-click Roles\DHCP Server*<Computer Name>*\IPv4, and then choose Properties.

2. In the Network Access Protection tab (as shown in Figure 8-9), click Enable On All Scopes. Confirm your choice, and then select one of the following options:
 - **Full Access** Enables NAP for monitoring only. Noncompliant clients will be granted full network access.
 - **Restricted Access** Enables NAP enforcement. Noncompliant clients will be assigned an IP address configuration that grants access only to servers listed in the remediation server group.
 - **Drop Client Packet** Ignores DHCP requests from noncompliant clients. Windows clients will then automatically assign themselves an Automatic Private IP Addressing (APIPA) address in the 169.254.0.0/16 network, where they will be able to communicate only with other APIPA computers.

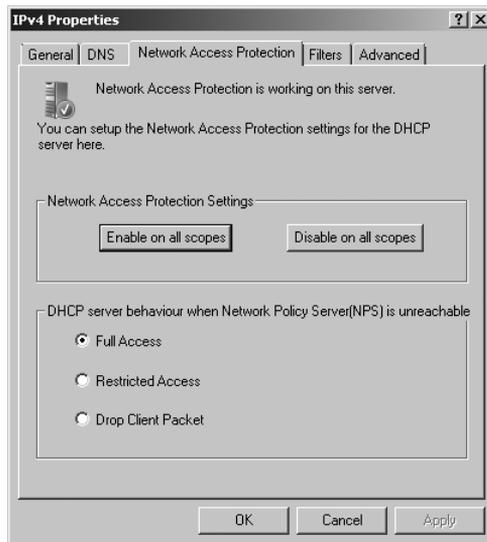


FIGURE 8-9 Configuring NAP on a DHCP server

3. Click OK.

ENABLING NAP ON A SINGLE DHCP SCOPE

To enable NAP for a single DHCP scope, follow these steps:

1. In Server Manager, right-click Roles\DHCP Server*<Computer Name>*\IPv4*<Scope Name>*, and then choose Properties.
2. In the Network Access Protection tab, select Enable For This Scope. Then, click OK.

Repeat these steps for each scope that you want to protect using NAP. For more information, read Chapter 4.

Configuring VPN Enforcement

Configuring VPN enforcement requires the following high-level steps:

1. Use the Configure NAP Wizard to configure the connection request policy, network policy, and NAP health policy, as described in the section of this chapter entitled “Using the Configure NAP Wizard.” Although you can configure these elements individually, it is much easier to use the wizard.
2. Configure remediation servers to define the computers that noncompliant clients can access. For more information, read “Configuring Remediation” later in this lesson.
3. Configure your VPN servers to perform PEAP-based authentication (either PEAP-MS-CHAP v2 or PEAP-TLS) and submit RADIUS requests to your NAP server. For more information, refer to Chapter 7.
4. If you plan to use certificates for authentication (using either PEAP-TLS or EAP-TLS), deploy a PKI such as the Certificate Services role and distribute certificates to client computers using a mechanism such as Active Directory autoenrollment. For more information, refer to Chapter 7. If you plan to use PEAP-MS-CHAP v2 domain authentication, use a PKI to issue server certificates to the NAP server.
5. Enable the NAP Remote Access Quarantine Enforcement Client and start the NAP service on NAP-capable client computers. For more information, read “Configuring NAP Clients” later in this chapter.

Configuring RD Gateway Enforcement

Configuring RD Gateway enforcement requires the following high-level steps:

1. Use the Configure NAP Wizard to configure the connection request policy, network policy, and NAP health policy, as described in the section of this chapter entitled “Using the Configure NAP Wizard.” Although you can configure these elements individually, it is much easier to use the wizard.
2. If you plan to use certificates for authentication (using either PEAP-TLS or EAP-TLS), deploy a PKI such as the Certificate Services role and distribute certificates to client computers using a mechanism such as Active Directory autoenrollment. For more information, refer to Chapter 7. If you plan to use PEAP-MS-CHAP v2 domain authentication, use a PKI to issue server certificates to the NAP server.
3. Enable NAP health policy checks on your RD Gateway server using the RD Gateway Manager snap-in. In Server Manager, right-click Roles\Remote Desktop Services\RD Gateway Manager*<computer_name>*, and then click Properties. On the RD CAP Store tab, verify that the Request Clients To Send A Statement Of Health check box is selected, which it is by default. If NPS is running on a different server, select the Central Server Running NPS check box, and then select your NPS server.
4. On NAP-capable client computers, enable the NAP RD Gateway Enforcement Client and the EAP Enforcement Client. Then, start the NAP service. For more information, read “Configuring NAP Clients” later in this chapter.

Configuring NAP Components

Depending on the NAP enforcement type and your organization's specific requirements, you will need to configure SHVs, NAP client settings, and health requirement policies. Additionally, during the initial deployment phase, you will need to configure NAP for monitoring only. The sections that follow describe these tasks in detail.

Configuring NAP Clients

After configuring the NPS server, you must configure client computers for NAP. The easiest way to do this is to use GPO settings in the Computer Configuration\Policies\Windows Settings\Security Settings\Network Access Protection\NAP Client Configuration node. You can configure client NAP settings using the three subnodes:

- **Enforcement Clients** You must enable one policy to configure clients to use that enforcement type.
- **User Interface Settings** Configure the User Interface Settings policy to provide customized text (and, optionally, an image) that users will see as part of the NAP client interface.
- **Health Registration Settings** Use the Request Policy subnode to configure cryptographic settings for NAP clients (the default settings are typically fine). Use the Trusted Server Group subnode to configure an HRA for IPsec NAP clients to use.

Additionally, you must start the Network Access Protection Agent service on all client computers. You can do this manually, but it is easiest to use Group Policy settings. In your GPO, select the Computer Configuration\Policies\Windows Settings\Security Settings\System Services node. Then, double-click the Network Access Protection Agent service. Define the policy in the properties dialog box, and set it to start automatically, as shown in Figure 8-10.

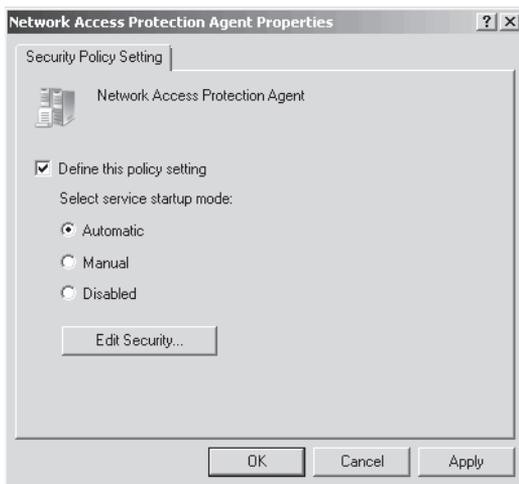


FIGURE 8-10 Starting the Network Access Protection Agent service automatically

Finally, to allow managed clients to use the default Windows SHV, you must enable Security Center by enabling the Computer Configuration\Policies\Administrative Templates\Windows Components\Security Center\Turn On Security Center policy.

NOTE CONFIGURING A WORKING NAP ENVIRONMENT

NAP configuration is complex, and this lesson has shown you many ways to configure NAP. Be sure to complete the practice at the end of this lesson to complete a NAP implementation from start to finish.

You can quickly verify a client's configuration by running the following command at a command prompt:

```
netsh nap client show state
```

The following output shows a client that has the Network Access Protection Agent service started and only the IPsec enforcement agent enabled:

Client state:

```
-----  
Name = Network Access Protection Client  
Description = Microsoft Network Access Protection Client  
Protocol version = 1.0  
Status = Enabled  
Restriction state = Not restricted  
Troubleshooting URL =  
Restriction start time =
```

Enforcement client state:

```
-----  
Id = 79617  
Name = DHCP Quarantine Enforcement Client  
Description = Provides DHCP based enforcement for NAP  
Version = 1.0  
Vendor name = Microsoft Corporation  
Registration date =  
Initialized = No  
  
Id = 79618  
Name = Remote Access Quarantine Enforcement Client  
Description = Provides the quarantine enforcement for RAS Client  
Version = 1.0  
Vendor name = Microsoft Corporation  
Registration date =  
Initialized = No  
  
Id = 79619  
Name = IPSec Relying Party  
Description = Provides IPSec based enforcement for Network Access Protection  
Version = 1.0  
Vendor name = Microsoft Corporation  
Registration date =  
Initialized = Yes
```

Id = 79621
Name = TS Gateway Quarantine Enforcement Client
Description = Provides TS Gateway enforcement for NAP
Version = 1.0
Vendor name = Microsoft Corporation
Registration date =
Initialized = No

Id = 79623
Name = EAP Quarantine Enforcement Client
Description = Provides EAP based enforcement for NAP
Version = 1.0
Vendor name = Microsoft Corporation
Registration date =
Initialized = No

System health agent (SHA) state:

Id = 79744
Name = Windows Security Health Agent

Description = The Windows Security Health Agent checks the compliance of a computer with an administrator-defined policy.

Version = 1.0

Vendor name = Microsoft Corporation

Registration date =
Initialized = Yes
Failure category = None
Remediation state = Success
Remediation percentage = 0
Fixup Message = (3237937214) - The Windows Security Health Agent has finished updating its security state.

Compliance results =
Remediation results =

Ok.

If applying Group Policy settings is not convenient, you can use the SHA ID numbers to enable a NAP client at the command line (or from within a script). For example, to enable the DHCP Quarantine enforcement client (which has an ID of 79617), run the following command:

```
netsh nap client set enforcement 79617 enable
```

Configuring a Health Requirement Policy

Health requirement policies determine which clients must meet health requirements, what those health requirements are, and what happens if a client cannot comply. A health requirement policy is a combination of the following:

- **Connection request policy** Determines whether a request should be processed by NPS.
- **System health validators** Define which health checks a client must meet to be considered compliant. For example, with the default Windows SHV, you can configure whether not having a firewall enabled makes a client noncompliant.
- **Remediation server group** A group of servers that noncompliant clients can access. These servers should provide clients with DNS and Active Directory services, as well as access to resources that will allow the client to become compliant, such as an update server.
- **Health policy** Defines health requirements using SHV settings. Separate health policies must exist for both compliant and noncompliant clients.
- **Network policy** Defines the level of network access clients get based on which health policy they match. You also use network policies to define the remediation servers that clients with limited access can connect to. As shown in Figure 8-11, you can specify network policy conditions that cause the network policy to apply to a client based on matching a specific health policy, operating system, or whether the client supports NAP.

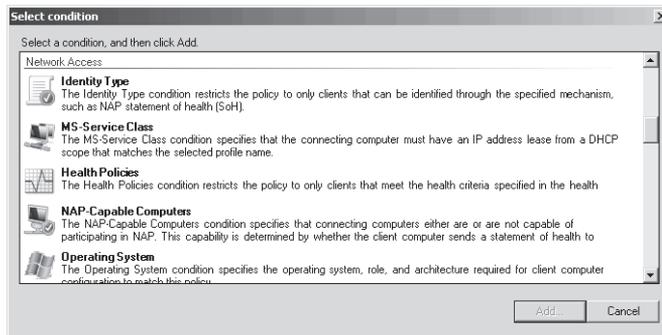


FIGURE 8-11 Configuring conditions for a network policy

CONFIGURING SHVS

Windows Server 2008 R2 includes only the Windows Security Health Validator SHV. Either Microsoft or third parties can supply additional SHVs that you would need to install on every NPS server.

After installing SHVs, configure the defaults (including the Windows SHV, described in the next section, “Configuring the Windows Security Health Validator”) by following these steps:

1. In Server Manager, select the Roles\Network Policy And Access Services\NPS\Network Access Protection\System Health Validators node.
2. In the details pane, right-click the SHV, and then choose Properties.

3. First, configure the error code resolution settings, as shown in Figure 8-12. In Server Manager, right-click Roles\Network Policy And Access Services\NPS\Network Access Protection\System Health Validators\<SHV_Name>\Error Codes, and then click Properties. For each of the five settings, you can define whether clients are compliant or noncompliant. Leave these set to Noncompliant for best security. However, if you experience a problem with clients receiving an error code when they should be compliant (for example, if an SHV or SHA needs to contact external services and cannot because of intermittent connectivity problems), you can change the error code resolution to Compliant. This could allow clients who would otherwise fail a health check to connect to your network, however.

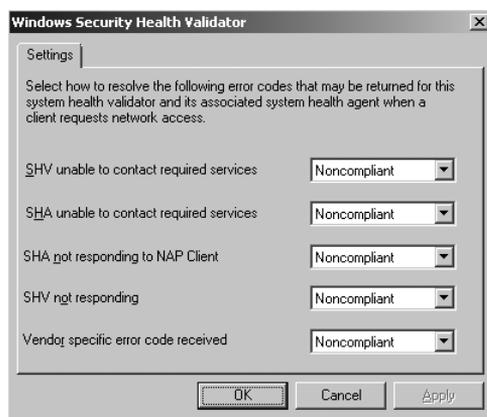


FIGURE 8-12 Configuring SHV error code resolution

4. Select the Roles\Network Policy And Access Services\NPS\Network Access Protection\System Health Validators\<SHV_Name>\Settings node in Server Manager to configure settings specific to that SHV, and then click OK. This dialog box is different for every SHV.

CONFIGURING THE WINDOWS SECURITY HEALTH VALIDATOR

By default, Windows Server 2008 R2 includes a single SHV: the Windows SHV. The Windows SHV performs many of the same checks as the Security Center:

- Verifies that a firewall (such as Windows Firewall) is enabled for all network connections. Windows XP, Windows Vista, and Windows 7 include Windows Firewall, which fulfills this requirement.
- Verifies that antivirus software is present and that the signatures are up to date. Because Windows does not include antivirus software, this check will cause Windows computers to fail by default.
- For Windows Vista and Windows 7 computers, verifies that antispyware software is present and the signatures are up to date. Windows Vista and Windows 7 include Windows Defender, which fulfills this requirement. You can also install Windows Defender on Windows XP computers, but the Windows Security Health Validator does not support checking antispyware software for computers running Windows XP.
- Automatic Updating is enabled.

Additionally, you can restrict access for clients that do not have all recent security updates installed and establish what level of security updates are required: Critical Only, Important And Above, Moderate And Above, Low And Above, or All. Figure 8-13 shows the Windows Security Health Validator properties with its default settings. The Windows XP node applies only to Windows XP clients with Service Pack 3 installed.

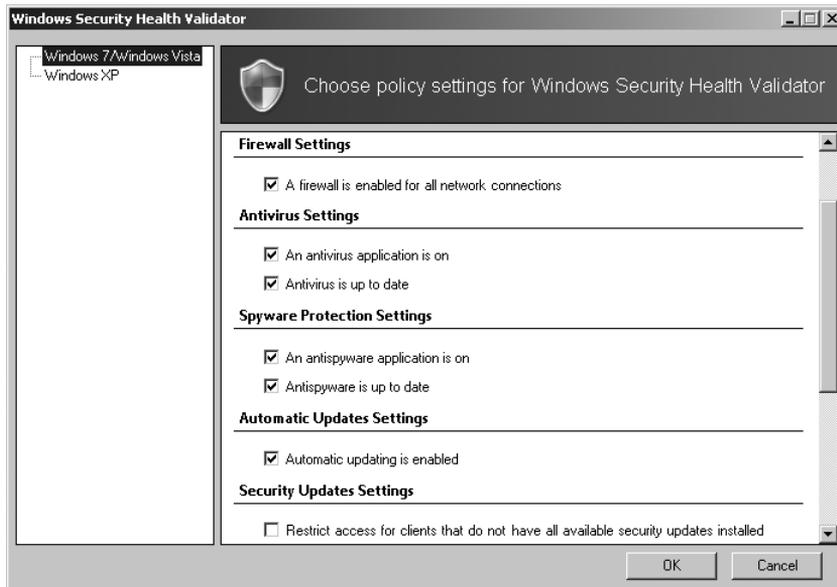


FIGURE 8-13 Editing the Windows SHV properties

To configure the Windows SHV, select NPS\Network Access Protection\System Health Validators\Windows Security Health Validator\Settings in the Network Policy And Access Services snap-in. Then, in the details pane, double-click Default Configuration. Alternatively, you can create additional configurations by clicking New in the Actions pane.

CONFIGURING REMEDIATION

Although NPS is designed to improve security by preventing noncompliant computers from connecting to your network, when it does detect a problem, it prevents legitimate users from their jobs. Therefore, you need resources so that those users can quickly and easily bring their computers into compliance and once again be productive.

To provide assistance to users of noncompliant computers when requiring NAP health enforcement, you can configure a remediation server group and troubleshooting URL that will be available to users if they fail the compliance check. The remediation server group is used only for DHCP and VPN enforcement types; 802.1X and IPsec enforcement use different technologies to limit network access. Remediation servers are not required if you are using reporting mode, because computers that fail the health check will still be allowed to connect to the network.

Although your exact remediation servers will vary depending on the requirements of your SHVs (the remediation servers should allow a noncompliant computer to enter compliance), remediation servers typically consist of the following:

- DHCP servers to provide IP configuration
- DNS servers, and optionally WINS servers, to provide name resolution
- Active Directory domain controllers, preferably configured as read-only, to minimize security risks
- Internet proxy servers so that noncompliant NAP clients can access the Internet
- HRAs so that noncompliant NAP clients can obtain a health certificate for the IPsec enforcement method
- A troubleshooting URL server, which provides a webpage users can access to view more information about the problem
- Antivirus update servers to retrieve updated antivirus signatures (if required by the health policy)
- Antispyware update servers to retrieve updated antispyware signatures (if required by the health policy)
- Software update servers

To configure these settings, follow these steps:

1. In Server Manager, select Roles\Network Policy And Access Services\NPS\Policies \Network Policies.
2. In the details pane, double-click the compliance policy that applies to noncompliant computers.
3. In the properties dialog box, click the Settings tab. In the Settings list, select NAP Enforcement. Then, click the Configure button.
4. In the Remediation Servers And Troubleshooting URL dialog box, do one or both of the following:
 - Use the Remediation Server Group list to select a remediation server group. If you haven't created a remediation server group, click the New Group button. Name the group, and then click the Add button to add each server that should be accessible to clients who fail the compliance check. One remediation server group might be enough, but you can create separate remediation server groups for noncompliant NAP clients and non-NAP-capable clients. Click OK.

NOTE UPDATING THE REMEDIATION SERVER GROUP

You can update your remediation server group later using Server Manager by selecting the Roles\Network Policy And Access Services\NPS\Network Access Protection\Remediation Server Groups node.

- In the Troubleshooting URL group, type the internal URL to a webpage that provides users with more information about why they can't connect to the network, how they can bring their computers into compliance, and whom they can call for assistance. A noncompliant computer visits this URL when a user clicks More Information in the Network Access Protection dialog box, which appears when a user attempts to troubleshoot a failed connection, as shown in Figure 8-14. On the webpage, you should provide information that the user can employ either to determine how to update the computer so that it is compliant or to troubleshoot network access. This URL is also visible when a user runs the **netsh nap client show state** command. The web server you specify in the URL should be part of the Remediation Server Group list so that the client computer can access it.



FIGURE 8-14 Information provided to a noncompliant NAP client

5. Click OK.

CONFIGURING NETWORK POLICIES

Network policies determine whether a connection request matches specific conditions (such as a health policy or a client operating system, or whether a computer is NAP-capable). They then grant full or limited network access to the client.

To add a network policy, follow these steps:

1. In Server Manager, right-click Roles\Network Policy And Access Services\NPS\Policies \Network Policies, and then choose New. The New Network Policy Wizard appears.
2. On the Specify Network Policy Name And Connection Type page, type a policy name, and then select a network access server type. For IPsec enforcement, select Health Registration Authority. For 802.1X or VPN enforcement, select Remote Access Server. If you plan to use the Health Credential Authorization Protocol (HCAP) to integrate with Cisco Network Access Control, select HCAP Server. Click Next.



EXAM TIP

For the exam, don't worry about HCAP. Instead, focus on the other enforcement types.

- On the Specify Conditions page, click the Add button to create any conditions you require, as shown in Figure 8-15, and then click Next. The most useful conditions for NAP are the following:
 - Health Policies** Specifies that a client must meet the conditions specified in a health policy.
 - NAP-Capable Computers** Allows you to match either computers that support NAP or computers that do not support NAP.
 - Operating System** Allows you to apply the network policy to NAP-capable computers with specific operating system version numbers or computer architectures (such as 32-bit or 64-bit computers). This condition is not used as frequently as Health Policies and NAP-Capable Computers.
 - Policy Expiration** Use this to apply different conditions based on the current date and time. For example, if you are creating a temporary policy that applies only for the next week, you would add the Policy Expiration condition. You should create a second network policy to apply after the Policy Expiration condition expires.
 - Windows Groups, Machine Groups, And User Groups** These conditions determine the computer or user's Active Directory group membership.

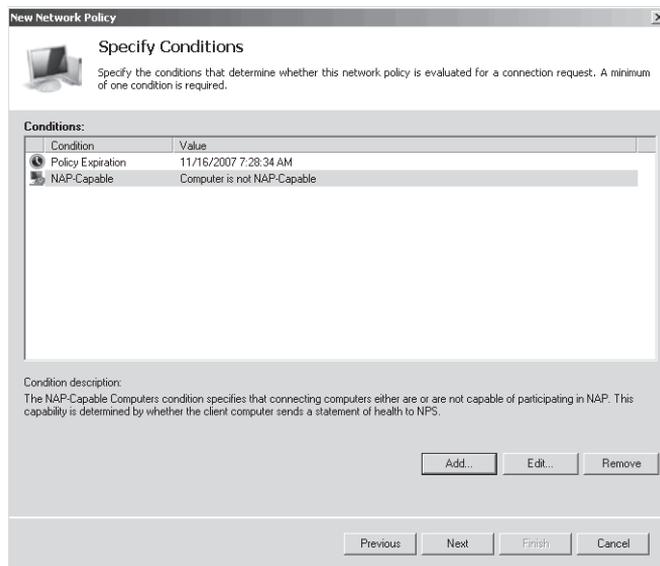


FIGURE 8-15 Specifying network policy conditions

4. On the Specify Access Permission page, select Access Granted. You should never select Access Denied for NPS policies because doing so prevents the health check from occurring. Click Next.
5. On the Configure Authentication Methods page, click Next. For NAP, authentication methods are selected in the Connection Request Policy.
6. On the Configure Constraints page, click Next. NAP rarely uses constraints, although you could use the Day And Time Restrictions constraints to apply the network policy at only specific times.
7. On the Configure Settings page, select NAP Enforcement. Then, select one of the following options and click Next:
 - **Allow Full Network Access** Grants full access. Use this option if you are creating a network policy for healthy computers.
 - **Allow Full Network Access For A Limited Time** Grants full access up to a specific date and then restricts access to the selected Remediation Server Group. Use this option during the initial NAP deployment if you want to offer a grace period for noncompliant computers. When selecting this option, click the Configure button to select a remediation server group and specify a troubleshooting URL. If you select this option when using VPN enforcement, VPN clients are disconnected when the expiration time is reached.
 - **Allow Limited Access** Limits access to the servers specified in the selected remediation server group. Use this option when creating a network policy for noncompliant computers. When selecting this option, click the Configure button to select a remediation server group and specify a troubleshooting URL.

NOTE THE EXTENDED STATE SETTING

This page also includes the Extended State setting. This setting is used only if you are using HCAP with Cisco Network Admission Control. Otherwise, leave this setting as the default.

8. On the Completing New Network Policy Wizard page, click Finish.
9. Right-click the network policy and choose Move Up or Move Down to prioritize it. Higher network policies are evaluated first, and the first network policy with criteria that match a client is applied.

Configuring NAP for Monitoring Only

During your initial NAP deployment, you should allow noncompliant computers to connect to all network resources, even if they fail the NAP health check. To do this, modify the non-compliant health policy to allow full network access by following these steps.

1. In Server Manager, select Roles\Network Policy And Access Services\NPS\Policies \Network Policies. In the details pane, double-click the noncompliant policy. For example, if you specified “NAP IPsec with HRA” as the name on the Select Network Connection Method For Use With NAP page of the NAP Wizard, the network policy for noncompliant NAP clients would have the name “NAP IPsec with HRA Noncompliant.”
2. Click the Settings tab, and then select NAP Enforcement.
3. In the network policy properties dialog box, in the details pane, select Allow Full Network Access, and then click OK.

To re-enable NAP enforcement, change the setting to Allow Limited Access.



REAL WORLD

Tony Northrup

Security risks have an annual cost, which is the potential damage if something bad happens, multiplied by the chance that it'll happen within a year. Countermeasures such as NAP reduce some security risks, which saves companies money. All countermeasures, including NAP, carry their own costs, though, and those costs must not exceed the money saved by reducing risk.

NAP might not actually cost you any cash; if you've already deployed the necessary infrastructure, you don't have to buy anything new. It still has the potential to cost your organization dearly, however. For example, imagine staying up late one night to enable NAP and prevent noncompliant computers from connecting to your network. The next morning, the president of your company has an important presentation in front of investors, but she can't access the network because her laptop is non-compliant. Not properly impressing the investors could have a significant negative consequence for your company.

In the real world, the costs will be less dramatic and more nagging. An administrative assistant will log on from home to file his expense report and discover that he hasn't installed the latest updates. A salesperson will attempt to connect to your VPN to send an urgent technical question to an engineer, but she'll give up when the remediation server directs her to update her antivirus definitions. Little amounts of lost productivity and missed opportunities add up in a way that you can't calculate.

When implementing any countermeasure, including NAP, do so gracefully and slowly. All countermeasures include some inconvenience, and that inconvenience has a cost. Always remember that the primary goal of information technology is to make employees more productive.

NAP Logging

NAP logging allows you to identify noncompliant computers. This is particularly important during the initial stages of a NAP deployment, when you will be using NAP only to gather information about the compliance level of the computers on your network. Using NAP logging, you can identify computers that are not compliant and resolve the problem before you enable NAP enforcement and prevent the computer from connecting to your network. NAP logging also enables you to identify computers that would be unable to connect to the network if NAP enforcement were enabled.

To configure NAP logging, right-click Roles\Network Policy And Access Services\NPS, and then choose Properties. On the General tab, select or clear the Rejected Authentication Requests and Successful Authentication Requests check boxes, as shown in Figure 8-16.

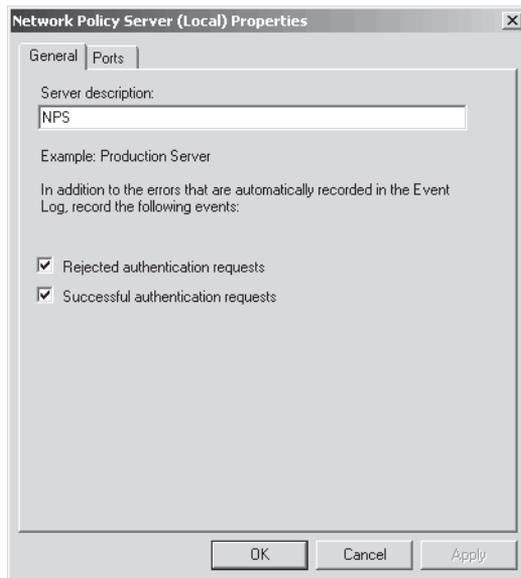


FIGURE 8-16 Configuring NPS logging

On the NAP server, you can use the Windows Logs\Security event log, available in Server Manager at Diagnostics\Event Viewer\Windows Logs\Security, to view NPS events. These events will reveal which NAP clients are not compliant. Figure 8-17 shows an event that indicates a computer that failed to pass the NAP health check. Figure 8-18 shows a computer that passed the NAP health check.

On clients running Windows Vista, Windows 7, Windows Server 2008, and Windows Server 2008 R2 NAP, use the Event Viewer console to examine the Applications and Services Logs\Microsoft\Windows\Network Access Protection\Operational log. On NAP clients running Windows XP With Service Pack 3, use the Event Viewer console to examine the System event log.

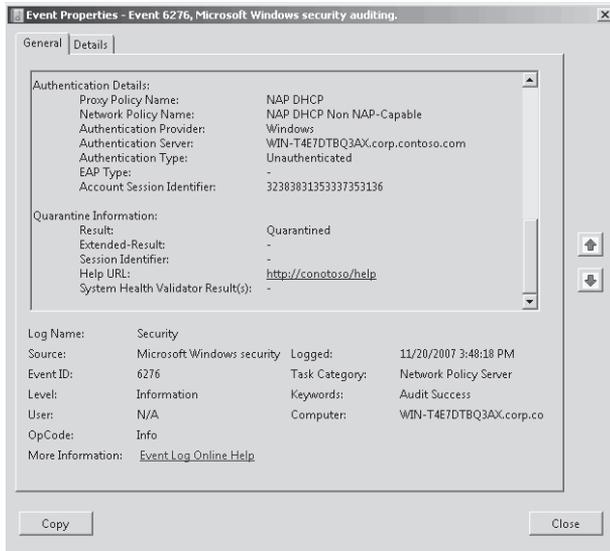


FIGURE 8-17 A failed NAP health check

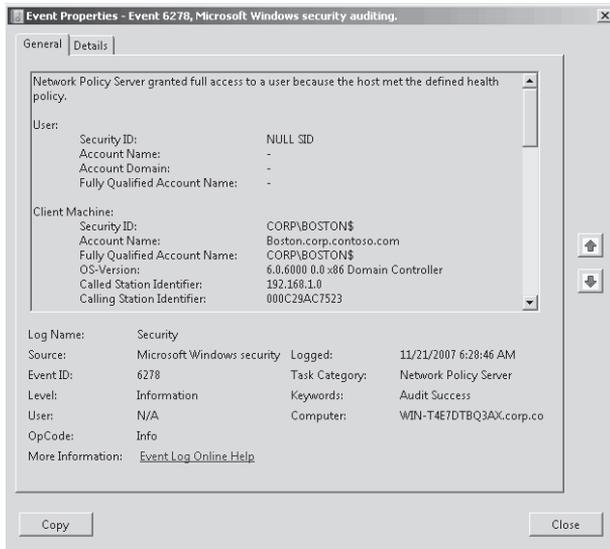


FIGURE 8-18 A successful NAP health check

Additionally, you can enable tracing for the Network Access Protection Agent service to gather extremely detailed information, which is typically required only when troubleshooting complex network problems. To enable tracing, run the following command:

```
netsh nap client set tracing enable level=verbose
```

The trace log files are stored in the %SystemRoot%\Tracing folder.

For more information about NAP logging, refer to Chapter 7. (NAP performs the same logging when used as a RADIUS server.)

PRACTICE Configuring DHCP NAP Enforcement

In this practice, you configure DHCP NAP enforcement and test it with both a compliant and noncompliant NAP client. Although DHCP NAP enforcement is the least secure, it is used as an example here because the configuration is the easiest to demonstrate. To prepare for the exam, you should configure each of the different NAP enforcement types in a lab environment.

Configuring NAP DHCP enforcement is a common scenario for networks with hardware that does not support 802.1X and where IPsec is not available. Although DHCP enforcement does not prevent knowledgeable attackers from connecting to your network, it does inform users who are unaware that their computers do not meet your security requirements of the problem. In production environments, you would typically implement NAP for monitoring only before enabling NAP enforcement.

EXERCISE 1 Adding the NPS and DHCP Server Roles

In this exercise, you add the Network Policy And Access Services and DHCP Server roles to Dcsv1. If either of these roles already exists (for example, if you added one or both in a previous exercise), remove the roles before continuing.

1. Configure Dcsv1 with a static IP address of 192.168.1.2, a subnet mask of 255.255.255.0, and a DNS server address of 192.168.1.2. You can use a different IP address for Dcsv1 as long as you replace all instances of 192.168.1.2 in this practice with Dcsv1's IP address. Start Hartford, and verify that it is a member of the domain and can communicate with Dcsv1.
2. In Server Manager, on Dcsv1, select Roles. In the details pane, click Add Roles. The Add Roles Wizard appears.
3. If the Before You Begin page appears, click Next.
4. On the Select Server Roles page, select the Network Policy And Access Services and DHCP Server check boxes. If the roles are already installed, remove them first, and then return to this step. Click Next.
5. On the Network Policy And Access Services page, click Next.
6. On the Select Role Services page, select the Network Policy Server check box. Click Next.
7. On the DHCP Server page, click Next.
8. On the Network Connection Bindings page, click Next.
9. On the IPv4 DNS Settings page, click Next.
10. On the IPv4 WINS Settings page, click Next.

11. On the DHCP Scopes page, click Add. Complete the Add Scope dialog box, as shown in Figure 8-19. Name the scope **NAP Clients**. Provide an IP address range of 192.168.1.10 to 192.168.1.100. If you are using a different IP address for Dcsrv1, specify an IP address range on the same subnet. In the Subnet Mask box, type **255.255.255.0**. In the Default Gateway box, type **192.168.1.1** (even though that IP address does not exist). In the Subnet Type list, select Wireless. Selecting Wireless simply specifies a shorter lease duration, which requires NAP clients to process any health policy updates more regularly. Click OK, and then click Next.

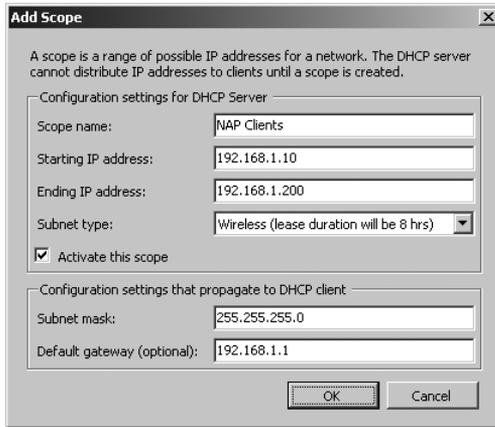


FIGURE 8-19 Configuring a DHCP scope

12. On the Configure DHCPv6 Stateless Mode page, click Next.
13. On the IPv6 DNS Settings page, click Next.
14. On the Authorize DHCP Server page, click Next.
15. On the Confirmation page, click Install.
16. On the Results page, click Close.

The DHCP and core NPS service are installed.

EXERCISE 2 Configuring NAP on the DHCP Server

In this exercise, you must configure NAP on the DHCP server to enforce health checks before assigning client computers an IP address that provides unlimited network access.

1. In Server Manager on Dcsrv1, select Roles\Network Policy And Access Services\NPS. If the node does not appear, close and re-open Server Manager.
2. In the details pane, under Standard Configuration, in the drop-down list, select Network Access Protection (NAP), and then click Configure NAP.
3. On the Select Network Connection Method For Use With NAP page, under Network Connection Method, select Dynamic Host Configuration Protocol (DHCP). Click Next.

4. On the Specify NAP Enforcement Servers Running DHCP Server page, click Add. In the New RADIUS Client dialog box, type **Dcsrv1** in the Friendly Name box and type Dcsrv1's IPv4 address (**192.168.1.2**) in the Address box. Click OK, and then click Next.
5. On the Specify DHCP Scopes page, click Next to apply NAP to all DHCP scopes.
6. On the Configure Machine Groups page, click Next to apply the policy to all users.
7. On the Specify A NAP Remediation Server Group And URL page, click New Group. In the New Remediation Server Group dialog box, type a Group Name of DHCP Remediation Servers. Then, click Add and provide a Friendly Name of NAP and Dcsrv1's IPv4 address (**192.168.1.2**). Click OK twice. Notice that you can also type a troubleshooting URL in this dialog box if you had set up a webpage for this purpose and added that server to the remediation server group. For now, type a troubleshooting URL of **http://contoso/help**. Although this URL will not work, it will allow you to see how the troubleshooting URL is used. Click Next.
8. On the Define NAP Health Policy page, click Next to accept the default settings.
9. On the Completing NAP Enforcement Policy And RADIUS Client Configuration page, click Finish.
10. In Server Manager, select Roles\Network Policy And Access Services\NPS\Policies\Connection Request Policies. Verify that the NAP DHCP policy exists and that it is the first policy listed. If other NAP connection request policies exist, remove them. Similarly, if other network policies exist, you should remove them, too.

Now you need to enable NAP enforcement on the DHCP server:

1. In Server Manager, select Roles\DHCP Server*<Computer Name>*\IPv4. Then right-click the node, and choose Properties.
2. In the Network Access Protection tab, click Enable On All Scopes, and then click Yes. Then select Restricted Access, and click OK.

EXERCISE 3 Configuring NAP Client Group Policy Settings

After configuring the NPS server, you must configure client computers for NAP by following these steps:

1. Click Start, Administrative Tools, and then Group Policy Management. The Group Policy Management console appears.
2. Right-click Group Policy Management\Forest\Domains*<Domain Name>*\Default Domain Policy, and then click Edit. The Group Policy Management Editor console appears.
3. Select the Computer Configuration\Policies\Windows Settings\Security Settings\Network Access Protection\NAP Client Configuration\Enforcement Clients node.
4. In the details pane, double-click DHCP Quarantine Enforcement Client. Select the Enable This Enforcement Client check box, and then click OK.

5. Select the Computer Configuration\Policies\Windows Settings\Security Settings\System Services node. Then, in the details pane, double-click Network Access Protection Agent. Select the Define This Policy Setting check box, and then select Automatic. Click OK.
6. Select the Computer Configuration\Policies\Administrative Templates\Windows Components\Security Center node. In the details pane, double-click Turn On Security Center. Select Enabled, and then click OK.

EXERCISE 4 Testing a Noncompliant Client

In this exercise, you connect a noncompliant computer to the network and determine whether it receives an IP address intended for compliant or noncompliant computers.

1. On Hartford, open a command prompt with administrative credentials and run the command **gpupdate /force**. This retrieves the updated Group Policy settings from the domain controller, verifying that the changes you made for NAP clients are applied correctly. Verify that the Network Access Protection Agent service is started.
2. On Hartford, run the command **netsh nap client show state** to verify that the DHCP Quarantine enforcement agent is enabled. If it is not, run the command **netsh nap client set enforcement 79617 enable** to manually enable it.
3. Disable any DHCP servers other than Dcsrv1. If you are using virtual machines, you can create a virtual network and connect both Dcsrv1 and Hartford to the virtual network.
4. Connect Hartford to the same network as Dcsrv1.
5. On Hartford, open a command prompt with administrative privileges. Then, run the following commands to retrieve new IP address settings from the DHCP server:

```
ipconfig /release  
ipconfig /renew
```

6. The client computer should display a new IP address configuration, with an IP address of 192.168.1.10 and a subnet mask of 255.255.255.255. Because the subnet mask is invalid (it should be 255.255.255.0), this indicates that the client computer failed the NAP health check.
7. At a command prompt, run the command **route print**. In the IPv4 Route Table, you should see a route with a Network Destination of 192.168.1.2. This address corresponds to the remediation server you configured.
8. At a command prompt, run the command **ping 192.168.1.2** (the IP address of Dcsrv1). Dcsrv1 should respond to the ping, verifying that the remediation server is accessible.
9. At a command prompt, run the command **ping 192.168.1.1**. The command fails with a Transmit Failed error because there is no valid route to the destination.
10. Notice that a notification bubble appears in the system tray, indicating that there was a problem. Click the link to view the details of the error. Notice that the error specifies that Windows did not detect an antivirus program. Click the More Information button to attempt to open the <http://contoso/help> page. Click Close.

11. On Dcsrv1, check the System event log. Find the event indicating that the client computer failed the NAP health check. If you had implemented NAP in monitoring-only mode, this would be the only sign that a computer did not meet the health requirements.

EXERCISE 5 Updating a Health Policy

In this exercise, you change the health policy to allow the client computer to pass the health check.

1. On Dcsrv1, in Server Manager, select Roles\Network Policy And Access Services \NPS\Network Access Protection\System Health Validators\Windows Security Health Validator\Settings. In the details pane, double-click Default Configuration.
2. On the Windows 7/Windows Vista tab, clear the An Antivirus Application Is On check box. Then, clear the Automatic Updating Is Enabled check box. Click OK.

The Hartford client computer will be able to pass the remaining health validation tests.

EXERCISE 6 Testing a Compliant Client

In this exercise, you connect a compliant computer to the network and determine whether it receives an IP address intended for compliant or noncompliant computers.

1. On Hartford, open a command prompt with administrative privileges. Then, run the following commands to retrieve new IP address settings from the DHCP server:

```
ipconfig /release  
ipconfig /renew
```

The client computer should display a new IP address configuration, with an IP address of 192.168.1.10, a subnet mask of 255.255.255.0 and a default gateway of 192.168.1.1. Because the subnet mask is now valid, it will be able to connect to other computers on the subnet (if any were available). A notification bubble will also appear, indicating that you have met the network's requirements.

2. On Hartford, open Event Viewer and view the Applications and Services Logs\Microsoft \Windows\Network Access Protection\Operational log. Examine the events for both the unsuccessful and successful NAP health checks.
3. On Dcsrv1, open Event Viewer and view the Windows Logs\Security log. Examine the events for both the unsuccessful and successful NAP health checks.

You can now remove NAP from Dcsrv1 and remove the DHCP enforcement client configuration from Hartford.

Lesson Summary

- Network Access Protection (NAP) allows you to verify that computers meet specific health requirements before granting them unlimited access to your internal network. You can enforce NAP by using IPsec, 802.1X access points, VPN servers, or DHCP servers.

- When deploying NAP, plan to implement it in monitoring-only mode first. This will allow you to identify and fix noncompliant computers before preventing them from connecting to your network.
- You can use Server Manager to install and configure Network Policy Server.
- Although the Configure NAP Wizard performs much of the configuration, each of the different NAP enforcement methods requires customized configuration steps.
- Before NAP takes effect, you must configure NAP clients. Additionally, when using IPsec enforcement, you must configure a health requirement policy.
- By default, NAP adds events to the Security event log on the NAP server each time a computer passes or fails a NAP health check. You can use the Security event log for auditing and to identify noncompliant computers that require manual configuration to become compliant.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 2, "Configuring Network Access Protection." The questions are also available on the companion CD if you prefer to review them in electronic form.

NOTE ANSWERS

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the "Answers" section at the end of the book.

1. You are currently configuring NAP enforcement in a lab environment. You need to create a network policy that prevents noncompliant computers from connecting to the network. How should you configure the network policy properties?
 - A. On the Settings tab, set NAP Enforcement to Allow Limited Access.
 - B. On the Overview tab, set Access Permission to Deny Access.
 - C. On the Constraints tab, set the Session Timeout to 0.
 - D. On the Settings tab, create an IP filter that drops all traffic.
2. You are a systems engineer developing NAP scenarios for future deployment within your organization. You want to configure a set of remediation servers that should be accessible for clients that do not support NAP. Which of the following do you need to do? (Choose all that apply.)
 - A. Create a health policy and set it to Client Fails All SHV Checks.
 - B. Create a network policy with a Condition type of NAP-Capable Computers.
 - C. Create a remediation server group with the servers that should be accessible.
 - D. Create a connection request policy with a Condition type of NAP-Capable Computers.

3. You are a systems administrator configuring NAP using DHCP enforcement. You plan to run NPS and DHCP on separate computers. Which of the following requirements do you need to fulfill? (Choose all that apply.)
- A. Configure a RADIUS proxy on the DHCP server.
 - B. Install NPS on the DHCP server.
 - C. Install HRA on the DHCP Server.
 - D. Configure Certificate Services on the DHCP server.

Chapter Review

To further practice and reinforce the skills you learned in this chapter, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenarios. These scenarios set up real-world situations involving the topics of this chapter and ask you to create a solution.
- Complete the suggested practices.
- Take a practice test.

Chapter Summary

- Windows Firewall is enabled by default to block most unwanted incoming connections. With additional configuration, you can limit the incoming connections that are allowed to specific subnets, user groups, or computer groups. Additionally, you can control which applications can initiate outgoing connections.
- Network Access Protection (NAP) is not enabled by default and requires complex planning and configuration to implement. After you deploy it, however, NAP provides network-level protection by allowing only clients that pass a health check to connect to your network.

Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- defense-in-depth
- firewall
- host route
- worm

Case Scenarios

In the following case scenarios, you will apply what you've learned about how to plan and deploy Windows Firewall and NAP. You can find answers to these questions in the "Answers" section at the end of this book.

Case Scenario 1: Evaluating Firewall Settings

You are a systems administrator for Fabrikam, Inc. Recently, your IT development department created a new client/server application that uses a web service. Your manager asks you to interview key people and then come to his office to answer his questions about the changes you will need to make to the Windows Firewall configuration.

INTERVIEWS

Following is a list of company personnel interviewed and their statements:

- **Developer** "It's a web service application, but it doesn't use IIS. Instead, it's its own service and listens for connections on TCP port 81. We need the server part of the application installed on Server1, and all client computers in the Accounting department should receive the client application. The client application just connects to the server on TCP port 81."
- **Lead systems engineer** "We use the default settings for Windows Firewall, so just let me know what I need to change."

QUESTIONS

Answer the following questions for your manager:

1. What type of firewall rule will you need to create to Windows Firewall on Server1?
2. What type of firewall rule will you need to create on the Windows Vista client computers in the Accounting department?

Case Scenario 2: Planning NAP

You are a systems administrator at Contoso, Ltd., an enterprise that manufactures large-scale farm equipment. Last night the news carried a story of corporate espionage—and your organization was the victim. According to the story, an employee of your biggest competitor gained access to your internal network six months ago, stole confidential plans for new equipment, and used them to improve their own designs. Last week, a disgruntled employee contacted the media and told the entire story.

Apparently, your competitor's employee waited patiently at a coffee shop near your offices. When he saw someone come in with a laptop and a Contoso badge, he waited for the employee to connect to the wireless network. He then exploited a known network vulnerability (which had been fixed several months earlier but had not been updated on the employee's computer) in the user's computer running Windows XP to install a tool that would automatically gather and forward documents from your company's internal network.

Your Chief Executive Officer (CEO) blames your Chief Security Officer (CSO), who in turn holds your Chief Information Officer (CIO) responsible. The CIO blames your manager, and your manager needs your help to create a plan to prevent this from happening again.

QUESTIONS

Answer the following questions for your manager:

1. Why would the attacker have been able to exploit a network vulnerability? How can that be prevented?
2. Is there some way we could have prevented the malware application from transmitting the confidential documents to a server on the Internet?
3. We can never guarantee that mobile computers will receive updates and won't be infected. After all, some of our staffers stay disconnected from the internal network for weeks at a time. So how can we keep these computers from connecting to our internal network and potentially doing damage?
4. If we suddenly turn on NAP, won't that cause problems for many of our client computers? How can we prevent that?
5. Which NAP enforcement method should we use?

Suggested Practices

To help you successfully master the exam objectives presented in this chapter, complete the following tasks.

Configure Firewall Settings

For this task, you should complete all four practices to gain real-world experience working with Windows Firewall.

- **Practice 1** Configure outbound filtering to block requests by default. Then, create firewall rules to allow common applications, including Internet Explorer and Microsoft Office, to connect to the Internet. Verify that Windows Update can retrieve updates from Microsoft.
- **Practice 2** Using a computer that is connected to the public Internet, enable firewall logging. Wait several hours, and then examine the firewall log. What types of requests were dropped? What might have happened if the firewall were not enabled?
- **Practice 3** On your organization's production network, examine the inbound firewall rules. How can you adjust the scope of these rules to minimize security risks?
- **Practice 4** Register for and watch the "Windows Vista Firewall And IPSec Enhancements" presentation by Steve Riley at <https://msevents.microsoft.com/CUI/Register.aspx?EventID=1032298288>.

Configure Network Access Protection

For this task, you should complete all six practices to gain experience using Network Access Protection in a variety of scenarios.

- **Practice 1** In a lab environment, deploy NAP using 802.1X, VPN, and IPsec. First, deploy NAP in monitoring-only mode. Then, switch to NAP enforcement.
- **Practice 2** Create a webpage that you could specify in the Troubleshooting URL, providing all the information the user of a noncompliant computer needs to remedy a problem and connect to the network.
- **Practice 3** Create a NAP test environment, including remediation servers. Using a noncompliant computer and any NAP enforcement technique, verify that you can bring the computer into compliance using just the resources provided by your remediation servers.
- **Practice 4** Watch the “Security and Policy Enforcement: Network Access Protection” presentation by Graziano Galante at <http://www.microsoft.com/emea/spotlight/sessionh.aspx?videoid=491>.
- **Practice 5** Watch the “NAP using DHCP in Windows Server 2008 R2” presentation by Kunal D. Mehta at <http://www.youtube.com/watch?v=iRtsj3BbwVs>.
- **Practice 6** Watch the “NAP Network Access Protection Demo” at <http://www.youtube.com/watch?v=DoO-x5MSsKw>.

Take a Practice Test

The practice tests on this book’s companion CD offer many options. For example, you can test yourself on just the content covered in this chapter, or you can test yourself on all the 70-642 certification exam content. You can set up the test so that it closely simulates the experience of taking a certification exam, or you can set it up in study mode so that you can look at the correct answers and explanations after you answer each question.

MORE INFO PRACTICE TESTS

For details about all the practice test options available, see “How to Use the Practice Tests” in this book’s Introduction.

Index

A

- ABE (Access-based Enumeration), 573
- aborting TCP sessions, 12
- access control lists (ACLs) and NAP enforcement, 447
- Access-based Enumeration (ABE), 573
- ACK flag, 11–12
- AcknowledgmentNumber* fields in TCP headers, 12
- ACLs (access control lists) and NAP enforcement, 447
- Active Directory
 - DNS zone integration with, 179
 - partitions. *See* partitions
- Active Directory Diagnostics Data Collector Set, 543
- Active Directory Domain Services (AD DS)
 - as SRV-aware application, 195
 - backing up BitLocker with, 581
 - deploying DNS servers in, 137–138
 - user accounts, 185
- Active Directory Domain Services server role, 137
- Active Directory forests, and DNS forwarders, 148
- Active Directory replication
 - forcing, 176
 - scope, setting, 180–181
 - tools for, 176
- AD DS. *See* Active Directory Domain Services (AD DS)
- AD DS Installation Wizard, 137
- ad hoc wireless networks, 352, 362. *See also* wireless networks
- Add Roles Wizard, 139
- address blocks, 53
 - determining size of, 56–59
 - host capacity of, 59
 - host capacity of, determining (practice exercise), 80–81
 - reference chart for, 54–56
 - size requirements, determining, 59–61
 - subnets and, 53
- address classes, 43
- address exclusions, 270–271
- address leases, 256
 - deleting, 273
 - durations, changing, 273
 - for remote access clients, 257
 - reservations for, 271–272
- Address Resolution Protocol. *See* ARP (Address Resolution Protocol) protocol
- address space, 53
 - dividing, calculations for, 68–76
 - subnetting. *See* subnets
- aging
 - defined, 196
 - enabling, 196
 - properties, setting, 198
- AH (Authentication Header) protocol, 321
- alias (CNAME) resource records, 192–193
- aliases for Network Monitor IP addresses, 560
- alternate configurations, assigning (practice exercise), 37
- answer files, creating for Dcpromo, 140
- APIPA addresses
 - automatic updating and, 166
 - defaulting to, 27
 - defined, 24
 - as problem indicator, 28
 - repairing connection when present, 29
- application directory partitions. *See* partitions
- Application event log
 - troubleshooting WSUS with, 510
 - WSUS server health, adding to, 508
- application layer (Layer 7), 13
- application layer protocols, 13
- approving security updates
 - overview of, 504–506
 - practice exercise for, 515–516

ARP (Address Resolution Protocol) protocol

- ARP (Address Resolution Protocol) protocol
 - defined 33
 - monitoring traffic via, 35
- Arp utility, 33–34
- asymmetric encryption. *See* key pairs
- auditing events, 523
- authentication
 - with AH (Authentication Header) protocol, 321
 - block period, setting, 364
 - caching, preventing, 364
 - for DirectAccess, 407
 - event forwarding and, 534
 - exempting, 332
 - with IPsec (Internet Protocol Security); IPsec Policies
 - logging attempts, 359–361
 - preference order, setting, 333
 - requiring, 333
 - restricting connections by type of, 397
 - vs. securing connections, 320
 - specifying method, 333
 - for VPNs (virtual private networks), 392
 - for wired networks, 363–364
 - for wireless networks, 351, 353–356
- Authentication Exemption rule, 332
- Authentication Header (AH) protocol, 321
- authoritative DNS servers, 122
- authorizing connections, 436–438
- Automatic Private IP Addressing addresses. *See* APIPA (Automatic Private IP Addressing) addresses
- Automatic Updates client. *See* Windows Update client

B

- backing up. *See also* shadow copies; Windows Server
 - Backup
 - automatically, 618–619
 - certificates, 237
 - from command prompt, 620–621
 - DNS zones, 234
 - file encryption keys, 574
 - practice exercise for, 625–626
 - restoring after, 621–623
 - restoring after, when Windows will not start, 624
 - scheduling, 618–619
 - system state, 621
- BaseNetworkTShoot filter for Network Monitor, 560

- bidirectional model of networking, 3
- binary notation
 - compared to other notations, 48
 - converting to decimal notation, 44–46
 - for subnet masks, 43
- BitLocker
 - backing up computer before enabling, 581
 - configuring on computer without TPM chip, 579
 - enabling, 581–582
 - Group Policy settings, 579
 - modes for, 579
 - overview of, 578
 - performance impact of, 578
 - recovery password, setting, 582
 - turning off, 582
 - volume integrity, checking, 581
 - when to use, 578
 - Windows editions containing, 579
- BitLocker-To-Go, 582
- blocking outbound connections
 - overview of, 435
 - practice exercise for, 441–442
- blocking reauthentication, 364
- b-node (NetBIOS), 118
- booting from recovery DVD, 624
- BranchCache
 - client configuration, 610
 - configuring, 609–610
 - Distributed Cache mode, 607
 - Group Policy settings, 610
 - Hosted Cache mode, 608–609
 - overview of, 587, 606
 - when to use, 606
- bridging network connections, 18–20
- broadcasts
 - address for, 59
 - limiting, with subnets, 64
 - overview of, 25

C

- CA (Certification Authority), 456–457
- cache locking, 151
- cache poisoning
 - Arp, 34
 - preventing, 151

- Cache.dns file
 - overview of, 128
 - viewing contents of, 144
 - caching, DNS, 131–132
 - caching-only servers
 - configuring, 142–143
 - when to use, 143
 - calculating subnet addresses, 68–76
 - caller ID, restricting connections by, 397
 - canonical names, 192–193
 - capturing network traffic
 - from command prompt, 557–558
 - from command prompt, practice exercise for, 564
 - filtering when, 560–561
 - with friendly aliases, 560
 - without full Network Monitor installation, 558
 - with Network Monitor, 555–558
 - certificate revocation list. *See* CRL (certificate revocation list)
 - certificates
 - backing up, 237
 - DirectAccess setup, 413
 - DirectAccess setup, practice exercise for, 417–418
 - IPsec authentication with, 323
 - Certification Authority (CA), 456–457
 - Change.txt file for WSUS, 510
 - Checksum fields
 - in IP headers, 9
 - in TCP headers, 13
 - classification management, 596–598
 - Classless Inter Domain Routing (CIDR) notation, 43
 - Client For Microsoft Networks, 18
 - clients. *See* network clients
 - CNAME resource records, 192–193
 - collusion, 577
 - computer groups
 - client-side targeting, 502–503
 - creating, 502
 - server-side targeting, 502–503
 - for WSUS, configuring, 502–503
 - computer names, creating, 160
 - conditional forwarding, 149–150
 - Configure A DNS Server wizard, 183
 - connection request policies, 358–359
 - Connection Security Rules, 316, 320
 - configuring for NAP, 458–459
 - creating, 330–334
 - encryption, requiring, 336
 - exporting, 320
 - IPsec settings for, 334–336
 - naming, 334
 - practice exercise for, 341–342
 - profiles, 334
 - setting by protocols/ports, 334
 - types of, 331
 - connectionless communication, 13
 - connection-oriented communication, 11
 - connection-specific suffixes, 161–162
 - connectivity. *See also* network connections
 - configuring, 17
 - Network Connections area, 17
 - testing, 397–402
 - troubleshooting, 2
 - country domains, 121
 - Create IP Security Rule Wizard, 326–330
 - CRL (certificate revocation list)
 - overview of, 410
 - practice exercise for, 418
 - publishing, 419
 - CScript, as default application, 649
- ## D
- Data Collector Sets
 - built-in, 543–544
 - creating, 544–546
 - customizing, 546
 - data sources, adding to, 546
 - importing/exporting, 549
 - Logman command-line tool for, 549
 - overview of, 543
 - reports, viewing, 547
 - running, practice exercise for, 551–552
 - specifying location for, 545
 - starting, 544, 546
 - stopping, 544
 - templates for, 545
 - types of, 546
 - viewing reports for, 544
 - data link layer (Layer 2)
 - overview of, 6
 - standards defined at, 6
 - switches, and Network Monitor, 556–557
 - Data Recovery Agents. *See* DRAs (Data Recovery Agents)
 - datagrams. *See* packets; TCP packets

DataOffset fields in TCP headers

- DataOffset fields in TCP headers, 12
- day and time restrictions on network connections, 396
- Dcdiag tool, 176
- Dcpromo.exe
 - answer files, creating, 140
 - overview of, 137
 - on Server Core installations of Windows Server 2008, 140–141
- deadlines for security updates, defining, 505
- declining security updates, 506
- default gateway. *See also* routers
 - configuring, 24
 - defined, 53, 291
 - for DHCP scopes, defining, 262
 - network ID, 51
 - overview of, 50
 - in routing table, 303
 - unconfigured, 51
- Default Response Rule in IPsec, 326
- Default User class, 276
- defense-in-depth, 437
- delaying DHCP responses, 280
- delegating DNS zones
 - implementing, 214–215
 - structure of, 213–214
 - when to do, 213
- delegating subdomains, 123
- Delegation Signer (DS) records, 227–232
- demand-dial routing, 299–300
 - enabling, 386–387, 392
 - filters, configuring, 300
 - overview, 299–300
- deploying DNS servers
 - on domain controllers, 137–138
 - on Server Core installations of Windows Server 2008, 140–141
 - on stand-alone servers, 139
- deploying NAP (Network Access Protection), 450
- deprecated IPv6 address state, 91
- DestinationAddress fields
 - in Ethernet headers, 7
 - in IP headers, 9
- device map, 16
- devolution of DNS suffixes, 163
- DFS (Distributed File System)
 - command-line configuration, 602
 - failover clustering, configuring, 604
 - health monitoring, configuring, 603
 - installing, 600
 - namespace configuration, 600–601
 - overview of, 599
 - polling configurations, 601
 - practice exercise for, 610–611
 - referral order, 601
 - replication, generating health reports for, 603
 - root permissions, configuring, 602
 - testing propagation, 603
 - testing propagation, practice exercise for, 613
- DFSCmd command-line tool, 602
- DFSUtil command-line tool, 602
- DHCID records, 284
- DHCP (Dynamic Host Configuration Protocol)
 - address assignment process, 254–255
 - Delay Configuration setting, 280
 - enabling, practice exercise for, 268
 - ICS (Internet Connection Sharing) and, 374
 - MAC filtering, 279–280
 - name protection, 284
 - negotiation, example of, 255
 - options classes, 276
 - options for, 257
 - options inheritance, 274
 - scope options, 274–275
 - user classes, 276–279
 - vendor classes, 276
 - for VPNs (virtual private networks), 393
- DHCP clients
 - Default User class, 276
 - DNS address configuration, 158
 - network broadcasts and
 - overview of, 25
 - troubleshooting, 29
- DHCP messages, 254–256
- DHCP Server role
 - adding, practice exercise for, 266–268, 476–477
 - overview of, 258–265
- DHCP servers
 - acknowledgment message, 255
 - address leases. *See* address leases
 - authorizing, 265
 - backup, 281
 - discovery of, 255
 - DNS suffix, setting, 259
 - DnsUpdateProxy security group, adding, 284–285
 - dynamic DNS updates, 282–283
 - exclusion ranges, 270–271

- installing/configuring, 258–265
- IPv6 protocol and, 254
- NAP configuration, practice exercise for, 477–478
- NAP enforcement and, 448, 460–461
- for NAT (Network Address Translation), configuring, 377–378
- network connection bindings, configuring, 259
- options inheritance, 274
- releasing IP addresses, 256
- scopes, 257, 261–262
- on Server Core installation, 285
- split-scope configuration, 281
- WINS server configuration, 260
- DHCPv6
 - stateless mode, configuring, 262–264
 - stateless mode vs. stateful mode, 267
- dialog box notifications for events, 525
- dial-up connections
 - advantages and disadvantages, 383
 - architecture of, 382
 - caller ID, 397
 - configuring, 385–391
 - creating manually, 390–391
 - enabling ICS (Internet Connection Sharing) for, 375–376
 - encryption, configuring, 390
 - filtering traffic for, 389
 - IP configuration, 386
 - modem configuration, 387–388
 - multilink settings, 397
 - RADIUS server configuration, 388–390
 - routing, 299–300
 - Routing And Remote Access Services configuration, 385–386
- DifferentiatedService* fields in IP headers, 9
- digital signatures, 225
- DirectAccess
 - advantages of, 406, 411
 - authentication, 407
 - certificate setup, 413
 - certificate setup, practice exercise for, 417–418
 - client configuration, practice exercise for, 420–421
 - configuration videos, 415
 - configuring, practice exercise for, 415–418
 - connection types, 407–408
 - end-to-edge protection, 407
 - end-to-end protection, 407
 - firewall configuration, 412–413
 - firewall configuration, practice exercise for, 416
 - hardware and software requirements, 410–411
 - ICMP configuration, practice exercise for, 416
 - IPsec protection, enabling, 409
 - on IPv4 networks, 408–409
 - IPv6 protocol and, 406, 408–409
 - limitations of, 412
 - name resolution, 409
 - NLS (Network Location Server), 410
 - operating system compatibility, 406
 - overview of, 405–406
 - propagating settings, 414
 - server configuration, practice exercise for, 418–420
 - setting up, 413–414
 - setting up, practice exercise for, 419–420
- directory partitions. *See* partitions
- DirQuota command-line tool, 590
- disabling IPv6 protocol, 134
- disk quotas
 - command-line configuration, 590
 - configuring with Windows Explorer, 590–592
 - creating, 589–590
 - event triggers for, 591
 - Group Policy settings, 592
 - overview of, 587
 - performance impact of, 570
 - Quota Management console, 587–590
 - templates for, 588–589
 - thresholds, adding, 589
- Distributed File System. *See* DFS (Distributed File System)
- Distributed Scan Server role service, 633
- dividing address spaces, 68–76
- DNS cache locking, 151
- DNS client cache
 - clearing, 168
 - flushing, 168
 - overview of, 131–132
 - practice exercise for, 168–169
 - viewing, 168
- DNS clients
 - DNSSEC configuration, 240–241
 - DNSSEC configuration, practice exercise for, 245
 - dynamic update settings, 184–185
 - single-label name tags, 199
 - updating, 165–167
 - updating manually, 185
- DNS console, clearing DNS server cache with, 132

DNS forwarders

- DNS forwarders
 - conditions, configuring, 149–150
 - configuring, 145–148
 - NAT (Network Address Translation), 378
 - security and, 146
 - when to use, 146–148
- DNS Manager
 - aging/scavenging settings, 196
 - Conditional Forwarders container, 149–150
 - connecting to DNS server from, 141
 - opening, 138
 - partitions, browsing, 205
- DNS namespace. *See also* DNS zones
 - canonical names, 192–193
 - overview of, 121
 - private, 122
 - root hints, 127–128, 144
 - subdomains, 121
 - top-level domain names, 121
 - trailing dot, 121
 - zones, 123
- DNS queries
 - components of, 124
 - with delegated subdomains, 213–214
 - example of, 128–130
 - forwarding. *See* forwarders
 - IPsec protection, enabling, 409
 - process for, 124–127
- DNS recursion, 127–128, 144
- DNS resolution. *See also* DNS suffixes; name resolution via DNSSEC, 227–232
 - example of, 128–130
 - methods of, 124
 - NSEC (Next Secure) records, 233
 - overview of, 120
 - security and, 146, 151
 - WINS resolution as backup for, 195
- DNS resolvers, 123, 125
- DNS server cache, 132
- DNS Server Core role, 141, 378
- DNS servers. *See also* deploying DNS servers
 - adding to list of, 158
 - address configuration, 158
 - authoritative, 122
 - caching-only, configuring, 142–143
 - caching-only, when to use, 143
 - configuring manually, 23, 139
 - connecting to from DNS Manager, 141
 - defined, 122
 - deploying, 137–141
 - domain controller deployment, 183
 - dynamic update settings, 165, 184–185
 - forwarders, configuring, 145–148
 - function of, 122
 - IP configuration, 143
 - manually configuring, 23, 139
 - multihomed, 143
 - primary, naming, 186
 - promoting to domain controllers, 137
 - properties, configuring, 143–150
 - refresh interval, setting, 187
 - reviewing configuration, practice exercise for, 153–154
 - root hints, 127–128, 144
 - single-label names for, 199
 - source port randomization, 151
 - trust anchors, configuring, 238–240
 - trust relationships, 226
- DNS socket pooling, 151
- DNS suffixes, 160–161. *See also* DNS resolution
 - configuring for DHCP servers, 259
 - connection-specific, 161–162
 - custom search lists, 164
 - devolution, 163
 - Group Policy configuration, 164
 - search lists, configuring, 162–163
- DNS zone replication
 - default partitions, 205
 - partitions, pattern of, 206
 - scope, setting, 180–181, 207–209
 - to Windows 2000 Server, 207
- DNS zones. *See also* DNS namespace
 - @ symbol in, 193
 - Active Directory-integrated, 179
 - Active Directory-integrated, advantages of, 205
 - aging, enabling, 196
 - backing up, 234
 - creating, 177
 - creating, practice exercise for, 243
 - defined, 177
 - delegating, 213–215
 - forward lookup, 123, 139, 181–182
 - GlobalNames, 199–202
 - host records, 165
 - masters, 179
 - name server (NS) records, 188, 221

- naming, 183
 - overview of, 123
 - pointer records. *See* pointer records
 - primary, 178
 - public keys in, 224, 228. *See also* public keys
 - read-only domain controllers and, 179
 - refresh interval, setting, 187
 - replication. *See* DNS zone replication
 - resource records. *See* resource records
 - retry interval, setting, 187
 - reverse lookup, 123, 181–182
 - round robin distribution, 193
 - scavenging, enabling, 196
 - secondary, 178–179
 - secondary, setting expiration time for, 187
 - serial number for, 186
 - signing files, 237–238
 - signing files, practice exercise for, 244
 - standard, 180
 - structuring, 213
 - stubs, 179
 - time stamping records, 198
 - type, choosing, 178–180
 - verifying, practice exercise for, 156
 - WINS records in, 195
 - zone transfers, 186, 188, 210–212
- Dnscmd tool
- key generation with, 235–237
 - overview of, 132
 - signing zone files with, 237–238
 - subcommands for, 176
- DNSKEY resource records, 224
- DNSSEC
- certificates, backing up, 237
 - client configuration for, 240–241
 - configuring, 234–240
 - configuring, practice exercise for, 245
 - defined, 223
 - Delegation Signer (DS) records, 227–232
 - IPSec authentication, enforcing, 241
 - key generation switches/options, 235–237
 - key rollover, configuring, 234–235
 - name resolution with, 227–232
 - NSEC (Next Secure) records, 233
 - Resource Record Signature (RRSIG) records, 227–232
 - trust relationships, 226
 - validating remote domains with, 226
- DnsUpdateProxy security group, 284–285
- Documents folder, encrypting, 576
- domain administrator account, creating (practice exercise), 154–155
- domain controllers
- creating, practice exercise for, 152–153
 - deploying DNS servers on, 137–138, 183
 - firewall configuration, 432
 - partitions in, 205
 - read-only, 179
 - setting replication scope for, 180–181
- Domain firewall profile, 431
- domain isolation, 332
- domain names
- fully qualified (FQDNs), 121
 - fully qualified (FQDNs), generation of, 124
 - registrars, 122
 - top-level, 121
 - trailing dot, 121
- Domain network profile, 15
- domain suffixes. *See* DNS suffixes
- DomainDnsZones partition, 205
- domains, adding computers to (practice exercise), 155–156
- dotted-decimal notation
- block size, determining, 58–59
 - compared to other notations, 48
 - converting subnet masks into, 43
 - converting subnet masks into (practice exercise), 78
 - converting to binary notation, 44–46
 - converting to slash notation, 49–50
 - overview of, 40
- DRAs (Data Recovery Agents)
- adding manually, 578
 - assigning two people to, 577
 - configuring, 577–578
 - configuring, practice exercise for, 582–583
 - overview of, 577
 - practice exercise for, 584
- DS (Delegation Signer) records, 227–232
- dynamic DNS updates, 165, 184–185
- configuring for DHCP servers, 282–283
 - forcing with Ipconfig, 185
- Dynamic Host Configuration Protocol. *See* DHCP (Dynamic Host Configuration Protocol); DHCP clients; DHCP servers

E

- EAP-MSCHAP v2 authentication, 356
- EFS (Encrypting File System)
 - applying, 574
 - applying, practice exercise for, 583
 - encryption key, backing up, 574, 576
 - Group Policy configuration, 575–577
 - indexing encrypted files, 577
 - locking desktop to prevent bypass of, 576
 - operating system compatibility, 573
 - overview of, 573
 - performance impact of, 570
 - recovering encrypted files, 577–578
 - sharing files protected by, 574–575
 - smart cards, requiring, 576
- email notifications for events, 525
- Encapsulating Security Payload (ESP) protocol, 321
- Encrypting File System. *See* EFS (Encrypting File System)
- encryption. *See also* BitLocker; EFS (Encrypting File System); key pairs
 - for dial-up connections, 390
 - with ESP (Encapsulating Security Payload) protocol, 321
 - for event forwarding, 531
 - for HRA server, 454–455
 - IPsec and, 321–322
 - for Network Monitor, 557
 - Network Monitor and, 522
 - for wireless networks, 350
 - for WSUS, 509–510
- endpoints, specifying, 333
- enumerating subnets, 68–76, 96, 99–101
- ESP (Encapsulating Security Payload) protocol, 321
- Ethernet frames
 - analyzing, 559
 - example of, 6–7
 - filtering, 560–561
 - IP protocol in, 8–10
 - vs. packets, 559
- Ethernet protocol, 6–7
- EthernetType* fields in Ethernet headers, 7
- event forwarding
 - configuring collecting computer, 528
 - configuring forwarding computer, 527
 - firewall exceptions for, 534
 - HTTPS protocol for, 531
 - normal delay in, 532
 - overview of, 526
 - practice exercise for, 536
 - privileges, verifying, 533
 - protocols and ports for, 526
 - query validity, verifying, 534–535
 - troubleshooting, 532–535
 - on Windows XP/Windows Server 2003 computers, 526
- event logs, filtering, 524–525
- event monitoring. *See also* event forwarding
 - event subscriptions, creating, 528–531
 - overview of, 523
 - practice exercises for, 535–537
- Event Viewer
 - Applications And Services Logs, 523
 - automatically responding to events in, 525–526
 - custom views, 525
 - event subscriptions, creating, 528–531
 - event subscriptions, practice exercise for, 536–537
 - event triggers, practice exercise for, 535–536
 - filtering event logs, 524–525
 - opening, 523
 - saving filters, 525
 - viewing events in, 523–524
- exclusion ranges
 - creating, 270–271
 - creating, practice exercise for, 286
- exporting
 - certificates, 237
 - DNS zones, 234
 - printers, 647
 - security update metadata, 508

F

- failover clustering for DFS servers, 604
- fault tolerance
 - DHCP server backups and, 281
 - NAP (Network Access Protection) and, 450
- fields. *See* IP headers; TCP headers; *specific fields*
- File And Printer Sharing
 - as default service, 18
 - enabling/disabling locally, 15

- file classification
 - properties for, 596
 - rules, creating, 597–598
 - scheduling, 598
- file permissions
 - Access-based Enumeration (ABE), 573
 - applying, 572–573
 - command-line configuration, 573
 - default, 571
 - precedence when in conflict, 572
- file screening
 - configuring, 594
 - performance impact of, 570
- file security, 570. *See also* BitLocker; EFS (Encrypting File System); NTFS file permissions
- File Services server role, 586–587
- file sharing, 38. *See also* folder sharing
- Filter Action Wizard, 339–340
- filtering
 - event logs, 524–525
 - inbound traffic, 432–434
 - inbound traffic, practice exercise for, 440–441
 - IP traffic. *See* IPsec (Internet Protocol Security); IPsec Policies; IPsec rules
 - network captures, 560–561
 - outbound traffic, 434–435
 - outbound traffic, practice exercise for, 441–442
- filters, 393–394
- FIN flag, 12
- firewall exceptions
 - for DirectAccess, 412–413
 - enabling, 32, 393–394
 - for event forwarding, 527, 534
 - ICMP and, 38
 - overview of, 32
 - practice exercise for, 38
 - for WSUS, 499–500
- firewall rules
 - actions available for, 433
 - automatically enabled, 441
 - connection authorization, configuring, 436–438
 - creating, 433–434
 - default, 432
 - for inbound traffic, 433–434
 - for inbound traffic, practice exercise for, 440–441
 - for outbound traffic, 434–435
 - for outbound traffic, practice exercise for, 441–442
 - port-based, 439–440
 - scope, configuring, 435–436
 - types of, 433
- firewalls. *See also* Windows Firewall with Advanced Security (WFAS)
 - defined, 431
 - domain controller configuration, 432
 - Group Policy configuration, 438
 - importance of, 431
 - profiles for, 431–432
 - scope, configuring, 435–436
 - server configuration, 432
 - verifying, 467
- flags. *See specific flags*
- flash drives, encrypting, 582
- flushing DNS client cache, 168
- folder sharing. *See also* DFS (Distributed File System)
 - from command prompt, 594–595
 - connecting to folders, 595
 - File Servers server role for, 586–587
 - mapping network drive for, 595
 - with Provision A Shared Folder Wizard, 593
 - from Windows Explorer, 592–593
- ForestDnsZones partition, 205
- forward lookup zones, 123
 - adding, 139
 - creating, 181–182
 - GlobalNames. *See* GlobalNames zone
 - naming, 183
 - overview of, 123
 - WINS servers in, 195
- forwarders
 - conditions, configuring, 149–150
 - configuring, 145–148
 - NAT (Network Address Translation), 378
 - security and, 146
 - when to use, 146–148
- forwarding events. *See* event forwarding
- FQDNs (fully qualified domain names), 121, 124
- fragmentation of packets, 9
- FragmentedFlag* fields in IP headers, 9
- frames. *See* Ethernet frames
- fully qualified domain names. *See* FQDNs (fully qualified domain names)

G

- gateway. *See* default gateway; routers
- gateway-to-gateway tunneling, 322
- geographical domains, 121
- global IPv6 addresses, 87–88
- GlobalNames zone
 - adding records, practice exercise for, 202
 - creating, practice exercise for, 201
 - deploying, 200
 - deploying, practice exercise for, 201
 - overview of, 199
 - populating, 200
 - testing, practice exercise for, 202
- glue records, 213
- GPOs (Group Policy Objects)
 - client-side targeting with, 503
 - creating Connection Security Rules with, 330–334
 - creating, practice exercise for, 338
 - defining Connection Security Rules with, 320
 - IPsec Policy in, 317, 324–326
 - NAP client configuration with, 463–465
 - WMI filters and, 438
- Group Policy
 - BitLocker policies, 579
 - BranchCache policies, 610
 - disk quota policies, 592
 - DNS suffix search lists in, 164
 - DNSSEC configuration policies, 240–241
 - EFS policies, 575–577
 - firewall policies, 438
 - LLMNR disabling policies, 116
 - printer policies, 646
 - refreshing, 414
 - SNMP policies, 563
 - Windows Update policies, 490–491
 - wired network policies, 363–364
 - wireless connection policies, 361–362
 - wireless connection policies, practice exercise for, 368–369
 - WSUS configuration, verifying, 511
- Group Policy Objects. *See* GPOs (Group Policy Objects)
- groups, computer, 502–503

H

- handshake process for TCP protocol, 11–12
 - headers. *See* *EthernetType* fields in Ethernet headers; IP headers; *specific headers*; TCP headers
 - health certificates, requiring for servers, 458–459
 - Health Policies template (NPS), 365
 - Health Registration Authority role service. *See* HRA (Health Registration Authority) role service
 - configuring, 456–457
 - installing, 453–455
 - health requirement policies
 - updating, 466–472
 - updating, practice exercise for, 480
 - hexadecimal numbering
 - overview of, 87
 - translating powers of 2 into, 97–98
 - h-node (NetBIOS), 119
 - HopLimit* field in IP headers, 10
 - host ID, 41–42. *See also* interface ID (IPv6)
 - host names, retrieving, 160
 - host records
 - defined, 165
 - function of, 190
 - host routes, 448
 - hostname command, 160
 - hosts, 445
 - Hosts file
 - vs. Lmhosts file, 132
 - location of, 131
 - updating, 131
 - HRA (Health Registration Authority) role service
 - configuring, 456–457
 - installing, 453–455
 - HTTPS protocol
 - for event forwarding, 531
 - tunneling protocol for (IP-HTTPS), 95
- I**
- IANA (Internet Assigned Numbers Authority), 121
 - IASNAP.log file, 361
 - lcacls command-line tool, 573

- ICANN (Internet Corporation for Assigned Names and Numbers), 121–122
- ICMP (Internet Control Message Protocol)
 - defined, 31
 - DirectAccess configuration, practice exercise for, 416
 - enabling request response, 295
 - exempting from authentication, 336
 - firewall exceptions, practice exercise for, 38
 - traffic, defined, 327
- ICS (Internet Connection Sharing)
 - architecture of, 374
 - configuring NAT with, 374–375
 - DHCP service for, 374
 - enabling for remote access connections, 375–376
 - internal network interface address, 374
 - for intranet servers, 375
 - IP addresses for, 375
 - overview of, 373
 - port number settings, 375
 - vs. Routing And Remote Access Services, 376
- Identification* fields in IP headers, 9
- idle time-outs, setting, 397
- IEEE 802.1X standards, 363–364
 - configuring, practice exercise for, 366–367
 - NAP enforcement, 446–447, 459–460
 - overview of, 363–364
- IGMP (Internet Group Management Protocol), 296, 298–299
- IIS (Internet Information Services), 496
- IKE (Internet Key Exchange) protocol, 321–322, 394
- importing printers, 647–648
- inbound traffic
 - filtering, 432–434
 - filtering, practice exercise for, 440–441
- indexing encrypted files, 577
- installing DNS servers. *See also* deploying DNS servers
 - on domain controllers, 137–138
 - on Server Core installations of Windows Server 2008, 140–141
 - on stand-alone servers, 139
- installing Network Monitor, 554
- installing printers
 - with Control Panel, 635–636
 - overview of, 634
 - practice exercise for, 651
 - with Print Management snap-in, 636–638
- installing WSUS (Windows Server Update Services)
 - planning for, 495–496
 - practice exercise for, 513–514
 - process for, 499–500
- interface ID (IPv6), 88
- international domains, 121
- Internet
 - preventing access to, Open Systems Interconnect (OSI) model, 4
 - overview of, 2–4
- Internet Assigned Numbers Authority (IANA), 121
- Internet Connection Sharing. *See* ICS (Internet Connection Sharing)
- Internet Control Message Protocol. *See* ICMP (Internet Control Message Protocol).
- Internet Group Management Protocol (IGMP), 296, 298–299
- Internet Information Services (IIS), 496
- Internet Key Exchange (IKE) protocol, 321–322, 394
- Internet layer (Layer 3), 7–8
- Internet printing
 - managing, 643–644
 - practice exercise for, 653
- Internet Printing Protocol (IPP), 633
- Internet Printing role service, 643
- Internet Protocol Security. *See* IPsec (Internet Protocol Security)
- intranet
 - example of, 292
 - servers, accessing Internet with, 375
- Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), 92–93
- IP addresses. *See also* IP protocol; IPv6 addresses; subnet masks
 - address block reference chart, 54–56
 - address blocks, 53
 - calculating for subnets, 68–76, 99–101
 - compared to ZIP codes, 41–42
 - DHCP assignment of, 254–255
 - DHCP configuration, 24–25
 - for dial-up connections, 386
 - for DNS requests, specifying, 143
 - exclusion ranges, 270–271
 - filtering Network Monitor by, 560
 - firewall scope, configurin, 436
 - for ICS (Internet Connection Sharing), 375

IP configuration

- IP addresses (*continued*)
 - leases, ending, 273
 - leasing, 256
 - leasing duration, 273
 - leasing, for remote access clients, 257
 - MAC address associations, 271–272
 - network and host IDs, 41–42
 - for network broadcasts, 59
 - private, 52, 372
 - public, 52
 - ranges, defining, 262
 - releasing, 256, 400
 - renewing, 400
 - reservations for, 271–272
 - resolving names. *See* DNS resolution; name resolution
 - for routers, 293, 301
 - scopes, 257
 - scopes, splitting, 281
 - static, 22–24
 - static, practice exercise for, 36
 - structure of, 40–41
 - subnet verification, 76
 - for subnets, 59
 - unicast, 51
- IP configuration
 - manually assigning, 22–24
 - manually assigning, practice exercise for, 36
 - practice exercise for, 36
 - renewing, 29
 - viewing, 20–21
- IP filter lists
 - adding IP addresses to, 327
 - filter actions, 328, 330
 - managing, 330
 - mirrored filters, 328
 - overview of, 318–319
 - predefined filters, 328
- IP Filters template (NPS), 365
- IP headers
 - Checksum* fields, 9
 - DestinationAddress* fields, 9
 - DifferentiatedService* fields, 9
 - FragmentedFlag* fields, 9
 - HopLimit* fields, 10
 - Identification* fields, 9
 - NextProtocol* fields, 9
 - overview of, 8–10
 - PayloadLength* fields, 10
 - SourceAddress* fields, 9
 - TimeToLive* fields, 9
 - TotalLength* fields, 9
- IP mapping, repairing, 34IP packets. *See* packets; TCP packets
- IP protocol. *See also* IPv6 protocol
 - alternate configurations, assigning
 - alternate configurations, practice exercise for, 37
 - communication with IPv6, 92–93
 - in Ethernet headers, 8–10
 - loopback addresses, 92
 - manually configuring, 22–24
 - overview of, 8
- IP registries, 52
- IP routing. *See* routing
- IP-address-to-name resolution, 121
- Ipconfig, 20
 - /release argument, 400
 - /renew argument, 29, 400
 - clearing/viewing DNS client cache with, 168
 - forcing DNS update with, 185
 - reading output (practice exercise), 101–102
 - renewing connection with, 29
 - running, 398–400
- IP-HTTPS protocol, 95
- IPP (Internet Printing Protocol), 633
- IPsec (Internet Protocol Security)
 - Authentication Header (AH) protocol, 321
 - authentication, 323, 329
 - client configuration for NAP, 457
 - connection establishment, 321–322
 - Connection Security Rules, 316, 320, 334–336
 - Default Response Rule, 326
 - default settings, 335
 - DNSSEC enforcement, 227, 241
 - enabling for NAP, 457
 - Encapsulating Security Payload (ESP) protocol, 321
 - encryption, 321–322
 - Internet Key Exchange (IKE) protocol, 321–322
 - Kerberos authentication, 323
 - NAP enforcement, 446
 - NAP enforcement, configuring, 453–459
 - negotiating security, 318
 - on non-Windows computers, 317
 - overview of, 317

- preshared keys, 323
 - requirements for, 446
 - Security Associations (SAs), 321
 - services provided by, 317
 - transport mode, 322
 - tunnel authorization, 336
 - tunnel mode, 322
 - for WSUS, 496
 - IPsec Policies. *See also* IPsec rules
 - assigning, 317
 - assigning multiple, 324
 - configuring, 326
 - creating, 325–326
 - creating, practice exercise for, 338
 - example of, 319
 - IP filter lists, 318–319, 330
 - Negotiate Security option, 318
 - predefined, 324
 - predefined filters, 328
 - testing, practice exercise for, 340–341
 - IPsec rules. *See also* Connection Security Rules; IPsec Policies
 - creating, 326–330
 - creating, practice exercise for, 338–340
 - filter actions, 339–340
 - filter actions, practice exercise for, 339–340
 - overview of, 318–319
 - IPv4 protocol. *See* IP protocol
 - IPv6 addresses. *See also* IP addresses
 - global, 87–88
 - link-local, 88–89
 - link-local, pinging (practice exercise), 102
 - overview of, 86–87
 - shortening, 87
 - site-local, as deprecated, 91
 - states of, 91
 - structure of, 87
 - syntax of, 86–87
 - unique local, 90–91
 - unique local, assigning (practice exercise), 102–103
 - IPv6 hosts
 - autoconfiguration flags, 264
 - stateful addressing, 264
 - IPv6 protocol. *See also* IP protocol
 - Arp cache poisoning and, 35
 - communication with IPv4, 92–93
 - configuring, 87
 - DHCP servers and, 254
 - DirectAccess and, 406
 - disabling, 134
 - DNS server settings, 264
 - header example, 9–10
 - ICMP and, 31
 - LLMNR (Link Local Multicast Name Resolution)
 - compatibility with, 116
 - loopback addresses, 92
 - manually configuring, 23
 - NAT and, 372
 - NetBIOS incompatibility with, 110, 120
 - overview of, 8
 - purpose of, 86
 - renewing configuration, 29
 - router compatibility, 92
 - subnets and, 96–101
 - transition technologies, 92–95, 408–409
 - ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) protocol, 92–93
 - Isolation rules, 332
 - iteration, 127
- ## K
- Kerberos authentication, 333
 - key pairs
 - digital signatures and, 225
 - generating, 235–237
 - Key Signing Keys and Zone Signing Keys, 228
 - overview of, 224
 - rollover configuration, 234–235
 - Key Signing Keys (KSKs), 228
 - generating, 235
 - rollover configuration, 234–235
 - keys. *See* private keys; public keys
- ## L
- L2TP (Layer Two Tunneling Protocol), 391
 - languages for security updates, configuring, 496
 - LANs, viewing device map, 16
 - latency
 - event subscriptions and, 530
 - VPNs and, 385

Layer 1 (physical layer)

- Layer 1 (physical layer), 6
- Layer 2 (data link layer)
 - overview of, 6
 - standards defined at, 6
 - switches, and Network Monitor, 556–557
- Layer 3 (Internet layer), 7–8
- Layer 4 (transport layer), 10
- Layer 7 (application layer), 13
- Layer Two Tunneling Protocol (L2TP), 391
- layered networking, pinging (practice exercise), 102
- leases
 - deleting, 273
 - durations, changing, 273
 - for remote access clients, 257
 - reservations for, 271–272
- limited connectivity with APIPA address, 29
- Line Printer Daemon (LPD) protocol, 633
- Link Layer Topology Discovery (LLTD) Mapper, 16
- Link Local Multicast Name Resolution. *See* LLMNR (Link Local Multicast Name Resolution)
- link-local addresses, 88–90
- LLMNR (Link Local Multicast Name Resolution)
 - advantages and disadvantages, 116
 - disabling, 116
 - Network Discovery and, 112
 - overview of, 110, 112
 - practice exercise for, 133–134
 - responses over IP protocols, 114–115
 - when to use, 112
- Lmhosts files
 - vs. Hosts file, 132
 - overview of, 117
- load balancing for SMTP mail servers, 193–194
- local IP addresses, configuring firewall scope for, 436.
 - See also* IP addresses
- location-aware printing, 639
- logging. *See also* monitoring
 - NAP (Network Access Protection), 474–476
 - Windows Firewall traffic, 439
- Logman command-line tool, 549
- loopback addresses
 - overview of, 92
 - in routing table, 303
- LPD (Line Printer Daemon) protocol, 633
- LPR Port Monitor feature, 634

M

- MAC addresses
 - associating with IP addresses, 271–272
 - in Ethernet headers, 7
 - filtering DHCP traffic by, 279–280
 - routing with, 294
- MAC filtering, 279–280
- mail exchange (MX) resource records, 193–194
- man in the middle attacks, preventing, 409
- manual backups, 617–618
- manually configuring IP addresses, 22–24
- mapping
 - network devices, 16
 - network drives, 595
- masks. *See* subnet masks
- master zones, 179
- Maximum Transmission Unit (MTU) packet size, 9
- MBSA (Microsoft Baseline Security Analyzer), verifying updates with, 497
- Media Streaming, enabling/disabling locally, 15
- Microsoft Baseline Security Analyzer (MBSA), verifying updates with, 497
- Microsoft Report Viewer Redistributable 2008, 499
- Microsoft Update, 489. *See also* WSUS (Windows Server Update Services)
- Microsoft Update Improvement Program, 501
- migrating printers, 647–648
- mirrored filters, 328
- m-node (NetBIOS), 119
- modem banks, 388
- modem configuration, 387–388
- monitoring. *See also* event monitoring; Network Monitor;
 - Performance Monitor
 - with Data Collector Sets, 543–549
 - NAP (Network Access Protection), 450, 472–473
 - ports, 556
 - printers, 650
 - with Reliability Monitor, 542–543
 - with Resource Monitor, 541
- MTU (Maximum Transmission Unit) packet size, 9
- multicast addresses in routing table, 303
- multihomed servers, configuring, 143
- Multilink And Bandwidth Allocation Protocol (BAP), 397
- MX (mail exchange) resource records, 193–194
- My Network Places, 116

N

- name protection, 284
- name resolution. *See also* DNS resolution
 - defined, 111
 - DirectAccess and, 409
 - with DNSSEC, 227–232
 - filtering Network Monitor by, 560
 - LLMNR (Link Local Multicast Name Resolution), 112–116
 - NetBIOS, 110, 116–120
 - NSEC (Next Secure) records, 233
 - overview of, 109
 - Windows methods for, 111–112
- Name Resolution Policy Table (NRPT)
 - adding rules to, 240–241
 - DirectAccess configuration, 408–409
 - DNSSEC requests through, 227–232
 - rules, viewing, 409
- name root servers, public, 2
- name server (NS) records
 - creating, 188
 - creating, practice exercise for, 221
- name servers. *See* DNS servers
- namespace, DNS. *See also* DNS zones
 - canonical names, 192–193
 - overview of, 121
 - private, 122
 - root hints, 127–128, 144
 - subdomains, 121
 - top-level domain names, 121
 - trailing dot, 121
 - zones, 123
- naming DNS zones, 183
- NAP (Network Access Protection)
 - 802.1X access point enforcement, 446–447
 - 802.1X configuration, 459–460
 - ACLs, quarantining with, 447
 - analyzing costs of, 473
 - architecture of, 445
 - autoremediation, enabling, 452
 - client configuration, 463–465
 - Configure NAP Wizard, 455–456
 - configuring, 451–453
 - connection process, 449
 - deploying, 450
 - DHCP configuration, practice exercise for, 477–478
 - DHCP enforcement, 448
 - DHCP enforcement, configuring, 460–461
 - enforcement, disabling, 472–473
 - enforcement types, 446
 - error code resolution settings, 467
 - fault tolerance, 450
 - health requirement policies, 466–472
 - health requirement policies, practice exercise for, 480
 - health state changes, 448
 - health validation, 448–449
 - implementation of, 445
 - installing, 450–451
 - IPsec enforcement, 446
 - IPsec enforcement, configuring, 453–459
 - logging, 361, 474–476
 - monitoring, 450
 - monitoring-only status, 472–473
 - network policy configuration, 470
 - noncompliant computers, identifying, 474–476
 - RADIUS servers as, 450
 - remediation, configuring, 468–470
 - Remote Desktop Gateway enforcement, 448
 - security updates, enforcing, 468
 - testing, practice exercise for, 479–480
 - verifying updates with, 497
 - VLANs, quarantining with, 447
 - VPN server enforcement, 447, 462
- NAS (Network Access Server), restricting connections
 - by port type, 397
- NAT (Network Address Translation). *See also* private IP addresses
 - client computers, configuring, 378
 - client configuration, practice exercise for, 380
 - configuring as server, practice exercise for, 379–380
 - configuring, with Internet Connection Sharing (ICS), 374–375
 - configuring, with Routing And Remote Access Services, 376–377
 - DHCP server, configuring, 377–378
 - DNS forwarding with, 378
 - as hardware vs. as server, 373
 - IPv6 and, 372
 - logging, configuring, 379
 - overview of, 372
- Neighbor Solicitation message, 91

Net Share command

Net Share command, 594–595

NetBIOS

- advantages and disadvantages, 120

- broadcasts over IPv4, 117

- configuring, 118

- defined, 116

- enabling/disabling, 110, 118

- history of, 110

- IPv6 incompatibility, 120

- Lmhosts files, 117

- methods of name resolution, 117

- node types, 118–119

- practice exercise for, 133–134

- responses over IP protocols, 116–117

- when to use, 116

- WINS servers, 117

NetBT. *See* NetBIOS

Netsh utility

- enabling wireless Single Sign On, 353

- IP configuration with, 24–25

- NAP client configuration, viewing, 464

- practice exercise for, 38

- syntax, 24

Netstat, identifying ports with, 439–440

Network Access Protection. *See* NAP (Network Access Protection)

Network Access Protection Agent service

- overview of, 463

- trace logging, 475–476

Network Access Server (NAS), restricting connections by port type, 397

network adapters, configuring for DHCP servers, 259

Network Address Translation. *See* NAT (Network Address Translation)

Network and Sharing Center

- Change Advanced Sharing Settings option, 15

- opening, 15

- See* Full Map option, 16

network authentication. *See* authentication

network broadcasts

- address for, 59

- limiting, with subnets, 64

- overview of, 25

network clients

- overview of, 18

- viewing, for connections, 17

Network Connection Details dialog box, 20

network connections

- alternate configurations

- authenticating. *See* authentication

- bindings, configuring, 259

- bridging, 18–20

- day and time restrictions, 396

- default components of, 17–18

- default configuration of, 23

- idle time-outs, setting, 397

- IP configuration, 20–21, 24–25

- IP configuration, practice exercise for, 36

- on-demand, 301

- properties of, 17, 20–21

- repairing, 29

- restrictions, configuring, 395–397

- security, 363–364

- sharing. *See* ICS (Internet Connection Sharing)

- testing connectivity, 397–402

- troubleshooting, 29–35

- verifying, practice exercise for, 38

Network Connections area, 17

network connectivity

- configuring, 17

- Network Connections area, 17

- testing, 397–402

- troubleshooting, 2

Network Discovery

- enabling/disabling locally, 15

- LLMNR (Link Local Multicast Name Resolution) and, 112

network ID, 41–42. *See also* subnet ID

- for default gateways, 51

- identifying, 42

- IP routing and, 50

- subnet verification, 76

network interface layer

- overview of, 6

- standards defined at, 6

- switches, and Network Monitor, 556–557

network layer (Layer 3), 7–8

network layers. *See also specific layers*

- numbering scheme, 5

- overview of, 2

Network Location Server (NLS), 410

Network Map, 16

- Network Monitor
 - analyzing data from, 558–559
 - capture file size, configuring, 556
 - capturing network traffic with, 555–557
 - command-line administration, 557–558
 - command-line administration, practice exercise for, 564
 - DHCP exchange display, 255
 - encryption, 522, 557
 - filtering data from, 560–561
 - installing, 554
 - IP addresses, translating, 560
 - Layer 2 switches and, 556–557
 - OneClick tool, 558
 - P-Mode capturing, 558
 - practice exercise for, 563–564
 - starting, 555
 - stopping data capture, 556
 - unscrambling packets, 522
 - Network Monitor 3 Driver, 554
 - network policies
 - configuring, 396–397
 - for NAP, 470
 - Network Policy And Access Services role (practice exercise), 476–477
 - Network Policy Server (NPS)
 - installing, practice exercise for, 366–368
 - Security event log, 359
 - templates for, 365
 - network prefix notation, 43, 97
 - network printers. *See* printers
 - network profiles, 15
 - network protocols, 18. *See* protocols; *specific protocols*
 - network segments, 6. *See also* packets; TCP packets
 - network services
 - overview of, 18
 - viewing, for connections, 17
 - networks. *See also* subnets; wireless networks
 - /n, determining size of, 56–57
 - host capacity, determining, 59
 - host capacity, determining (practice exercise), 80–81
 - size requirements, determining, 59–61
 - size requirements, determining (practice exercise), 82–83
 - New Connection Security Rule Wizard, 331–334
 - New Resource Record dialog box, 189
 - New Zone Wizard, 177–185
 - Next Generation TCP/IP stack, 5, 8, 92
 - Next Secure (NSEC) records, 233
 - NextProtocol* fields in IP headers, 9
 - NFS (Network File System) services, enabling, 587
 - NLS (Network Location Server), 410
 - NMCap command-line tool, 557–558
 - No-Broadcasts filter for Network Monitor, 560
 - node types for NetBIOS, 118–119
 - notifications
 - for disk quota thresholds, 589
 - for events, 525
 - for printer events, 644–645
 - for Windows Update, configuring, 490–491
 - for zone transfers, 212
 - NPS (Network Policy Server)
 - installing, practice exercise for, 366–368
 - Security event log, 359
 - templates for, 365
 - NRPT. *See* Name Resolution Policy Table (NRPT)
 - NS (name server) records, 188, 221
 - NSEC (Next Secure) records, 233
 - NTFS file permissions
 - Access-based Enumeration (ABE), 573
 - applying, 572–573
 - command-line configuration, 573
 - default, 571
 - precedence when in conflict, 572
 - numbering of network layers, 5
- O**
- octet values for subnet masks, 49, 54–55
 - octets, in IP addresses, 40
 - Offline Files
 - configuring, 604–606
 - encrypting, 577
 - Open Shortest Path First (OSPF) protocol, 296
 - Open Systems Interconnect (OSI) model, 4
 - Operational event log, troubleshooting with, 532
 - organizational domains, 121
 - OSI (Open Systems Interconnect) model, 4
 - OSPF (Open Shortest Path First) protocol, 296
 - outbound traffic, filtering, 434–435

P

- packet filtering
 - configuring, 397
 - for VPNs (virtual private networks), 393–394
- packet forwarding. *See also* routing
 - Layer 2 and Layer 3 addresses, 294
 - process for, 293
 - tracing, 294–295
- packet loss, tracking, 33
- packets. *See also* Ethernet frames; TCP packets
 - vs. Ethernet frames, 559
 - fragmentation, 9
 - Maximum Transmission Unit (MTU) size, 9
 - overview of, 3
- page file
 - configuring, 550–551
 - monitoring, 550
 - overview of, 549
- partitions
 - creating, 206, 209
 - creating, practice exercise for, 218
 - default, 205
 - defined, 205
 - domain, storing DNS data in, 207
 - FQDNs for, 205
 - permissions for creating, 209
 - referencing local server name for, 209
 - replication pattern, 206
 - setting replication scope for, 207–209
 - storing zones in, 209
 - storing zones in, practice exercise for, 218
 - for virtual memory, avoiding, 551
- passwords for wireless networks, 350–351
- PathPing utility
 - output example, 33, 294–295
 - practice exercise for, 306
 - testing network connectivity with, 33
 - as Tracert alternative, 402
 - Windows Firewall and, 32
- PayloadLength* field in IPv6 header, 10
- PEAP (Protected EAP) authentication, 355
- Performance Monitor
 - configuring, 540–541
 - counters, bolding, 539
 - overview of, 539
- page file, monitoring, 550
- virtual memory configuration, 549–551
- permissions
 - NTFS, 571–573
 - printer, 640
- persistent connection, configuring VPN as, 300
- Pfirewall.log file, 439
- physical layer (Layer 1), 6
- physical memory, analyzing usage, 549
- Ping utility
 - output example, 400–401
 - practice exercise for, 38
 - running, 400–401
 - testing network connectivity with, 30–31
 - troubleshooting, 192
 - Windows Firewall and, 32
- pinging link-local addresses (practice exercise), 102
- p-node (NetBIOS), 118
- pointer records
 - 99 in, 194
 - defined, 165
 - referencing in zone files, 194
 - update behavior for, 166–167
- Point-to-Point Tunneling Protocol (PPTP), 391
- poisoned ARP cache example, 34
- ports. *See* TCP ports
- postal codes, compared to IP addresses, 41–42
- power management, and Windows Update client, 491
- powers of 2, translating hexadecimal into, 56, 97–98
- PPTP (Point-to-Point Tunneling Protocol), 391
- preferred IPv6 address state, 91
- preshared keys, 323
- primary DNS server, naming, 186
- primary DNS zones, 178
- Print And Document Services server role
 - installing, 633–634
 - installing, practice exercise for, 650
- Print Management snap-in
 - installing, 633
 - printer driver installation with, 640–641
 - printer installation with, 636–638
 - printer sharing with, 638
- print server permissions, 640
- PrintBRM command-line tool, 647–648
- printer drivers
 - digitally signed, 641
 - finding, 641

- installing, 636–638, 640–641
- isolation of, 641
- for printer pools, 642
- printer pools, 632
 - configuring, 642
 - configuring, practice exercise for, 651
 - drivers for, 642
 - printing to, practice exercise for, 652
- printers
 - best practices for, 632
 - command-prompt management, 648–649
 - counters, viewing, 650
 - exporting, 647
 - filters for, 644–645
 - Group Policy settings, 646
 - Group Policy settings, practice exercise for, 652
 - importing, 647–648
 - installing, 634–636
 - installing, practice exercise for, 651
 - Internet management of, 643–644
 - Internet management of, practice exercise for, 653
 - location-aware printing with, 639
 - LPR Port Monitor feature, 634
 - monitoring, 650
 - network, finding, 635
 - notifications, configuring, 644–645
 - permissions, configuring, 640
 - print queues, 643
 - priorities, configuring, 643
 - scripts for, 648–649
 - sharing, 636, 638
- private domain namespace, 122
- Private firewall profile, 432
- private IP addresses, 52. *See also* NAT (Network Address Translation)
 - address ranges for, 372
 - unique local addresses, 90–91, 102–103
- private keys. *See also* public keys
 - digital signatures and, 225
 - location of, 224
 - overview of, 224
 - rollover configuration, 234–235
- Process Monitor, 522
- profiles, network, 15
- Protected EAP (PEAP) authentication, 355
- protected wireless networks, 362–363. *See also* wireless networks

- protocols. *See also specific protocols*
 - application layer, 13
 - defined, 2
 - network, 18
 - viewing, for connections, 17
- Provision A Shared Folder Wizard, 593
- PSH flag, 12
- Public firewall profile, 432
- Public Folder Sharing, 15
- public IP addresses, 52. *See also* global IPv6 addresses
- public key cryptography, 224–226. *See also* private keys; public keys
 - public keys. *See also* private keys
 - configuring for wireless networks, 352–353
 - digital signatures and, 225
 - Key Signing Keys and Zone Signing Keys, 228
 - overview of, 224
 - rollover configuration, 234–235
 - trust anchors, 226
 - trust anchors, configuring, 238–240
- publishing CRL (certificate revocation list), 419

Q

- QoS Packet Scheduler, setting as default, 18
- querying DNS servers. *See* DNS queries
- Quota Management console, 587–590
- quotas. *See* disk quotas

R

- RAC (Reliability Analysis Component), 543
- RADIUS Clients template (NPS), 365
- RADIUS servers
 - backups for, 360
 - configuring Windows Server as, 354–355
 - configuring, practice exercise for, 366–368
 - dial-up configuration, 388–390
 - log files, 360–361
 - logons, monitoring, 359–361
 - modem banks, connecting to, 388–390
 - as NAP (Network Access Protection) servers, 450
 - proxy configuration, 356–359
 - Security event log, 359

RD (Remote Desktop) Gateways and NAP enforcement

- RADIUS servers (*continued*)
 - Security event log, practice exercise for, 370
 - server groups, adding, 358
 - trace logging, enabling, 361
 - UDP port for, 356
 - wireless network authentication, 351, 354–356
- RD (Remote Desktop) Gateways and NAP enforcement, 448, 462
- read-only domain controllers (RODCs), 179
- recovering encrypted files. *See* DRAs (Data Recovery Agents)
- recovery DVD, booting from, 624
- recursion, DNS, 127–128, 144
- refreshing
 - Group Policy, 414
 - resource record timestamps, 198
- releasing IP addresses, 400
- Reliability Analysis Component (RAC), 543
- Reliability Monitor, 542–543
- reloading DNS zones, 212
- remediation networks. *See also* NAP (Network Access Protection)
 - automatic connection to, 452
 - infrastructure for, 445
- Remediation Server Groups template (NPS), 365
- remediation servers, 469
- remote access. *See also* dial-up connections; VPNs (virtual private networks)
 - clients, address leasing, 257. *See also* address leases
 - clients, filtering traffic to, 389
 - granting for VPN users, 394
 - overview of, 382
- Remote Desktop (RD) Gateways, and NAP enforcement, 448, 462
- remote hosts, pinging, 400–401
- Remote RADIUS Servers template (NPS), 365
- renewing
 - IP addresses, 400
 - IPv6 configuration, 29
- Repadmin tool, 176
- repairing network connection, 29
- replication, Active Directory
 - forcing, 176
 - scope, setting, 180–181
 - tools for, 176
- replication, DNS zone
 - default partitions, 205
 - partitions, pattern of, 206
 - scope, setting, 180–181, 207–209
 - to Windows 2000 Server, 207
- reports
 - for Data Collector Sets, viewing, 547
 - for WSUS, viewing, 506–508
- reservations, 271–272
- resolvers, DNS, 123, 125
- resolving names. *See* DNS resolution; name resolution
- Resource Monitor, 541
- Resource Record Signature (RRSIG) records, 227–232
- resource records, 123, 185–195. *See also* name server (NS) records; *specific resource records*
 - caching, 132
 - CNAME, 192–193
 - created automatically, 189
 - creating, 189
 - deleting outdated. *See* scavenging
 - dynamic update settings, 165
 - host (A or AAAA), 190–191
 - MX (mail exchange), 193
 - name protection, 284
 - owner of, 185
 - pointer (PTR). *See* pointer records
 - responsible person (RP), 187
 - scavenging, 196, 198–199
 - security permissions, 184
 - service (SRV), 195
 - Start of Authority (SOA) records, 185–186
 - time stamping, 196, 198
- responsible person (RP) resource record, 187
- restarting computer after updates, 490
- restoring backups
 - from shadow copies, 624–625
 - of files, 621–622
 - of volumes, 622–623
 - when Windows will not start, 624
- restrictions on network connections, day and time, 396
- reverse domains, 121
- reverse lookup zones, 123
 - creating, 181–182
 - overview of, 123
- reverse lookups, 121
- RIP (Routing Internet Protocol), 296–298
- RODCs (read-only domain controllers), 179
- role services, adding, 366
- roles, server. *See* server roles
- root hints, 127–128, 144
- root servers

- locating, 127–128
 - public, for troubleshooting, 2
 - updating list of, 144
 - rootkits, 578
 - round robin DNS, 193
 - Route Add command, 304
 - Route command, 301–304
 - Route Print command, 301, 303–304
 - routers. *See also* default gateway
 - advertisements, filtering, 298
 - configuring Windows Server as, 296–299
 - defined, 7, 291
 - hardware-based, benefits of, 291
 - IP addresses for, 293, 301
 - neighbors, listing, 298
 - routing protocols, 295–296
 - subnet configuration, 304
 - routing
 - demand-dial, 299–300, 386–387, 392
 - example of, 292
 - Layer 2 and Layer 3 addresses, 294
 - overview of, 50–51, 293
 - persistent vs. nonpersistent routes, 304
 - static, 301–305
 - Routing And Remote Access Services
 - advantages of, 376
 - configuring static routing with, 305
 - dial-up connection configuration, 385–386
 - enabling NAT (Network Address Translation) with, 376–377
 - installing, 296–297
 - overview of, 373
 - Routing Internet Protocol (RIP), 296–298
 - routing table, 293, 301, 305
 - RP (responsible person) resource record, 187
 - RST flag, 12
 - rules. *See specific rules*
 - rules, classification, 597–598
 - rules, security. *See* Connection Security Rules; IPsec Policies
- S**
- SAs (Security Associations)
 - details, viewing, 342
 - establishing, 321–322
 - overview of, 321
 - Scan Management snap-in, 634
 - scavenging
 - automatic, enabling, 198
 - defined, 196
 - enabling, 196
 - manually, 199
 - properties, setting, 198
 - scheduling backups, 618–619
 - Schtasks command-line tool, 526
 - scope, firewall, 435–436
 - scopes, DHCP, 257, 261–262
 - secondary DNS zones. *See also* zone transfers
 - creating, practice exercise for, 219
 - deploying, practice exercise for, 219
 - enabling transfers to, 188
 - enabling transfers to, practice exercise for, 220
 - expiration time, setting, 187
 - manually updating, 212
 - overview of, 178–179
 - Secure Socket Tunneling Protocol (SSTP), 391
 - Secured Password authentication, 356
 - Security Associations (SAs)
 - details, viewing, 342
 - establishing, 321–322
 - overview of, 321
 - Security event log
 - for Network Policy Server, 359
 - viewing, practice exercise for, 370
 - Security Rule Wizard, 326–330
 - security rules. *See* Connection Security Rules; IPsec Policies
 - security updates. *See also* NAP (Network Access Protection); WSUS (Windows Server Update Services)
 - approving, 504–506
 - approving, practice exercise for, 515–516
 - auditing, 496–497
 - automatic, configuring, 490, 501
 - auto-restart after, configuring, 490
 - deadlines, defining, 505
 - declining, 505–506
 - deployment process, 502
 - enforcing, 468
 - importance of, 488
 - installed, viewing, 512
 - Knowledge Base (KB) number, finding, 513
 - metadata, exporting, 508
 - for off-network computers, 495

segments

- security updates (*continued*)
 - removing, 506, 513
 - reports, viewing, 506–508
 - sorting list of, 505
 - storing locally, 495, 501
 - troubleshooting, 510–512
- segments, 6. *See also* packets; TCP packets
- SequenceNumber* fields in TCP headers, 12
- serial number for DNS zones, 186
- Server Cleanup Wizard (WSUS), 501
- Server Core installation of Windows Server 2008, 285
- server roles
 - Active Directory Domain Services, 137
 - DHCP Server, 258–268, 476–477
 - DNS Server Core, 141
 - File Services, 586–587
 - Print And Document Services, 633–634
 - Print And Document Services, practice exercise for, 650
- server firewall configuration, 432
- servers, DNS. *See* DNS servers
- Server-To-Server rule, 333
- service (SRV) resource records, 195
- services, network
 - overview of, 18
 - viewing, for connections, 17
- shadow copies
 - command-prompt management, 616
 - IDs, viewing list of, 616
 - overview of, 616
 - restoring from, practice exercise for, 624–625
 - Windows Explorer management, 616
- Shared Secrets template (NPS), 365
- sharing
 - EFS-protected files, 574–575
 - folders. *See* folder sharing
 - Internet connections. *See* ICS (Internet Connection Sharing)
 - printers, 636, 638
- SHAs (System Health Agents), 448
- SHVs. *See* System Health Validators (SHVs)
- signatures, digital, 225
- signing zone files
 - overview of, 237–238
 - practice exercise for, 244
- Simple Network Management Protocol (SNMP),
 - configuring, 561–563
 - single-label computer names, resolving.
 - See* GlobalNames zone
 - site-local addresses (deprecated), 91
 - 6to4 protocol, 93
 - slash notation for subnet masks, 43
 - compared to other notations, 48
 - converting to (practice exercise), 79
 - converting to dotted-decimal, 49–50
 - sliding windows in TCP protocol, 13
 - smart cards
 - requiring for EFS (Encrypting File System), 576
 - for wireless security, 351, 356
 - SMTP servers, 193
 - SNMP (Simple Network Management Protocol),
 - configuring, 561–563
 - SOA (Start of Authority) records. *See* Start of Authority (SOA) records
 - socket pooling, 151
 - sockets, 151
 - software notifications for updates, 491
 - SoftwareDistribution.txt file for WSUS, 510
 - SourceAddress* fields
 - in Ethernet headers, 7
 - in IP headers, 9
 - split-scope DHCP configuration, 281
 - SRV (service) resource record, 195
 - SSL certificates
 - installing, 455
 - for WSUS, 509–510
 - SSTP (Secure Socket Tunneling Protocol), 391
 - stability, monitoring with Reliability Monitor, 542–543
 - stack, TCP/IP, 5, 8, 92
 - stale resource records. *See* aging; scavenging
 - standard DNS zones, 180
 - Start of Authority (SOA) records, 185
 - opening, 186
 - serial number, changing, 186
 - stateless mode
 - defined, 262
 - for DHCPv6, 262–264
 - vs. stateful mode, 267
 - static DNS servers, 23
 - static IP addresses
 - assigning, 22–24
 - practice exercise for, 36
 - static keys for wireless encryption, 350–351

- static routing
 - configuring, 301–304
 - configuring, practice exercise for, 306–309
 - in routing table, 303
 - stub DNS zones
 - example of, 215–217
 - name resolution improvement with, 217
 - overview of, 179
 - uses for, 216
 - subdomains, 121. *See also* delegating DNS zones
 - subnet ID. *See also* network ID
 - in IPv6 addresses, 88, 96
 - number of bits, determining, 96–97
 - overview of, 65
 - subnet masks
 - address classes, as deprecated, 43
 - assigning (practice exercise), 77
 - commonly used, 46
 - for DHCP scopes, defining, 262
 - dotted-decimal, converting to (practice exercise), 78
 - dotted-decimal, determining size of, 58–59
 - enumerating address ranges, 68–76
 - identifying network ID with, 42
 - midrange values for, 46–47
 - notation comparison, 48
 - notations for, 43
 - octet values, 48–49
 - octet values, calculating, 54–55
 - overview of, 42
 - restricting broadcast traffic with, 64
 - size required, determining, 59–61
 - size required, determining (practice exercise), 82–83
 - slash notation, converting to (practice exercise), 79
 - slash notation, converting to binary, 43
 - slash notation, converting to dotted-decimal, 49–50
 - variable-length, creating, 67–68
 - subnets
 - address space calculations, 68–76
 - creating, 61–62
 - defined, 53
 - determining number of, 66
 - effect on address space, 66
 - enumerating address ranges, 68–76, 96, 99–101
 - equally-sized, creating, 66–67
 - hexadecimal increment, determining, 97–99
 - in IPv6 protocol, 96–101
 - IP address for, 59
 - network prefix, determining, 97
 - vs. networks, 54
 - ownership verification, 76
 - physical topology considerations, 64
 - reference table, creating, 68
 - routers on, 304
 - for variable numbers of hosts, 67–68
 - vs. virtual LAN (VLAN) switches, 64
 - subscribing to events
 - overview of, 528–531
 - practice exercise for, 536–537
 - suffixes, DNS. *See* DNS suffixes
 - switches as layer 2 devices, 6
 - SYN flag, 11
 - sysdm.cpl command, 160
 - System Diagnostics Data Collector Set, 544
 - System event log, 378–379, 395
 - system file encryption. *See* BitLocker
 - System Health Agents (SHAs), 448
 - System Health Validators (SHVs)
 - configuring, 466–467
 - overview of, 448
 - Windows SHV, 467–468
 - System Image Recovery wizard, 624
 - System Performance Data Collector Set, 543
 - System Properties dialog box, changing computer name in, 160
 - system reliability, monitoring, 542
 - System Stability Report, 542
 - system state, backing up, 621
- ## T
- Task Manager, viewing physical memory usage in, 549
 - Task Scheduler, 525–526
 - TCP flags, 12
 - TCP headers
 - AcknowledgmentNumber* fields, 12
 - Checksum* fields, 13
 - DataOffset* fields, 12
 - DstPort* fields, 12
 - SequenceNumber* fields, 12
 - SrcPort* fields, 12
 - vs. UDP headers, 13
 - Window* fields, 13
 - for WSUS, 509

TCP packets

TCP packets. *See also* packets
encapsulation of, 14
header example, 12
number of protocols in, 14

TCP ports
configuring, for DNS, 151
for event forwarding, 526
for HTTP traffic, 12
for ICS (Internet Connection Sharing), 375
identifying, 439–440
monitoring, 556
overview of, 10
for VPNs (virtual private networks), 392–393

TCP protocol
connection-oriented nature of, 11
handshake process, 11–12
overview of, 10

TCP sessions, aborting, 12

TCP/IP layers. *See* network layers

TCP/IP protocol encapsulation, 14

TCP/IP stack, 5, 8, 92

Telnet Server, installing (practice exercise), 440–441

Telnet services, installing (practice exercise), 337

templates, NPS, 365. *See also* NPS (Network Policy Server)

tentative IPv6 address state, 91

Teredo protocol, 94–95

Terminal Services Gateway. *See* RD (Remote Desktop)
Gateways, and NAP enforcement

testing network connectivity, 397–402

thumb drives, encrypting, 582

Time to Live (TTL) values, 132, 187

time stamping resource records, 198

TimeToLive fields in IP headers, 9

top-level domain names, 121

TotalLength fields in IP headers, 9

TPM (Trusted Platform Module) chips, 579

trace logging
enabling, 361
for Network Access Protection Agent service, 475–476

Tracert utility
-d switch, 32
DNS lookups, turning off, 401
output example, 32

practice exercise for, 306
running, 401–402
testing network connectivity with, 32
Windows Firewall and, 32

tracing packets, 294–295. *See also* PathPing utility;
Tracert utility

trailing dot in domain names, 121

transition technologies, 92–95

transport layer (Layer 4), 10

troubleshooting
DHCP connections, 29
network connections, 29–35
TCP/IP connectivity, 2

trust anchors
configuring, 238–240
location of, 238
overview of, 226

Trusted Platform Module (TPM) chips, 579

TTL (Time to Live) values, 132, 151

tunneling protocols, 92–95
authorizing, 336
overview of, 92–95

U

UDP headers, vs. TCP headers, 13

UDP protocol
header example, 13
overview of, 10

unicast addresses, 51

unique local addresses
assigning, 90–91
assigning (practice exercise), 102–103

updates. *See* security updates; WSUS (Windows Server Update Services)

updating DNS zones, 212

URG flag, 12

UrgentPointer field in TCP headers, 13

user classes
implementing, 276–279
populating, 279
predefined, 277

user credentials, preventing caching of, 364

user permissions. *See* NTFS file permissions

V

- validating digital signatures, 225
- validating DNS data, 226
- virtual LAN (VLAN)
 - configuring, 356
 - NAP enforcement and, 447
 - switches, vs. subnets, 64
- virtual machines, configuring, 307
- virtual memory
 - configuring, 550–551
 - overview of, 549
 - partitioning and, 551
- virtual private networks (VPNs)
 - advantages and disadvantages, 384–385
 - architecture of, 384
 - authentication, 392
 - client configuration, 394
 - client configuration, practice exercise for, 403
 - configuring router, 299–300
 - connecting to, 394
 - DHCP relay agent configuration, 393
 - enabling ICS (Internet Connection Sharing) for, 375–376
 - latency and, 385
 - NAP enforcement, 462
 - packet filter configuration, 393–394
 - port configuration, 392–393
 - protocols for, 391
 - reconnecting automatically, 394–395
 - routing configuration and, 301
 - server configuration, 392–393
 - server configuration, practice exercise for, 402–403
 - server enforcement of NAP, 447
 - troubleshooting, 395
- Visual Basic scripts for printing, 648–649
- VLAN (virtual LAN)
 - configuring, 356
 - NAP enforcement and, 447
 - switches, vs. subnets, 64
- volume backups, restoring, 622–623
- volume encryption. *See* BitLocker
- volume integrity, checking, 581
- Volume Shadow Copy. *See* shadow copies
- VPN Reconnect, 394–395
- VPNs. *See* virtual private networks (VPNs)
- VSSAdmin command-line tool, 616

W

- WANs (wireless area networks)
 - BranchCache and, 606
 - WSUS architecture for, 492
- WAPs (wireless access points). *See also* wireless networks
 - bridging into single network, 18–20
 - configuring, practice exercise for, 368
 - connecting to, practice exercise for, 369
 - as RADIUS clients, 355–356
 - security standards, 350–352
- Wbadmin command-line tool, 620–621
- Web pages, filtering Network Monitor by, 560
- WEP (Wired Equivalent Protection), 350
- Windows Firewall with Advanced Security. *See* WFAS (Windows Firewall with Advanced Security)
- wi-fi. *See* wireless networks
- Wi-Fi Protected Access (WPA), 351
- Window* fields in TCP headers, 13
- Windows 2000 Server, replicating DNS zones to, 207
- Windows Defender, 467
- Windows Event Collector
 - for event forwarding, 526, 528
 - practice exercise for, 536
- Windows Firewall with Advanced Security (WFAS), 527.
 - See also* firewalls
 - connection authorization, configuring, 436–438
 - Connection Security Rule configuration, 320
 - Connection Security Rule export, 320
 - creating inbound rules in, 433–434
 - creating outbound rules in, 435
 - default behavior of, 432, 434
 - DirectAccess configuration, 412–413
 - DirectAccess configuration, practice exercise for, 416
 - Group Policy configuration, 438
 - ICMP and, 32
 - IPsec settings for Connection Security Rules, 334–336
 - logging, enabling, 439
 - rule scope, configuring, 435–436
- Windows Internal Database, security update storage in, 495
- Windows logs, viewing, 523
- Windows PowerShell, configuring WSUS APIs with, 509
- Windows Remote Management
 - configuration, verifying, 532–533
 - for event forwarding, 526–527

Windows Search Service

- Windows Search Service, 587
- Windows Server
 - dashboard for, 15–16
 - Server Core installation, deploying DNS servers on, 140–141
- Windows Server 2003 File Services, 587
- Windows Server 2008
 - Name Protection feature, 284
 - as RADIUS proxy, 356–359
 - as RADIUS server, 354–355
 - as router. *See* routers
 - Server Core installation, installing DHCP on, 285
- Windows Server Backup
 - features, installing, 617
 - manually backing up with, 617–618
 - overview of, 617
 - scheduling backups with, 618–619
- Windows Server Update Services (WSUS). *See* WSUS (Windows Server Update Services)
- Windows System Health Validator (SHV), 467–468
- Windows Update client. *See also* security updates;
WSUS (Windows Server Update Services)
 - automatic updates, configuring, 490
 - blocking user access to, 491
 - client-side targeting, 491, 502–503
 - Group Policy settings, 490–491
 - history, viewing, 512
 - Install Updates And Shut Down option, disabling, 491
 - notification settings, configuring, 490–491
 - power management and, 491
 - server-side targeting, 502–503
 - troubleshooting, 511–512
 - verifying updates with, 496
 - WSUS server location, specifying, 490
- Windows XP computers, on Network Map, 16
- WindowsImageBackup folder, 618
- WinRM command-line tool, 527
- WINS servers
 - configuring, 260
 - DHCP specification for, 257
 - in DNS zones, 195
 - overview of, 117
- Wired Equivalent Protection (WEP), 350
- wired networks, requiring authentication for, 363–364
- wireless access points. *See* WAPs (wireless access points)
- wireless area networks. *See* WANs (wireless area networks)
- Wireless Diagnostics Data Collector Set, 544
- wireless networks. *See also* WAPs (wireless access points)
 - ad hoc mode, 352
 - authentication, 353–356
 - configuring, 361–362
 - connecting automatically, 361–362
 - deploying, 362–363
 - Group Policy settings, 361–362
 - Group Policy settings, practice exercise for, 368–369
 - infrastructure mode, 352
 - logons, monitoring, 359–361
 - password protecting, 350–351
 - public key infrastructure (PKI) configuration, 352–353
 - RADIUS authentication, 351
 - security standards, 350–352
 - Single Sign On, 353
 - troubleshooting, 544
 - unprotected, 350
 - WPA-EAP deployment, 362–363
- wizards. *See specific wizards*
- worms
 - defined, 431
 - propagation of, 434
- WPA (Wi-Fi Protected Access), 351
- WPA2 (IEEE 802.11i), 351
- WPA-EAP (Extensible Authentication Protocol), 351
 - autoenrollment, configuring, 352–353
 - configuring, practice exercise for, 366–370
 - deploying wireless networks with, 362–363
 - Security event log, 359
- WPA-Personal (WPA-PSK), 351
- WSUS (Windows Server Update Services). *See also* security updates
 - APIs, accessing, 509
 - approving updates, 504–506
 - approving updates, practice exercise for, 515–516
 - architecture of, 492–494
 - auditing updates, 496–497
 - client computers, configuring, 504, 515
 - command-line administration, 508–509
 - computer groups, configuring, 502–503
 - configuration overview, 500
 - database location, 496

- declining updates, 505–506
- exporting data from, 508
- firewall exceptions, 499–500
- health monitoring, configuring, 508
- IIS website for, 496
- inactive approvals, managing, 509
- installation planning, 495–496
- installing, 499–500
- installing, practice exercise for, 513–514
- languages, configuring, 496
- moving files, 508
- for multiple IT departments, 494
- for multiple offices, 492–494
- new features in, 489
- notification settings, configuring, 501
- overview of, 489
- personalizing, 501
- ports, configuring, 509
- products updated, configuring, 496, 501
- redundancy and, 492
- removing updates, 506, 513
- replication, configuring, 495
- reports, viewing, 506–508
- requirements for, 494–495
- secure communications for, 496
- Server Cleanup Wizard, 501
- Server Configuration Wizard, 501
- server registry key, updating, 508
- servers, viewing list of, 508
- for single offices, 492
- sorting updates, 505
- source and proxy server, configuring, 501
- SSL certificate configuration, 509–510
- synchronization, scheduling, 501

- synchronization, viewing report on, 508
- troubleshooting, 510
- update source, configuring, 495
- update storage, configuring, 495, 501
- verifying connection to, 511
- WAN architecture, 492
- Windows Internal Database storage, 495
- Windows Update client, 490–491
- WSUSUtil tool, 508–509
- WSUSUtil command-line tool, 508–509

Z

- ZIP codes, compared to IP addresses, 41–42
- zone ID (IPv6), 89–90
- zone replication
 - default partitions, 205
 - partitions, pattern of, 206
 - scope, setting, 180–181, 207–209
 - to Windows 2000 Server, 207
- Zone Signing Keys (ZSKs), 228, 234–235
- zone transfers, 179
 - enabling, 188, 211
 - forcing, 186
 - initiation of, 211
 - notifications for, 212
 - options for, 211
 - practice exercise for, 220
 - retry interval, setting, 187
 - when to use, 210
- zones, 176. *See also* DNS zones
- ZSKs (Zone Signing Keys), 228, 234–235

About the Authors



TONY NORTHRUP, MCITP, MCTS, MCPD, MCSE, and CISSP, is a Windows consultant and author living in Waterford, Connecticut, in the United States. Tony started programming before Windows 1.0 was released, but has focused on Windows administration and development for the last 15 years. He has written more than two dozen books covering Windows development, networking, and security. Among other titles, Tony is coauthor of the *Windows 7 Resource Kit* (Microsoft Press, 2009) and *Windows Server 2008 Networking and Network Access Protection (NAP)* (Microsoft Press, 2008). Tony has a technology blog at www.vistaclues.com.



J.C. MACKIN, MCITP, MCTS, MCSE, and MCDST, is a writer and Microsoft Certified Trainer. He is coauthor of many titles from Microsoft Press, including the *Self-Paced Training Kits* for Exams 70-291, 70-622, 70-643, and 70-685.