

# Windows Server® 2008 Networking and Network Access Protection (NAP)



Joseph Davies and  
Tony Northrup with the  
Microsoft® Networking Team

**PUBLISHED BY**

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

Copyright © 2008 by Microsoft Corporation

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: 2007940507

Printed and bound in the United States of America.

1 2 3 4 5 6 7 8 9 QWT 2 1 0 9 8 7

Distributed in Canada by H.B. Fenn and Company Ltd.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at [www.microsoft.com/mspress](http://www.microsoft.com/mspress). Send comments to [mspinput@microsoft.com](mailto:mspinput@microsoft.com).

Microsoft, Microsoft Press, Active Directory, ActiveX, Internet Explorer, MSDN, Outlook, SQL Server, Visual Basic, Visual Studio, Windows, Windows Media, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other product and company names mentioned herein may be the trademarks of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, e-mail 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:** Martin DelRe

**Developmental Editor:** Karen Szall

**Project Editor:** Maria Gargiulo

**Editorial Production:** Interactive Composition Corporation

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

**Cover:** Tom Draper Design

Body Part No. X14-31173

# Table of Contents

|                                  |       |
|----------------------------------|-------|
| Acknowledgments. ....            | xxiii |
| Introduction. ....               | xxv   |
| Document Conventions. ....       | xxv   |
| Reader Aids. ....                | xxv   |
| About the Companion CD-ROM. .... | xxvi  |
| System Requirements. ....        | xxvii |
| Technical Support. ....          | xxvii |

## Part I Addressing and Packet Flow Infrastructure

|   |    |
|---|----|
| 1 IPv4. ....                                  | 3  |
| Concepts. ....                                | 3  |
| Network Layers. ....                          | 3  |
| IPv4 Addressing. ....                         | 4  |
| Private IPv4 Addresses. ....                  | 7  |
| Automatic Private IP Addressing (APIPA). .... | 7  |
| Multicast Addresses. ....                     | 8  |
| Network Address Translation. ....             | 8  |
| Layer 2 and Layer 3 Addressing. ....          | 10 |
| Layer 4 Protocols: UDP and TCP. ....          | 12 |
| Planning and Design Considerations. ....      | 13 |
| Designing Your Internet Connection. ....      | 13 |
| Creating an IPv4 Addressing Scheme. ....      | 14 |
| Planning Host Addresses. ....                 | 15 |
| Using VPNs. ....                              | 16 |
| Planning Redundancy. ....                     | 17 |
| Using Multihomed Computers. ....              | 19 |

 **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/](http://www.microsoft.com/learning/booksurvey/)

|  |           |
|--|-----------|
| Deployment Steps . . . . .   | 20        |
| Manually Configuring IPv4 Clients . . . . .                              | 20        |
| Configuring Client Behavior When a DHCP Server Is Not Available. . . . . | 20        |
| Adding Routes to the Routing Table. . . . .                              | 21        |
| Ongoing Maintenance . . . . .  | 21        |
| Troubleshooting . . . . .  | 21        |
| ARP. . . . .   | 21        |
| Ipconfig . . . . .   | 22        |
| Netstat. . . . .   | 23        |
| PathPing . . . . .   | 24        |
| Performance Monitor . . . . .  | 25        |
| Ping . . . . .   | 26        |
| Task Manager . . . . .   | 27        |
| Windows Network Diagnostics . . . . .                                    | 27        |
| Chapter Summary . . . . .  | 28        |
| Additional Information . . . . .   | 28        |
| <b>2 IPv6 . . . . .</b>  | <b>29</b> |
| Concepts . . . . .   | 29        |
| Changes from IPv4 to IPv6. . . . .                                       | 30        |
| IPv6 Addressing . . . . .  | 31        |
| IPv6 Autoconfiguration. . . . .  | 36        |
| DHCPv6. . . . .  | 39        |
| Neighbor Discovery. . . . .  | 39        |
| IPv6 Security. . . . .   | 39        |
| IPv6 Transition Technologies . . . . .                                   | 39        |
| Planning and Design Considerations . . . . .                             | 46        |
| Migrating to IPv6. . . . .   | 47        |
| Acquiring IPv6 Addresses. . . . .  | 48        |
| Planning Network Infrastructure Upgrades. . . . .                        | 48        |
| Planning for IPv6 Transition Technologies. . . . .                       | 49        |
| Deployment Steps . . . . .   | 50        |
| How to Disable IPv6 . . . . .  | 50        |
| How to Manually Configure IPv6 . . . . .                                 | 51        |
| How to Configure IPv6 from a Script . . . . .                            | 52        |
| How to Enable ISATAP. . . . .  | 52        |



|   |           |
|---|-----------|
| How to Enable 6to4 .....                                  | 54        |
| How to Enable Teredo .....                                | 55        |
| How to Configure a Computer as an IPv6 Router .....       | 56        |
| Ongoing Maintenance .....                                 | 59        |
| Troubleshooting .....                                     | 60        |
| Netsh .....   | 60        |
| Ipconfig .....  | 61        |
| Nslookup .....  | 61        |
| Troubleshooting Teredo .....                              | 62        |
| Chapter Summary .....                                     | 63        |
| Additional Information .....                              | 63        |
| <b>3 Dynamic Host Configuration Protocol .....</b>        | <b>65</b> |
| Concepts .....  | 65        |
| The DHCP Address Assignment Process .....                 | 65        |
| DHCP Life Cycle .....                                     | 67        |
| Planning and Design Considerations .....                  | 67        |
| DHCP Servers .....  | 68        |
| DHCP Relay Agents .....                                   | 68        |
| DHCP Lease Durations .....                                | 69        |
| Designing Scopes .....                                    | 70        |
| Server Clustering for DHCP .....                          | 71        |
| Dynamic DNS .....   | 71        |
| Deployment Steps .....                                    | 71        |
| DHCP Servers .....  | 72        |
| DHCP Relay Agents .....                                   | 80        |
| DHCP Client Configuration .....                           | 81        |
| Ongoing Maintenance .....                                 | 82        |
| Monitoring DHCP Servers .....                             | 82        |
| Manually Backing Up and Restoring a DHCP Server .....     | 84        |
| Troubleshooting .....                                     | 85        |
| Troubleshooting DHCP Clients .....                        | 85        |
| Troubleshooting DHCP Servers .....                        | 85        |
| Using Audit Logging to Analyze DHCP Server Behavior ..... | 86        |
| Chapter Summary .....                                     | 86        |
| Additional Information .....                              | 87        |

|          |   |            |
|----------|---|------------|
| <b>4</b> | <b>Windows Firewall with Advanced Security</b>          | <b>89</b>  |
|          | Concepts  | 89         |
|          | Filtering Traffic by Using Windows Firewall             | 90         |
|          | Protecting Traffic by Using IPsec                       | 90         |
|          | Planning and Design Considerations                      | 95         |
|          | Planning Windows Firewall Policies                      | 95         |
|          | Protecting Communications with IPsec                    | 98         |
|          | Deployment Steps  | 105        |
|          | Firewall Settings with Group Policy                     | 105        |
|          | IPsec Connection Security Rules                         | 112        |
|          | Ongoing Maintenance                                     | 116        |
|          | Troubleshooting   | 117        |
|          | Windows Firewall Logging                                | 118        |
|          | Monitoring IPsec Security Associations                  | 120        |
|          | Using Network Monitor                                   | 121        |
|          | Chapter Summary   | 122        |
|          | Additional Information                                  | 122        |
| <b>5</b> | <b>Policy-Based Quality of Service</b>                  | <b>123</b> |
|          | Concepts  | 123        |
|          | The Causes of Network Performance Problems              | 123        |
|          | How QoS Can Help  | 125        |
|          | QoS for Outbound Traffic                                | 126        |
|          | QoS for Inbound Traffic                                 | 128        |
|          | QoS Implementation                                      | 128        |
|          | Planning and Design Considerations                      | 128        |
|          | Setting QoS Goals                                       | 129        |
|          | Planning DSCP Values                                    | 129        |
|          | Planning Traffic Throttling                             | 131        |
|          | Hardware and Software Requirements                      | 131        |
|          | Planning GPOs and QoS Policies                          | 132        |
|          | QoS Policies for Mobile Computers Running Windows Vista | 134        |
|          | Deployment Steps  | 134        |
|          | How to Configure QoS by Using Group Policy              | 134        |
|          | How to Configure System-Wide QoS Settings               | 138        |
|          | Ongoing Maintenance                                     | 140        |
|          | Removing QoS Policies                                   | 140        |

|                |   |            |
|----------------|---|------------|
|                | Editing QoS Policies .....                        | 141        |
|                | Monitoring QoS .....                              | 141        |
|                | Troubleshooting .....                             | 143        |
|                | Analyzing QoS Policies .....                      | 143        |
|                | Verifying DSCP Resilience .....                   | 146        |
|                | Isolating Network Performance Problems .....      | 147        |
|                | Chapter Summary .....                             | 148        |
|                | Additional Information .....                      | 148        |
| <b>6</b>       | <b>Scalable Networking .....</b>                  | <b>149</b> |
|                | Concepts .....                                    | 149        |
|                | TCP Chimney Offload .....                         | 150        |
|                | Receive-Side Scaling .....                        | 151        |
|                | NetDMA .....                                      | 153        |
|                | IPsec Offload .....                               | 154        |
|                | Planning and Design Considerations .....          | 155        |
|                | Evaluating Network Scalability Technologies ..... | 155        |
|                | Load Testing Servers .....                        | 156        |
|                | Monitoring Server Performance .....               | 157        |
|                | Deployment Steps .....                            | 159        |
|                | Configuring TCP Chimney Offload .....             | 160        |
|                | Configuring Receive-Side Scaling .....            | 160        |
|                | Configuring NetDMA .....                          | 161        |
|                | Configuring IPsec Offload .....                   | 161        |
|                | Ongoing Maintenance .....                         | 162        |
|                | Troubleshooting .....                             | 162        |
|                | Troubleshooting TCP Chimney Offload .....         | 162        |
|                | Troubleshooting IPsec Offload .....               | 163        |
|                | Chapter Summary .....                             | 164        |
|                | Additional Information .....                      | 165        |
| <b>Part II</b> | <b>Name Resolution Infrastructure</b>             |            |
| <b>7</b>       | <b>Domain Name System .....</b>                   | <b>169</b> |
|                | Concepts .....                                    | 169        |
|                | DNS Hierarchy .....                               | 169        |
|                | DNS Zones .....                                   | 170        |
|                | DNS Records .....                                 | 170        |

|  |            |
|--|------------|
| Dynamic DNS Updates . . . . .                          | 171        |
| DNS Name Resolution . . . . .                          | 172        |
| Planning and Design Considerations . . . . .           | 173        |
| DNS Zones . . . . .                                    | 173        |
| DNS Server Placement . . . . .                         | 174        |
| DNS Zone Replication . . . . .                         | 176        |
| DNS Security . . . . .                                 | 178        |
| The GlobalNames Zone . . . . .                         | 179        |
| Deployment Steps . . . . .                             | 180        |
| DNS Server Configuration . . . . .                     | 180        |
| DHCP Server Configuration . . . . .                    | 190        |
| DNS Client Configuration . . . . .                     | 192        |
| Configuring Redundant DNS Servers . . . . .            | 193        |
| Ongoing Maintenance . . . . .                          | 193        |
| Adding Resource Records . . . . .                      | 194        |
| Maintaining Zones . . . . .                            | 194        |
| Automated Monitoring . . . . .                         | 195        |
| Promoting a Secondary Zone to a Primary Zone . . . . . | 197        |
| Troubleshooting . . . . .                              | 198        |
| Event Logs . . . . .                                   | 198        |
| Using Nslookup . . . . .                               | 199        |
| Debug Logging at the Server . . . . .                  | 201        |
| Using DNSLint . . . . .                                | 202        |
| Using DCDiag . . . . .                                 | 203        |
| Using Network Monitor . . . . .                        | 205        |
| Chapter Summary . . . . .                              | 205        |
| Additional Information . . . . .                       | 206        |
| <b>8 Windows Internet Name Service . . . . .</b>       | <b>207</b> |
| Concepts . . . . .                                     | 207        |
| History . . . . .                                      | 207        |
| NetBIOS Names . . . . .                                | 208        |
| WINS Name Resolution . . . . .                         | 209        |
| WINS Client Registrations . . . . .                    | 210        |
| Planning and Design Considerations . . . . .           | 211        |
| WINS Server Placement . . . . .                        | 211        |
| WINS Replication . . . . .                             | 212        |

|   |     |
|---|-----|
| Deployment Steps .....                    | 213 |
| Configuring a WINS Server .....           | 213 |
| Configuring WINS Replication .....        | 214 |
| WINS Client Configuration .....           | 214 |
| Ongoing Maintenance .....                 | 217 |
| Backing Up the WINS Server Database ..... | 217 |
| Compacting the WINS Database .....        | 217 |
| Performing Consistency Checking .....     | 218 |
| Monitoring a WINS Server .....            | 219 |
| Adding a Static WINS Record .....         | 220 |
| Deleting a WINS Record .....              | 222 |
| Troubleshooting .....                     | 222 |
| Troubleshooting WINS Servers .....        | 222 |
| Troubleshooting WINS Clients .....        | 224 |
| Chapter Summary .....                     | 228 |
| Additional Information .....              | 228 |

## **Part III Network Access Infrastructure**

|  |            |
|--|------------|
| <b>9 Authentication Infrastructure .....</b>               | <b>231</b> |
| Concepts .....   | 231        |
| Active Directory Domain Services .....                     | 231        |
| Public Key Infrastructure .....                            | 235        |
| Group Policy .....   | 240        |
| RADIUS .....   | 243        |
| Planning and Design Considerations .....                   | 248        |
| Active Directory .....                                     | 248        |
| PKI .....  | 249        |
| Group Policy .....   | 251        |
| RADIUS .....   | 252        |
| Deployment Steps .....                                     | 260        |
| Deploying Active Directory .....                           | 260        |
| Deploying PKI .....  | 261        |
| Group Policy .....   | 269        |
| RADIUS Servers .....                                       | 270        |
| Using RADIUS Proxies for Cross-Forest Authentication ..... | 277        |
| Using RADIUS Proxies to Scale Authentications .....        | 284        |

|           |  |            |
|-----------|--|------------|
|           | Ongoing Maintenance .....                                  | 288        |
|           | Active Directory .....                                     | 289        |
|           | PKI .....  | 289        |
|           | Group Policy .....   | 289        |
|           | RADIUS .....   | 290        |
|           | Troubleshooting Tools .....                                | 291        |
|           | Active Directory .....                                     | 291        |
|           | PKI .....  | 292        |
|           | Group Policy .....   | 292        |
|           | RADIUS .....   | 292        |
|           | Chapter Summary .....                                      | 294        |
|           | Additional Information .....                               | 294        |
| <b>10</b> | <b>IEEE 802.11 Wireless Networks .....</b>                 | <b>297</b> |
|           | Concepts .....   | 297        |
|           | Support for IEEE 802.11 Standards .....                    | 298        |
|           | Wireless Security .....                                    | 300        |
|           | Components of 802.11 Wireless Networks .....               | 304        |
|           | Planning and Design Considerations .....                   | 305        |
|           | Wireless Security Technologies .....                       | 305        |
|           | Wireless Authentication Modes .....                        | 308        |
|           | Intranet Infrastructure .....                              | 309        |
|           | Wireless AP Placement .....                                | 311        |
|           | Authentication Infrastructure .....                        | 316        |
|           | Wireless Clients .....                                     | 317        |
|           | PKI .....  | 328        |
|           | 802.1X Enforcement with NAP .....                          | 332        |
|           | Deploying Protected Wireless Access .....                  | 332        |
|           | Deploying Certificates .....                               | 332        |
|           | Configuring Active Directory for Accounts and Groups ..... | 334        |
|           | Configuring NPS Servers .....                              | 335        |
|           | Deploying Wireless APs .....                               | 336        |
|           | Configuring Wireless Clients .....                         | 339        |
|           | Ongoing Maintenance .....                                  | 345        |
|           | Managing User and Computer Accounts .....                  | 345        |
|           | Managing Wireless APs .....                                | 345        |
|           | Updating Wireless XML Profiles .....                       | 346        |



|  |            |
|--|------------|
| Troubleshooting .....  | 346        |
| Wireless Troubleshooting Tools in Windows .....              | 347        |
| Troubleshooting the Windows Wireless Client .....            | 355        |
| Troubleshooting the Wireless AP .....                        | 356        |
| Troubleshooting the Authentication Infrastructure .....      | 361        |
| Chapter Summary .....  | 367        |
| Additional Information .....                                 | 367        |
| <b>11 IEEE 802.1X–Authenticated Wired Networks.....</b>      | <b>369</b> |
| Concepts .....   | 369        |
| Components of Wired Networks With 802.1X Authentication..... | 369        |
| Planning and Design Considerations .....                     | 370        |
| Wired Authentication Methods .....                           | 371        |
| Wired Authentication Modes .....                             | 374        |
| Authentication Infrastructure .....                          | 375        |
| Wired Clients .....  | 376        |
| PKI .....  | 382        |
| 802.1X Enforcement with NAP .....                            | 385        |
| Deploying 802.1X-Authenticated Wired Access .....            | 386        |
| Deploying Certificates .....                                 | 386        |
| Configuring Active Directory for Accounts and Groups .....   | 388        |
| Configuring NPS Servers .....                                | 388        |
| Configuring 802.1X-Capable Switches .....                    | 390        |
| Configuring Wired Clients .....                              | 392        |
| Ongoing Maintenance .....                                    | 395        |
| Managing User and Computer Accounts .....                    | 395        |
| Managing 802.1X-Capable Switches .....                       | 396        |
| Updating Wired XML Profiles .....                            | 396        |
| Troubleshooting .....  | 397        |
| Wired Troubleshooting Tools in Windows .....                 | 397        |
| Troubleshooting the Windows Wired Client .....               | 402        |
| Troubleshooting the 802.1X-Capable Switch .....              | 403        |
| Troubleshooting the Authentication Infrastructure .....      | 407        |
| Chapter Summary .....  | 413        |
| Additional Information .....                                 | 413        |

|           |   |            |
|-----------|---|------------|
| <b>12</b> | <b>Remote Access VPN Connections</b>                      | <b>417</b> |
|           | Concepts  | 417        |
|           | Components of Windows Remote Access VPNs                  | 420        |
|           | Planning and Design Considerations                        | 421        |
|           | VPN Protocols   | 421        |
|           | Authentication Methods                                    | 426        |
|           | VPN Servers   | 428        |
|           | Internet Infrastructure                                   | 431        |
|           | Intranet Infrastructure                                   | 433        |
|           | Concurrent Intranet and Internet Access for VPN Clients   | 437        |
|           | Authentication Infrastructure                             | 439        |
|           | VPN Clients   | 441        |
|           | PKI   | 445        |
|           | VPN Enforcement with NAP                                  | 448        |
|           | Additional Security Considerations                        | 449        |
|           | Strong Link Encryption                                    | 449        |
|           | VPN Traffic Packet Filtering on the VPN Server            | 450        |
|           | Firewall Packet Filtering for VPN Traffic                 | 450        |
|           | Multi-Use VPN Servers                                     | 460        |
|           | Blocking Traffic Routed from VPN Clients                  | 462        |
|           | Concurrent Access   | 463        |
|           | Unused VPN Protocols                                      | 463        |
|           | Deploying VPN-Based Remote Access                         | 463        |
|           | Deploying Certificates                                    | 464        |
|           | Configuring Internet Infrastructure                       | 467        |
|           | Configuring Active Directory for User Accounts and Groups | 468        |
|           | Configuring RADIUS Servers                                | 469        |
|           | Deploying VPN Servers                                     | 470        |
|           | Configuring Intranet Network Infrastructure               | 474        |
|           | Deploying VPN Clients                                     | 477        |
|           | Ongoing Maintenance                                       | 482        |
|           | Managing User Accounts                                    | 482        |
|           | Managing VPN Servers                                      | 483        |
|           | Updating CM Profiles                                      | 485        |
|           | Troubleshooting   | 485        |
|           | Troubleshooting Tools                                     | 485        |
|           | Troubleshooting Remote Access VPNs                        | 489        |

|   |            |
|---|------------|
| Chapter Summary . . . . .   | 496        |
| Additional Information . . . . .                                    | 496        |
| <b>13 Site-to-Site VPN Connections . . . . .</b>                    | <b>499</b> |
| Concepts . . . . .  | 499        |
| Demand-Dial Routing Overview . . . . .                              | 500        |
| Components of Windows Site-to-Site VPNs . . . . .                   | 505        |
| Planning and Design Considerations . . . . .                        | 506        |
| VPN Protocols . . . . .   | 506        |
| Authentication Methods . . . . .                                    | 510        |
| VPN Routers . . . . .   | 511        |
| Internet Infrastructure . . . . .                                   | 515        |
| Site Network Infrastructure . . . . .                               | 517        |
| Authentication Infrastructure . . . . .                             | 520        |
| PKI . . . . .   | 522        |
| Deploying Site-to-Site VPN Connections . . . . .                    | 525        |
| Deploying Certificates . . . . .                                    | 525        |
| Configuring Internet Infrastructure . . . . .                       | 529        |
| Configuring Active Directory for User Accounts and Groups . . . . . | 530        |
| Configuring RADIUS Servers . . . . .                                | 530        |
| Deploying the Answering Routers . . . . .                           | 532        |
| Deploying the Calling Routers . . . . .                             | 538        |
| Configuring Site Network Infrastructure . . . . .                   | 544        |
| Configuring Intersite Network Infrastructure . . . . .              | 546        |
| Ongoing Maintenance . . . . .                                       | 548        |
| Managing User Accounts . . . . .                                    | 549        |
| Managing VPN Routers . . . . .                                      | 549        |
| Troubleshooting . . . . .   | 551        |
| Troubleshooting Tools . . . . .                                     | 551        |
| Troubleshooting Site-to-Site VPN Connections . . . . .              | 551        |
| Chapter Summary . . . . .   | 561        |
| Additional Information . . . . .                                    | 561        |

## **Part IV Network Access Protection Infrastructure**

|  |            |
|--|------------|
| <b>14 Network Access Protection Overview . . . . .</b>   | <b>565</b> |
| The Need for Network Access Protection . . . . .         | 565        |
| Malware and Its Impact on Enterprise Computing . . . . . | 565        |
| Preventing Malware on Enterprise Networks . . . . .      | 567        |

|  |            |
|--|------------|
| The Role of NAP . . . . .  | 571        |
| Business Benefits of NAP . . . . .   | 574        |
| Components of NAP . . . . .  | 575        |
| System Health Agents and System Health Validators . . . . .                  | 577        |
| Enforcement Clients and Servers . . . . .                                    | 577        |
| NPS . . . . .  | 578        |
| Enforcement Methods . . . . .  | 579        |
| IPsec Enforcement . . . . .  | 579        |
| 802.1X Enforcement . . . . .   | 580        |
| VPN Enforcement . . . . .  | 580        |
| DHCP Enforcement . . . . .   | 580        |
| How NAP Works . . . . .  | 581        |
| How IPsec Enforcement Works . . . . .  | 582        |
| How 802.1X Enforcement Works . . . . .                                       | 583        |
| How VPN Enforcement Works . . . . .  | 584        |
| How DHCP Enforcement Works . . . . .   | 585        |
| Chapter Summary . . . . .  | 586        |
| Additional Information . . . . .   | 587        |
| <b>15     Preparing for Network Access Protection . . . . .</b>              | <b>589</b> |
| Evaluation of Your Current Network Infrastructure . . . . .                  | 589        |
| Intranet Computers . . . . .   | 589        |
| Networking Support Infrastructure . . . . .                                  | 592        |
| NAP Health Policy Servers . . . . .  | 593        |
| Planning and Design Considerations . . . . .                                 | 593        |
| Deployment Steps . . . . .   | 596        |
| Ongoing Maintenance . . . . .  | 596        |
| Health Requirement Policy Configuration . . . . .                            | 597        |
| Components of a Health Requirement Policy . . . . .                          | 597        |
| How NAP Health Evaluation Works . . . . .                                    | 605        |
| Planning and Design Considerations for Health Requirement Policies . . . . . | 609        |
| Remediation Servers . . . . .  | 611        |
| Remediation Servers and NAP Enforcement Methods . . . . .                    | 612        |
| Planning and Design Considerations for Remediation Servers . . . . .         | 613        |
| Chapter Summary . . . . .  | 614        |
| Additional Information . . . . .   | 614        |

|           |  |            |
|-----------|--|------------|
| <b>16</b> | <b>IPsec Enforcement</b>                                   | <b>617</b> |
|           | Understanding IPsec Enforcement                            | 617        |
|           | IPsec Enforcement Logical Networks                         | 618        |
|           | Communication Initiation Processes with IPsec Enforcement  | 619        |
|           | Connection Security Rules for IPsec Enforcement            | 623        |
|           | Planning and Design Considerations                         | 625        |
|           | Active Directory   | 625        |
|           | PKI  | 626        |
|           | HRAs   | 631        |
|           | IPsec Policies   | 638        |
|           | NAP Clients  | 639        |
|           | Deploying IPsec Enforcement                                | 640        |
|           | Configuring Active Directory                               | 641        |
|           | Configuring PKI  | 641        |
|           | Configuring HRAs   | 644        |
|           | Configuring NAP Health Policy Servers                      | 651        |
|           | Configuring Remediation Servers on the Boundary Network    | 654        |
|           | Configuring NAP Clients                                    | 655        |
|           | IPsec Enforcement Deployment Checkpoint for Reporting Mode | 658        |
|           | Configuring and Applying IPsec Policies                    | 659        |
|           | Ongoing Maintenance  | 665        |
|           | Adding a NAP Client  | 665        |
|           | Adding a New SHA and SHV                                   | 666        |
|           | Managing NAP CAs   | 666        |
|           | Managing HRAs  | 667        |
|           | Troubleshooting  | 669        |
|           | Troubleshooting Tools                                      | 669        |
|           | Troubleshooting IPsec Enforcement                          | 672        |
|           | Chapter Summary  | 678        |
|           | Additional Information                                     | 678        |
| <b>17</b> | <b>802.1X Enforcement</b>                                  | <b>681</b> |
|           | Overview of 802.1X Enforcement                             | 681        |
|           | Using an ACL   | 684        |
|           | Using a VLAN   | 685        |

|   |            |
|---|------------|
| Planning and Design Considerations . . . . .  | 686        |
| Security Group for NAP Exemptions. . . . .  | 686        |
| 802.1X Authentication Methods . . . . .   | 686        |
| Type of 802.1X Enforcement . . . . .  | 687        |
| 802.1X Access Points . . . . .  | 687        |
| NAP Clients. . . . .  | 688        |
| Deploying 802.1X Enforcement. . . . .   | 690        |
| Configuring Active Directory. . . . .   | 690        |
| Configuring a PEAP-Based Authentication Method . . . . .  | 690        |
| Configuring 802.1X Access Points. . . . .   | 691        |
| Configuring Remediation Servers on the Restricted Network . . . . .                               | 693        |
| Configuring NAP Health Policy Servers . . . . .   | 693        |
| Configuring NAP Clients. . . . .  | 701        |
| 802.1X Enforcement Deployment Checkpoint for Reporting Mode. . . . .                              | 704        |
| Testing Restricted Access . . . . .   | 704        |
| Configuring the Network Policy for Noncompliant NAP Clients<br>for Deferred Enforcement . . . . . | 706        |
| Configuring Network Policy for Enforcement Mode . . . . .   | 707        |
| Ongoing Maintenance . . . . .   | 708        |
| Adding a NAP Client . . . . .   | 708        |
| Adding a New SHA and SHV . . . . .  | 708        |
| Managing 802.1X Access Points . . . . .   | 709        |
| Troubleshooting . . . . .   | 709        |
| Troubleshooting Tools. . . . .  | 709        |
| Troubleshooting 802.1X Enforcement . . . . .  | 711        |
| Chapter Summary . . . . .   | 714        |
| Additional Information . . . . .  | 715        |
| <b>18 VPN Enforcement . . . . .</b>   | <b>717</b> |
| Understanding VPN Enforcement. . . . .  | 717        |
| Planning and Design Considerations . . . . .  | 720        |
| Use of Network Access Quarantine Control . . . . .  | 720        |
| Security Group for NAP Exemptions. . . . .  | 721        |
| Types of Packet Filtering . . . . .   | 722        |
| VPN Authentication Methods . . . . .  | 723        |
| VPN Servers . . . . .   | 723        |
| NAP Clients. . . . .  | 724        |



|   |            |
|---|------------|
| Deploying VPN Enforcement .....   | 726        |
| Configuring Active Directory .....  | 726        |
| Configuring VPN Servers .....   | 727        |
| Configuring a PEAP-Based Authentication Method .....                                  | 727        |
| Configuring Remediation Servers .....   | 727        |
| Configuring NAP Health Policy Servers .....   | 728        |
| Configuring NAP Clients .....   | 735        |
| VPN Enforcement Deployment Checkpoint for Reporting Mode .....                        | 737        |
| Testing Restricted Access .....   | 737        |
| Configuring Deferred Enforcement .....  | 739        |
| Configuring Network Policy for Enforcement Mode .....                                 | 739        |
| Ongoing Maintenance .....   | 741        |
| Adding a NAP Client .....   | 741        |
| Adding a New SHA and SHV .....  | 741        |
| Troubleshooting .....   | 741        |
| Troubleshooting Tools .....   | 742        |
| Troubleshooting VPN Enforcement .....   | 744        |
| Chapter Summary .....   | 747        |
| Additional Information .....  | 747        |
| <b>19 DHCP Enforcement .....</b>  | <b>749</b> |
| Understanding DHCP Enforcement .....  | 749        |
| Planning and Design Considerations .....  | 752        |
| Security Group for NAP Exemptions .....   | 752        |
| DHCP Servers .....  | 753        |
| NAP Health Policy Servers .....   | 753        |
| Health Requirement Policies for Specific DHCP Scopes .....                            | 754        |
| DHCP Options for NAP Clients .....  | 754        |
| DHCP Enforcement Behavior When the NAP Health Policy Server<br>Is Not Reachable ..... | 754        |
| NAP Clients .....   | 754        |
| Deploying DHCP Enforcement .....  | 756        |
| Configuring Remediation Servers .....   | 756        |
| Configuring NAP Health Policy Servers .....   | 757        |
| Configuring NAP Clients .....   | 761        |
| Configuring DHCP Servers .....  | 763        |
| DHCP Enforcement Deployment Checkpoint for Reporting Mode .....                       | 767        |

- Testing Restricted Access ..... 767
- Configuring Deferred Enforcement ..... 769
- Configuring Network Policy for Enforcement Mode..... 769
- Ongoing Maintenance ..... 771
  - Adding a NAP Client..... 771
  - Adding a New SHA and SHV..... 771
- Troubleshooting..... 772
  - Troubleshooting Tools..... 772
  - Troubleshooting DHCP Enforcement..... 774
- Chapter Summary ..... 777
- Additional Information..... 777
- Glossary .....779
- Index.....793

 **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/](http://www.microsoft.com/learning/booksurvey/)

## Chapter 5

# Policy-Based Quality of Service

This chapter provides information about how to design, deploy, maintain, and troubleshoot Quality of Service (QoS) in Windows Server 2008.

This chapter assumes the following:

- That you understand the role of Active Directory and Group Policy for managing Microsoft Windows computers.
- That you have a solid understanding of managing routers in your network infrastructure.

## Concepts

As more people and organizations begin to use real-time networking services, such as Voice over Internet Protocol (VoIP), multimedia streaming, and video conferencing, the impact of network performance problems becomes more significant. Ten years ago, a network performance problem would just cause a Web page to open slowly. Today, it can make your phone service unusable, interrupt critical video conference meetings, and prevent financial transactions from being completed.

Simply adding bandwidth won't solve most network performance problems. On many networks, a single large file transfer can monopolize the entire network connection, negatively impacting any real-time communications while the transfer occurs. Quality of Service (QoS) works with your servers, clients, and network infrastructure to prioritize network traffic. With QoS, phone calls and other real-time communications can be given priority over file transfers, e-mail, Web browsing, and other lower-priority communications.

## The Causes of Network Performance Problems

Some of the network conditions that can cause performance problems include latency, jitter, out-of-order delivery, and dropped packets. The sections that follow describe these conditions in more detail.

### Latency

*Latency* is the delay it takes for a packet to reach its destination, typically measured in milliseconds (ms). When planning for QoS, round-trip latency (the time it takes for a packet to be sent to a remote host and for a response to be returned) is the most important metric because it has the most significant impact on real-time two-way communications such as VoIP.

Latency comes from several sources:

- **Forwarding delay** When a router processes a packet and moves it from one network to the next, there is normally an insignificant delay of one or two milliseconds. However, if the router receives traffic faster than it can forward it to the destination network—for example, if it receives communications at 5 megabits per second (Mbps) and must forward it to an interface that supports only 1.54 Mbps—the router must store the packet a queue, causing additional latency.
- **Propagation delay** Communications take time to travel a distance. For most copper or fiber networks, the speed is about 2/3 the speed of light—around 125,000 miles per second. On local area networks, the propagation delay is negligible. However, for a transmission to travel halfway around the world would take close to 100 ms (because networks are never a straight line). Communications that pass through a satellite link incur a latency of about 500 ms, making them unusable for VoIP. Virtual private networks (VPNs) often cause network communications to take an extremely inefficient path between the source and destination, multiplying the propagation delay.
- **Host processing delay** If an incoming packet is VoIP or streaming media, the operating system or application will hold the packet in a jitter buffer to minimize the impact of jitter (discussed in the next section) and out-of-order delivery. Jitter buffers vary, but packets are typically held about 20 to 200 ms, adding latency. Once the packets have been held in a jitter buffer for a sufficient time, the application must process the data contained within the packets, which can incur an additional processing delay depending on the speed of the computer and the percentage of processing time currently dedicated to the network application.

## Jitter

*Jitter* is change in latency. For example, a one-way video stream that begins with 10 ms of latency might suddenly have 100 ms of latency if network conditions change. Software uses jitter buffers, as discussed in the previous section, to reduce the impact of jitter. However, the more jitter you have, the longer data must be held in the jitter buffer—increasing latency for all communications.

## Out-of-Order Delivery

In IP networks, two consecutive packets can take different routes between the source and the destination. This can cause packets to arrive out of order. Transmission Control Protocol (TCP) handles this automatically by waiting for out-of-order packets and reassembling them. For file transfers and most other communications, out-of-order packets do not cause a problem. However, with real-time communications, such as VoIP or streaming media (which typically use user datagram protocol or UDP), a packet that is received out of order and arrives after the jitter buffer has expired is useless and will be discarded by the client computer.

## Dropped Packets

Routers drop packets only when the router's queue is full and no more packets can be stored. When TCP communications are dropped, this can worsen the network congestion problems, because dropped packets must be retransmitted.

## How QoS Can Help

QoS can reduce the impact of network problems on high-priority traffic in several ways:

- **Reducing latency** By default, most routers forward traffic on a first-in, first-out (FIFO) basis. With QoS, routers can use a Differentiated Services Code Point (DSCP) value to forward high-priority traffic before low-priority traffic, even if the high-priority traffic arrived last. This increases the latency for the low-priority traffic but decreases it for the high-priority traffic. Adding bandwidth allows the router to empty the queue faster, but it cannot eliminate queuing. Computers that are transmitting traffic also queue packets sent by different applications and can use QoS priorities to transmit high-priority packets first.



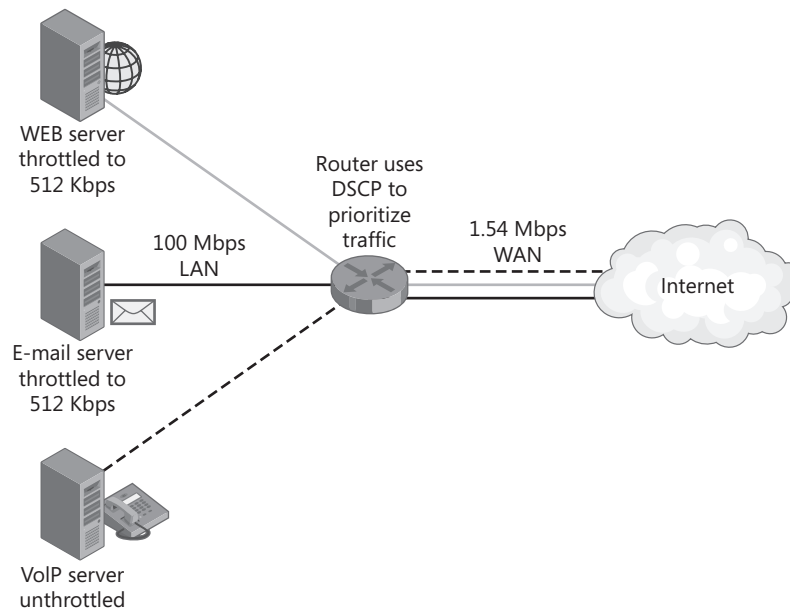
**Note** Even with QoS, bulk communications can still cause a slight forwarding delay at the router. For example, consider a user carrying on a VoIP conversation while trying to use HTTP to upload a large file across a 128 kilobits per second (Kbps) DSL link. After the router forwards all VoIP traffic, it will begin to transfer an HTTP packet, which will often be 1,500 bytes. If a VoIP packet arrives after the router begins forwarding the HTTP packet, it must finish sending the packet before it can forward the new VoIP packet, which would take about 100 ms in this example. That 100 ms forwarding delay could put the total latency above the acceptable latency for VoIP of about 150 ms. To minimize the impact of the latency incurred when transmitting a single large packet, some routers can be configured to fragment low-priority packets into smaller packets.

- **Reducing jitter** By using the same techniques used to reduce latency, QoS can also reduce jitter. Lower amounts of jitter allow you to decrease the time data is held in the jitter buffer, further decreasing total latency. Software often automatically adjusts the jitter buffer as network conditions change. Depending on the layout of your network, you might be able to reduce jitter by routing critical traffic across dedicated, low-latency links.
- **Reducing out-of-order delivery** If you mark communications with DSCP values, you can configure routers to use the same route for all communications of that type, minimizing both jitter and out-of-order delivery.
- **Reducing high-priority dropped packets** If you mark communications with DSCP values, you can configure many routers to drop lower-priority traffic before higher-priority traffic.

## QoS for Outbound Traffic

On computers running the Windows Vista and Windows Server 2008 operating systems, you can implement QoS by using two techniques: DSCP and traffic throttling. Ideally, you would combine both to provide the highest level of service.

Figure 5-1 demonstrates how DSCP and traffic throttling can work together. In this example, a high-speed 100-Mbps LAN is connected to a relatively low-speed 1.54-Mbps Internet connection. Server administrators have used traffic throttling to limit to 512 Kbps each the traffic that is bound for the Internet from the Web server and the e-mail server. Therefore, even if the Web site is busy and many users are downloading their e-mail simultaneously, there will still be 512 Kbps of bandwidth remaining for the VoIP server to use. To reduce latency that might occur when the router queues traffic, traffic from each server is labeled with a DSCP number. The router uses these DSCP numbers to prioritize the traffic so that latency-sensitive VoIP traffic always leaves the queue first.



**Figure 5-1** DSCP and traffic throttling working together

The sections that follow describe DSCP and traffic throttling in more detail.

### DSCP

RFC 2474 defines the Differentiated Services Code Point (DSCP), which adds a value to the IP header of outgoing datagrams that routers can use to prioritize traffic. For example, you could configure Windows Server 2008 to add the DSCP value 10 to FTP traffic and the DSCP value 46 to streaming media traffic. When a router forwards this traffic, it can place traffic with a DSCP value of 10 in the low-priority queue and traffic with a DSCP value of 46 in the high-priority queue.



DSCP is specified by a number from 0 to 63 in the IP header, where 0 indicates that no DSCP value has been provided. In IPv4, DSCP uses the type of service (TOS, defined in RFC 791) octet in the header, as shown in Figure 5-2. In IPv6, DSCP uses the traffic class octet in the header. Though DSCP uses a full eight-bit octet in both IPv4 and IPv6, DSCP uses only the first six bits. The remaining two bits are Explicit Congestion Notification (ECN) bits.



**More Info** For more information about ECN, visit <http://www.microsoft.com/technet/community/columns/cableguy/cg1006.mspx>.

|                        |               |               |                        |        |                      |
|------------------------|---------------|---------------|------------------------|--------|----------------------|
| 4                      | Header length | DSCP (6 bits) | 0                      | Length |                      |
| 16-bit identification  |               |               |                        | Flags  | Fragmentation offset |
| Time-to-live           |               | Protocol      | 16-bit header checksum |        |                      |
| Source IP address      |               |               |                        |        |                      |
| Destination IP address |               |               |                        |        |                      |
| Options                |               |               |                        |        |                      |
| Data                   |               |               |                        |        |                      |

Figure 5-2 DSCP in the IPv4 header

### Traffic Throttling

Whereas DSCP allows your routing infrastructure to prioritize traffic, Windows Vista and Windows Server 2008 can also use throttling to limit the amount of bandwidth used by specific applications and protocols. With throttling, a QoS policy limits outgoing network traffic that matches the specified criteria to a given rate.

For example, if you have a 10-Mbps LAN, you could prevent your e-mail servers from consuming all available network bandwidth by placing them into their own organizational unit (OU) and linking a Group Policy Object (GPO) to that OU. In the GPO, configure a QoS policy that throttles e-mail communications (which you can configure using TCP port numbers) to 4 Mbps.

## QoS for Inbound Traffic

DSCP values give you some control over how your network infrastructure handles outgoing traffic, whereas traffic throttling directly reduces the outgoing bandwidth used by different applications. Although you cannot directly control the rate of inbound traffic, Windows can adjust the TCP receive window to slow down the rate of incoming traffic.

By default, Windows will optimize the TCP receive window for maximum throughput. However, you can use system-wide QoS settings in computers running Windows Vista or Windows Server 2008 to limit incoming traffic by specifying the maximum size of the TCP receive window. The TCP receive window is the amount of data that a receiver allows a sender to transmit before being required to wait for an acknowledgement. A larger maximum window size means that the sender can send more data at a time, increasing network utilization and throughput. By limiting the maximum size of the TCP receive window, a receiver can indirectly control the incoming throughput for a TCP connection.

For more information about configuring the maximum TCP receive window size, read “How to Configure System-Wide QoS Settings” later in this chapter.

## QoS Implementation

On Windows Vista and Windows Server 2008, QoS is implemented in the Pacer.sys NDIS 6.0 lightweight filter driver, located in %SystemRoot%\System32\Drivers. Pacer.sys controls QoS packet scheduling for both policy-based QoS and applications that use the Generic QoS (GQoS), Traffic Control (TC), and qWAVE (Quality Windows Audio Video Experience) APIs. Pacer.sys is used only when the QoS Packet Scheduler component of a network connection or adapter is enabled (which it is by default). Pacer.sys replaces Psched.sys, which is used in the Windows Server 2003 and Windows XP operating systems.



**Note** Quality Windows Audio Video Experience (qWAVE) is a QoS API designed to improve the performance of audio and video streaming across home networks. On computers running Windows Server 2008, qWAVE provides only rate-of-flow and prioritization services. Because it is designed primarily for home use, it is not discussed further in this chapter.

## Planning and Design Considerations

QoS can be deployed to Windows computers with just a few clicks, but without proper planning, you might not realize any performance benefits. This section describes how to identify goals for your QoS implementation, specify DSCP values, plan traffic throttling, verify hardware and software requirements, and design QoS policies.

## Setting QoS Goals

When planning QoS policies for your network, determine the applications that require QoS, and set latency and bandwidth goals for each application as follows:

- **Latency goals** Minimizing latency is critical for VoIP and video conferencing. Although lower latency is always better, the maximum end-to-end delay that humans can tolerate is 150–200 ms. Above 200 ms, people will be irritated by the delay. Conversations with excess latency become awkward, and people frequently begin speaking at the same time.
- **Bandwidth goals** These goals will depend on your specific applications but might include a minimum number of simultaneous media streams or maximum time to transfer a network backup of a specific size.

After you have implemented QoS, you can use these goals to determine whether the implementation was a success or additional changes are required.

## Planning DSCP Values

Table 5-1 lists standard DSCP values for different applications in order from highest to lowest priority. Rows shown in bold are the most commonly used.

**Table 5-1 DSCP Interoperability Values**

| Purpose                      | Common Uses  | DSCP Value |
|------------------------------|--|------------|
| IP routing                   | Router-to-router communications.   | 48         |
| <b>VoIP</b>                  | <b>VoIP traffic, including signaling and control traffic.</b>  | <b>46</b>  |
| <b>Interactive video</b>     | <b>Two-way video conferencing.</b>   | <b>34</b>  |
| Streaming video              | One-way video streaming. Alternatively, you can classify streaming video as mission critical data.   | 32         |
| <b>Mission-critical data</b> | <b>Database queries, line-of-business communications, video streaming.</b>   | <b>26</b>  |
| Transactional data           | Database queries and transactions. Alternatively, you can classify transactional data as mission-critical data.  | 28         |
| Call signaling               | VoIP control traffic, which can also be classified as VoIP.  | 24         |
| Network management           | Network management protocols such as simple network management protocol (SNMP). Alternatively, you can classify network management traffic as mission-critical data. | 16         |
| <b>Best effort</b>           | <b>All other traffic, including e-mail and Web browsing.</b>   | <b>0</b>   |
| <b>Bulk data</b>             | <b>Backups, non-business applications, file transfers.</b>   | <b>10</b>  |
| Scavenger                    | Low-priority traffic. Alternatively, you can classify all low-priority traffic as bulk data.   | 8          |



**Note** It's common practice to remove or rewrite DSCP values from traffic originating from remote networks such as the Internet because the DSCP value might have a different meaning, or it might even be part of a denial-of-service (DoS) attack. Therefore, you can't be sure that DSCP values will be retained or respected when sending traffic to networks that you don't manage.

## Wireless Multimedia and DSCP Values

Wireless Multimedia (WMM) includes four access categories for prioritizing traffic on 802.11 wireless networks. WMM uses DSCP values to set priority, so you can automatically take advantage of WMM by specifying DSCP values. Table 5-2 shows how DSCP values correspond to WMM access categories.

**Table 5-2 DSCP Values and WMM Access Categories**

| DSCP Value | WMM Access Category |
|------------|---------------------|
| 48–63      | Voice (VO)          |
| 32–47      | Video (VI)          |
| 24–31, 0–7 | Best effort (BE)    |
| 8–23       | Background (BK)     |

You do not need to specify a separate DSCP number for every protocol on your network. Instead, you should specify DSCP numbers for a different traffic priority types. For example, a typical DSCP strategy includes the following five queues:

- **Control traffic** Communications transmitted between routers. Typically these communications require minimal bandwidth, but they should be assigned a high priority because quick transmission can reduce downtime in the event of a hardware failure. You should also use this priority for VoIP control traffic. Use DSCP values of 26 for control traffic.
- **Latency-sensitive traffic** Traffic, such as VoIP, that must be delivered as quickly as possible. Typically, you should assign this a DSCP value of 46, known as Expedited Forwarding (EF).
- **Business critical traffic, also known as Better than Best Effort (BBE)** Communications that should receive priority treatment, such as customer service database queries from a line-of-business (LOB) application or streaming video, but that are not highly sensitive to latency. Use a DSCP value of 34.
- **Best effort traffic** Standard traffic, including any traffic not marked with a DSCP number, that should be handled after either of the preceding two queues. This traffic should have a DSCP value of 0, which is the default if no DSCP value has been specified.
- **Scavenger traffic** Low-priority traffic, such as backups, downloading of updates, non-critical file synchronization, and non-work-related traffic that employees might generate. Use a DSCP value of 10 or 8.



**Note** If you mark traffic from too many applications as high-priority, the high-priority queue on routers can grow long enough to add significant latency. This defeats the purpose of QoS. Therefore, you should reserve the highest priority DSCP marking for real-time communications such as VoIP.

Many networks use an even simpler structure, with only two priorities: one for latency-sensitive traffic and another for best effort traffic. However, if you have third-party tools that can use DSCP values to report on network performance for different types of traffic, it is advantageous to define a larger number of DSCP values even if your network infrastructure isn't configured to handle each DSCP value uniquely.

## Planning Traffic Throttling

Ideally, networks should always be fully utilized—even if the traffic is considered low-priority. Prioritizing traffic by using DSCP values supports this philosophy by allowing lower-priority traffic to use all available bandwidth when no higher-priority traffic requires the bandwidth.

If sections of your network do not support prioritizing traffic with DSCP values, you can use traffic throttling to limit the amount of traffic being sent from your computers running Windows Vista and Windows Server 2008. Do not attempt to use traffic throttling to limit the bandwidth of every application or protocol; instead, use traffic throttling to limit only traffic from low-priority applications such as network backups or the downloading of large updates.

Remember, traffic throttling limits traffic on individual computers only. Traffic throttling cannot limit the aggregate bandwidth used by multiple computers. For example, if you have five FTP servers and you want to ensure that they never use more than half of your 500-kilobyte-per-second (KBps) link, you must configure the QoS policy to throttle traffic at 50 KBps for each of the five computers, which would total 250 KBps if all five servers were sending traffic at their throttled maximum.

## Hardware and Software Requirements

The sections that follow describe operating system and application support for QoS policies, backwards compatibility for QoS APIs used in earlier versions of Windows, and network infrastructure requirements for QoS support.

### Support for QoS Policies

You can apply QoS policies only to Windows Vista and Windows Server 2008. Earlier versions of Windows support QoS APIs that individual applications can use to set DSCP values and traffic throttling; however, they do not support the application of QoS policies through the use of GPOs. To implement QoS policies in a domain, your domain controller can be running Microsoft Windows 2000 Server, Windows Server 2003, or Windows Server 2008.

Applications running on Windows Vista and Windows Server 2008 do not need to support QoS to have their network traffic prioritized. You can use QoS policies to apply QoS to any network traffic, including network traffic generated by core operating system services (such as the Server service).

## Backward Compatibility for QoS APIs

Windows 2000, Windows XP, and Windows Server 2003 provide QoS capabilities by using the Generic QoS (GQoS), IP Type Of Service (TOS), and Traffic Control (TC) application programming interfaces (APIs). Developers needed to create an application specifically to take advantage of one of these APIs, and most developers did not add this capability to their applications. Therefore, most applications did not support QoS. If you have an application that uses one of these APIs, the application will still work on Windows Vista and Windows Server 2008.



**More Info** For more information about these APIs, see “The MS QoS Components” at <http://www.microsoft.com/technet/prodtechnol/windows2000serv/maintain/featusability/qoscomp.mspx>.



**Note** GQoS, IP TOS, and TC have been deprecated. Although they are still supported in Windows Vista, future versions of Windows might not support them. Therefore, if you have applications that use GQoS, TOS, or TC, you should encourage the developers to use the new QoS2 API instead. Developers can find QoS2 in the QoS2.h header file.

## Network Infrastructure Requirements

To fully support QoS policies, your network infrastructure must support the use of multiple queues to prioritize traffic based on DSCP value (as defined in RFC 2474). Traffic throttling does not have any network infrastructure requirements.

## Planning GPOs and QoS Policies

With Windows Vista and Windows Server 2008, you can use local or Active Directory GPOs to configure QoS for any application on your network. Combined with the flexibility of Group Policy, this allows you to:

- **Configure QoS policy for organizational units** Most organizations use Active Directory OUs to organize users and computers. By linking a GPO to an OU and setting the QoS policy in that GPO, you can apply different QoS policies to different computers in your organization. For example, if you have separate OUs for e-mail and Web servers, you could link different GPOs to each OU to configure e-mail traffic (such as Microsoft Office Outlook Web Access or OWA) as a higher priority than standard Web traffic.



- **Configure QoS policy based on user group membership** By configuring GPOs with access control lists (ACLs) that restrict access based on group membership, you can use User Configuration QoS policies to give specific groups a higher priority than other groups. For example, you might assign a higher priority to traffic generated by your customer service team.
- **Configure QoS policy for sites** If different sites have the network infrastructure configured differently (for example, if they use different DSCP values), you can link GPOs to Active Directory sites to apply settings that will work correctly with that site.
- **Configure QoS policy for stand-alone computers** Windows Vista and Windows Server 2008 support multiple local GPOs (MLGPOs). You can use MLGPOs to configure QoS policies on computers that are not a member of your Active Directory domain.

When you configure a QoS policy, you assign a DSCP value or throttle rate and then specify the criteria that Windows will use to identify the traffic that the QoS policy applies to. The criteria you can use are:

- Sending application (by file name, such as *application.exe*)
- Source or destination IPv4 or IPv6 network or address
- Source or destination TCP or UDP port number or range

### How it Works: QoS Policy Priorities

If a connection matches the criteria for multiple QoS policies, the most specific QoS policy is applied. For example, if you create a policy for an entire network (such as 192.168.1.0/24) and a second policy for a specific IP address (such as 192.168.1.5), the IP address policy, rather than the network policy, will be applied. The specific rules that Windows follows when applying QoS policies are:

1. User-level QoS policies take precedence over computer-level QoS policies.
2. QoS policies that identify applications take precedence over QoS policies that identify networks or IP addresses.
3. QoS policies that specify IP addresses and more-specific networks take precedence over QoS policies that specify less-specific networks.
4. QoS policies that specify port numbers take precedence over QoS policies that specify port ranges, which take precedence over QoS policies that do not specify a port number.
5. If multiple QoS policies still conflict, policies that specify source IP addresses take precedence over policies that specify destination IP addresses, and policies that specify a source port take precedence over policies that specify a destination port.

Only one QoS policy can be applied to any given connection. For example, if two traffic throttling policies apply to a single connection, the most specific policy will set the throttle rate—it is not cumulative.

To reduce the number of conflicts and simplify QoS deployment, design your QoS policies to be as specific as possible. For example, instead of applying a QoS policy to all traffic from a computer, apply the QoS policy to traffic with a specific port number from that computer.

## QoS Policies for Mobile Computers Running Windows Vista

Computers running Windows Server 2008 always apply QoS policies to all network interfaces. However, computers running Windows Vista operate slightly differently because they are designed to be mobile, so they might connect to networks outside of your organization. Different organizations might use different DSCP numbers or might not use QoS policies at all. Therefore, Windows Vista applies QoS policies only while connected to your internal network. Specifically, Windows Vista applies QoS policies only for domain network types. Windows Vista identifies a network as being part of a domain when it can contact a domain controller across that interface. Therefore, if a user connects to a wireless network at a coffee shop, Windows Vista will not apply your QoS policies. However, if the user connects to your internal network by using a VPN, Windows Vista will apply QoS policies to that VPN.



**More Info** For more information about network types in Windows Vista, refer to Chapter 26, “Configuring Windows Networking,” in the *Windows Vista Resource Kit* by Mitch Tulloch, Tony Northrup, and Jerry Honeycutt, with the Microsoft Windows Vista Team (Microsoft Press, 2007).

## Deployment Steps

After proper planning, it only takes a few minutes to implement QoS policies by using Group Policy settings. This section describes how to use GPOs to configure QoS policies that define DSCP values and traffic throttling. It also describes how to configure system-wide QoS settings such as the inbound TCP throughput level.

### How to Configure QoS by Using Group Policy

For QoS to be effective, you should configure as many computers as possible to assign DSCP values to traffic and, when prioritizing traffic is not sufficient, throttling outbound traffic.

#### To Configure QoS by Using Group Policy

1. In the console tree of the Group Policy Management Editor snap-in for the GPO to which you want to add the policy, open the Computer Configuration\Windows Settings\Policy-based QoS node or the User Configuration\Windows Settings\Policy-based QoS node.



**Note** User QoS policies are applied to user processes only when the user is logged on. For servers, you should almost always configure Computer QoS policies.

2. Right-click the Policy-Based QoS node, and then click Create New Policy.
3. The Policy-Based QoS Wizard appears. On the Create A QoS Policy page, shown in Figure 5-3, specify a unique name for the policy. Then, specify one of the DSCP values shown in Table 5-1 (which your network infrastructure can use to prioritize traffic) and a throttle rate in either Kbps or megabytes per second (MBps) (which Windows Vista will use to restrict outgoing bandwidth usage) as needed. Click Next.



**Note** Notice that throttle rate must be entered in kilobytes per second (KBps) or megabytes per second (MBps) rather than the more commonly used kilobits per second (Kbps) or megabits per second (Mbps)—notice the lowercase *b*. Eight bits equals one byte. Therefore, if you determine the Kbps or Mbps that you want to throttle at, divide that number by 8 when typing it into the Policy-Based QoS Wizard. For example, if you want to throttle at 128 Kbps, you would type **16 KBps**.

Policy-based QoS

Create a QoS policy  
A QoS policy applies a Differentiated Services Code Point (DSCP) value, throttle rate, or both to outbound TCP or UDP traffic.

Policy name:  
VoIP Server

☒ Specify DSCP Value:  
28

☐ Specify Throttle Rate:  
1 KBps

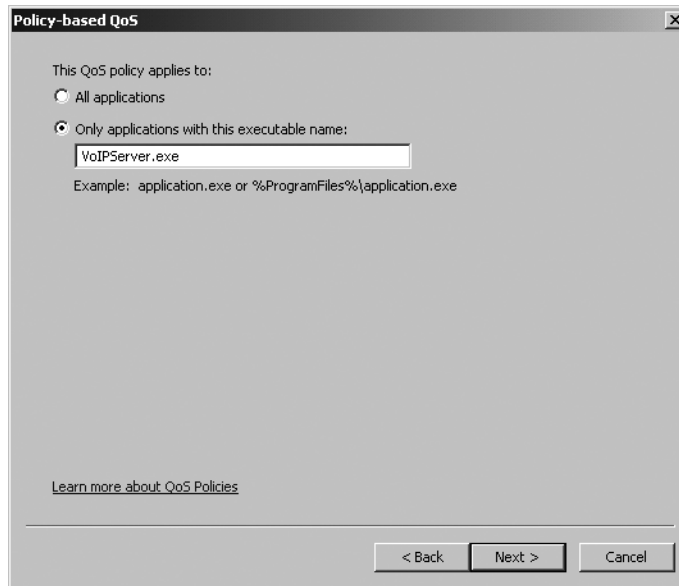
[Learn more about QoS Policies](#)

< Back   Next >   Cancel

**Figure 5-3** The Create a QoS Policy page

4. On the This QoS Policy Applies To page, shown in Figure 5-4, select either All Applications or Only Applications With This Executable Name. If you are specifying an application, Windows Vista will apply the DSCP value or throttle rate to network traffic generated by that application. If you must throttle a service, check the service properties

by viewing them in the Services snap-in. If the service has its own executable file (other than svchost.exe), you can specify that file. Otherwise, you can identify the traffic to apply the policy to by using the next two wizard pages. Click Next.



**Figure 5-4** The This QoS Policy Applies To page

5. On the Specify The Source And Destination IP Addresses page, shown in Figure 5-5, you can configure the policy to apply to traffic to or from a specific IP address or network. For example, if you want to configure a QoS policy that throttles traffic sent across a VPN, you would select Only For The Following Destination IP Address Or Prefix, and then type the destination network for the VPN. IPv4 and IPv6 addresses will both work, and you can use network prefix length representation to specify networks. With network prefix length representation, you would specify 192.168.1.0/24 to mean the entire 192.168.1.x network or 192.168.0.0/16 to mean the entire 192.168.x.x network. For more information, read “Subnets and Subnet Masks” at [http://www.microsoft.com/technet/prodtechnol/windows2000serv/reskit/cnet/cnbb\\_tcp\\_prux.msp](http://www.microsoft.com/technet/prodtechnol/windows2000serv/reskit/cnet/cnbb_tcp_prux.msp). Click Next.



**Note** QoS policies apply only to outgoing traffic. So the computer to which you’re applying the policy will always be identified by the source address, and the remote computer or network will always be identified by the destination address. Specify a source address if you want to identify the computers to which to apply the policy by using a technique other than the scope of the GPO.

**Policy-based QoS**

Specify the source and destination IP addresses.  
A QoS policy can be applied to outbound traffic that is from a source or to a destination IP (IPv4 or IPv6) address or prefix.

This QoS policy applies to:

☒ Any source IP address

☐ Only for the following source IP address or prefix:

Example for a host address: 192.168.1.1 or 3ffe:ffff::1  
Example for an address prefix: 192.168.1.0/24 or 3ffe:ffff::/64

This QoS policy applies to:

☒ Any destination IP address

☐ Only for the following destination IP address or prefix:

[Learn more about QoS Policies](#)

< Back   Next >   Cancel

**Figure 5-5** The Specify The Source And Destination IP Addresses page

- On the Specify The Protocol And Port Numbers page, shown in Figure 5-6, you can identify traffic based on TCP or UDP port numbers. For example, if you want to throttle all outgoing Web traffic from a Web server, you would select TCP, select From This Source Port Number Or Range, and then specify port 80 (the port number HTTP Web traffic uses). Click Finish.



**Note** When configuring QoS policies for servers, specify the source port number, and allow any destination port number. When configuring QoS policies for clients, specify the destination port number, and allow any source port number.

After creating a policy, you can edit it by right-clicking it in the details pane of the Group Policy Management Editor and then clicking Edit Existing Policy. The Edit An Existing QoS Policy dialog box has tabs that correspond to each of the pages in the Policy-Based QoS Wizard. For more information, see “Editing QoS Policies” and “Removing QoS Policies” later in this chapter.



**Note** Currently, using GPOs is the only way to configure QoS policies. Microsoft does not provide tools for configuring QoS policies by using scripts.

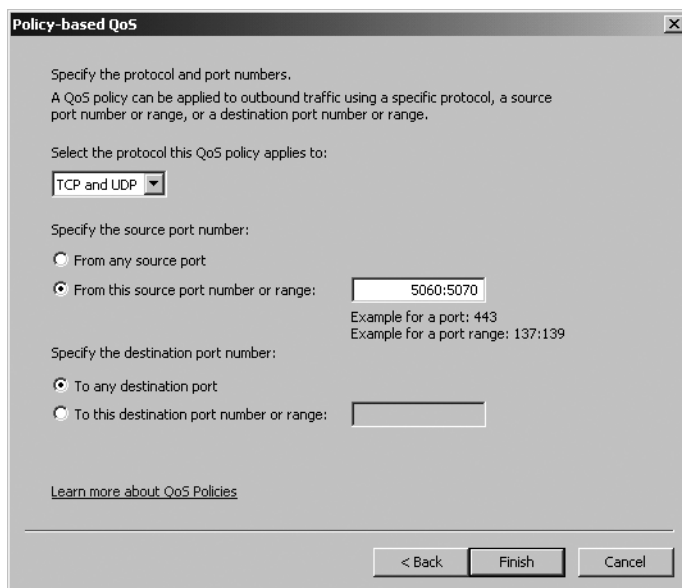


Figure 5-6 The Specify The Protocol And Port Numbers page

## How to Configure System-Wide QoS Settings

You can configure system-wide QoS settings within the Computer Configuration\Administrative Templates\Network\QoS Packet Scheduler node of Group Policy. You must modify these settings only if you must limit the outstanding packets, limit the bandwidth that can be reserved, or change the Packet Scheduler timer resolution. The policies available in the QoS Packet Scheduler node are as follows:

- **Limit outstanding packets** Specifies the maximum number of outstanding packets that can be issued to the network adapter at any given time. When this limit is reached, new packets are queued up in Pacer.sys until the network adapter completes a packet, at which point a previously queued packet is removed from the Pacer.sys queue and sent to the network adapter. This setting is disabled by default, and you should never need to enable this setting.
- **Limit reservable bandwidth** Controls the percentage of the overall bandwidth that the application can reserve. By default, this is set to 20%, which provides 80 percent of bandwidth to processes that do not have reserved bandwidth.
- **Set timer resolution** This value is not supported and should not be set.

The QoS Packet Scheduler node also has three sub-nodes that you can use to manually configure the standard DSCP values. The sub-nodes are:

- **DSCP value of conforming packets** These settings apply to packets that comply with flow specifications.

- **DSCP value of non-conforming packets** These settings apply to packets that do not comply with flow specifications.
- **Layer-2 priority value** These settings specify default link layer priority values for networks that support it.

You would need to change the values contained in these sub-nodes only if you have configured your network infrastructure to use non-standard DSCP values.

You can also configure advanced QoS settings for computers by using Group Policy. Within the Group Policy Management Editor, right-click the Computer Configuration\Windows Settings\Policy-based QoS node, and then click Advanced QoS Settings. You can use the resultant Advanced QoS Settings dialog box to:

- **Specify the inbound TCP throughput level** Most QoS policies relate to outbound traffic that the client computer sends. You can use this setting on the Inbound TCP Traffic tab to configure Windows so that it will attempt to throttle incoming traffic by adjusting the TCP receive window size, as discussed in “QoS for Inbound Traffic” earlier in this chapter. Table 5-3 lists the maximum TCP receive window for each inbound throughput level. By default, Windows will use the level 3 (maximum throughput) for TCP receive window size. Unlike policy-based QoS settings for outgoing traffic, this setting cannot control the rate of incoming traffic on a per-application, per-address, or per-port basis.

**Table 5-3 Maximum TCP Receive Windows**

| Inbound Throughput Level | Maximum |
|--------------------------|---------|
| 0                        | 64 KB   |
| 1                        | 256 KB  |
| 2                        | 1 MB    |
| 3                        | 16 MB   |



**Note** Because UDP traffic is not acknowledged, you cannot throttle UDP traffic from the receiving computer.

- **Control DSCP marking requests from applications** Applications can request their own DSCP values for outgoing network communications, but most applications do not specify a value. By default, Windows will use the DSCP value specified by an application. If you want Windows to ignore the DSCP value specified by the application and rely only on QoS policies to set DSCP values, select the Control DSCP Marking Requests From Applications check box, and then select Ignored, as shown in Figure 5-7.

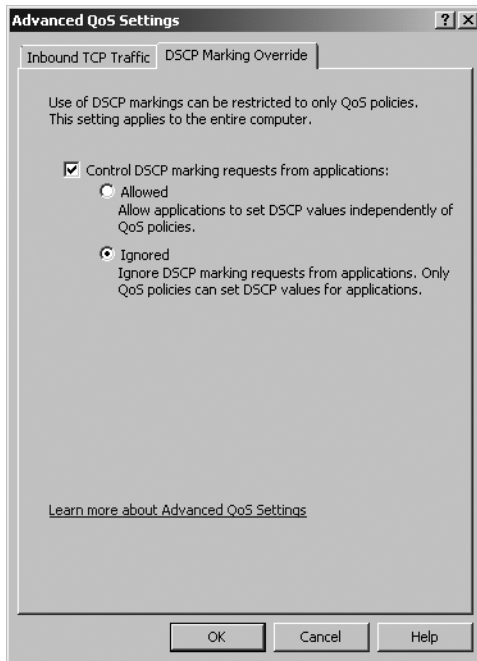


Figure 5-7 Ignoring application-specified DSCP values

## Ongoing Maintenance

Ongoing maintenance for QoS consists of updating or removing QoS policies as necessary and monitoring QoS policies to ensure that they are applied and functioning properly. The sections that follow describe ongoing maintenance in more detail.

## Removing QoS Policies

You can remove QoS policies by using the Group Policy Management console.

### To Remove a Policy

1. In Administrative Tools, open the Group Policy Management console.
2. Right-click the GPO containing the policy, and then click Edit.
3. In the Group Policy Management Editor, expand either User Configuration or Computer Configuration, expand Windows Settings, and then click Policy-Based QoS.
4. In the Details pane, right-click the policy you want to remove, and then click Delete Policy.
5. Click Yes when prompted.



## Editing QoS Policies

You can edit QoS policies by using the Group Policy Management console.

### To Edit a QoS Policy

1. In Administrative Tools, open the Group Policy Management console.
2. Right-click the GPO that you want to add the policy to, and then click Edit.
3. In the Group Policy Management Editor, expand either User Configuration or Computer Configuration, expand Windows Settings, and then click Policy-Based QoS.
4. In the Details pane, right-click the policy that you want to edit, and then click Edit Existing Policy.
5. Make any required changes, and then click OK.

The changes will be applied the next time computers refresh Group Policy. To immediately refresh Group Policy settings on a computer, run the command **gpupdate /force** from a command prompt with administrative privileges.

## Monitoring QoS

You can monitor QoS by using Performance Monitor (built into Windows), Network Monitor (a free download from Microsoft), or third-party monitoring tools. The sections that follow describe each of these tools.

### Performance Monitor


The Performance Monitor snap-in (available within the Computer Management console under Reliability And Performance\Monitoring Tools) provide some useful counters for gathering information about network performance:

- **Network Interface\Bytes Sent/sec** and **Network Interface\Packets Sent/sec** Show the total amount of traffic transmitted, including data that does and does not have QoS policies applied.
- **Pacer Flow\Bytes Transmitted** and **Pacer Flow\Packets Transmitted** Relative to the Network Interface counters, these counters are useful for gauging the portion of outgoing traffic that has a QoS policy applied.
- **Pacer Pipe\Nonconforming packets scheduled** and **Pacer Pipe\Nonconforming packets scheduled/sec** If these values are increasing, it means that QoS is enforcing traffic throttling by slowing down the transmission of packets.

### Network Monitor

Microsoft Network Monitor is a protocol analyzer, also known as a *sniffer*. Network Monitor captures raw network data and allows you to examine it. As shown in Figure 5-8, you can use

Network Monitor to examine the DSCP values in the IP header. In the figure, notice that the selected IPv4 packet has a DSCP value of 10 (bulk traffic). Therefore, you can use Network Monitor to verify that DSCP values are being applied and to perform detailed troubleshooting.



**Figure 5-8** Viewing the DSCP value in Network Monitor

You can also use Network Monitor to determine the TCP receive window being used, which you can configure by following the instructions in “How to Configure System-Wide QoS Settings” earlier in this chapter. After capturing traffic, examine the Window value in the TCP header, as shown in Figure 5-9. Windows will dynamically adjust this value, but it should always be below the value shown in Table 5-3 for the configured setting.

To download Network Monitor, visit <http://www.microsoft.com/downloads/>, and search for “Network Monitor.” For detailed instructions on how to use Network Monitor to capture and analyze network communications, refer to the Help site.

## Third-Party Monitoring Tools

Monitoring individual computers can provide some useful information about how QoS policies are being applied. However, only by monitoring your network infrastructure can you develop a comprehensive view of your network performance and the impact of QoS policies. Contact your network infrastructure provider for information about monitoring tools that provide insight into QoS performance.

You can also use third-party tools to monitor the performance of specific applications. For example, several developers (including Agilent and NetIQ) offer software that monitors VoIP performance. If you are implementing QoS to provide VoIP, use monitoring tools such as these to verify that you are meeting your performance requirements. If performance is low, increase bandwidth, reduce the amount of network traffic that QoS policies label as high priority, or both.

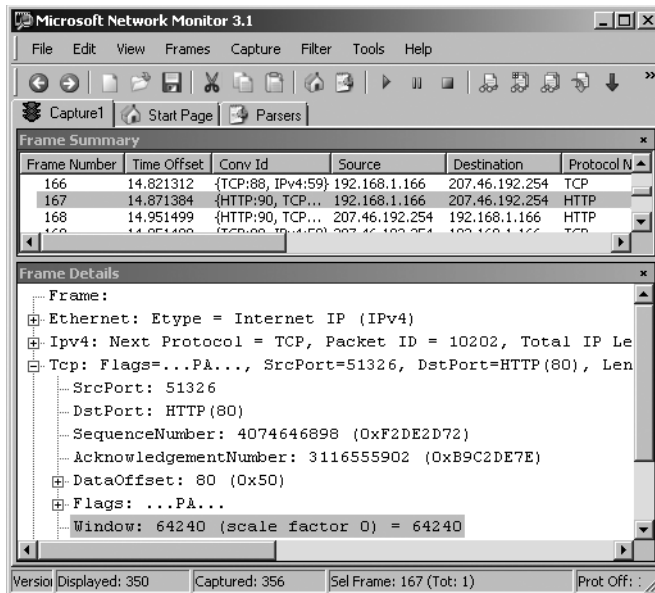


Figure 5-9 Viewing the TCP receive window in Network Monitor

## Troubleshooting

QoS policies should never cause outright connectivity problems. However, if QoS does not meet your performance expectations, you can analyze the policies and the configuration of your network infrastructure to verify that your implementation matches your design. The sections that follow describe techniques for troubleshooting problems with QoS policies and network performance.

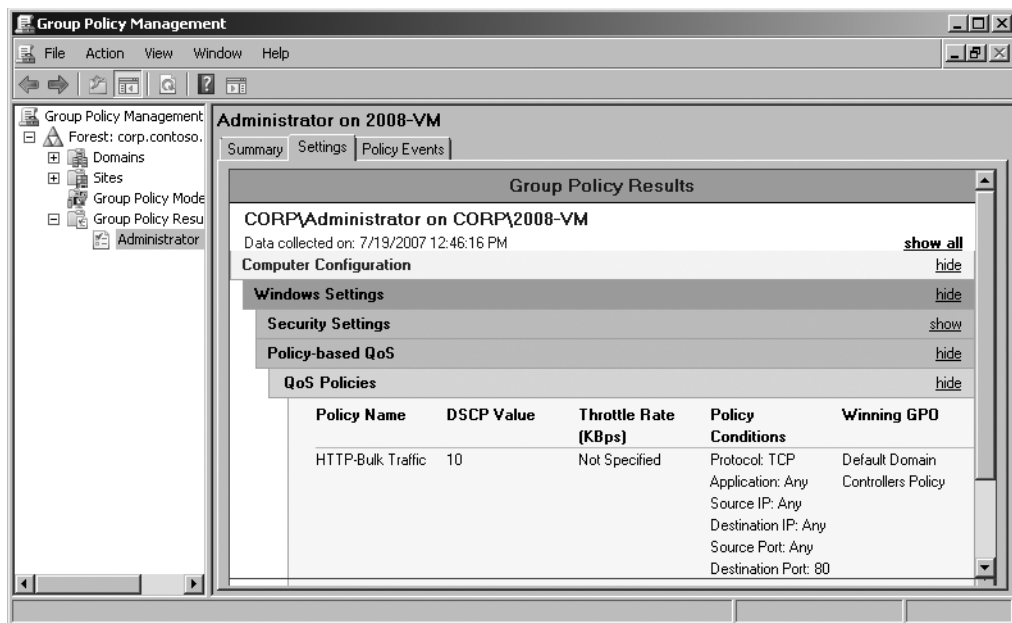
## Analyzing QoS Policies

You can use the Group Policy Results Wizard to generate a report of QoS policies applied to a computer or user.

### To Display QoS Policies

1. In Administrative Tools, open the Group Policy Management console.
2. Right-click the Group Policy Results node, and then click Group Policy Results Wizard.
3. On the Welcome To The Group Policy Results Wizard page, click Next.
4. On the Computer Selection page, accept the default setting by clicking Next.
5. On the User Selection page, accept the default setting by clicking Next.
6. On the Summary Of Selections page, click Next.
7. On the Completing The Group Policy Results Wizard page, click Finish.

8. In the Group Policy Management console, press Enter to accept the default name for report.
9. On the Settings tab, under both Computer Configuration and User Configuration, click Show For Policy-Based QoS. Then, click Show For QoS Policies.
10. As shown in Figure 5-10, the Group Policy Management console displays all QoS Policies that are applied to the computer or user.



**Figure 5-10** Viewing Group Policy Results

The Group Policy Management console shows the QoS policies with their DSCP value, throttle rate, policy conditions, and winning GPO (the GPO with the highest priority). For more information about QoS policy priorities, read “Planning GPOs and QoS Policies” earlier in this chapter.

### **Direct from the Source: Capturing QoS Tags with Network Monitor**

Consider a case where a network application calls Windows QoS APIs to add a layer-2 IEEE 802.1Q UserPriority tag (almost always referred to as 802.1p) to outgoing traffic. Ascertaining whether the tag was actually added to an outgoing packet is not as simple as it seems due to the nature of how the Windows network stack is designed and how framing actually occurs. From an internal implementation perspective, the QoS Packet Scheduler (Pacer.sys in Vista/2008 Server, and Psched.sys in XP/2003 Server) in the

network stack merely updates an out-of-band structure (not the actual formed packet) that an 802.1Q UserPriority tag should be added. The specific NDIS structure is `NDIS_NET_BUFFER_LIST_8021Q_INFO`, which contains member variables for both `VlanID` and `UserPriority` and is passed to the NDIS miniport driver for implementing both priority tagging (`UserPriority`) and VLAN (`VlanId`). It is up to the NDIS miniport driver to actually insert the 802.1Q tag into the frame based on these values before transmitting on the wire. A miniport driver will only insert this tag if the feature is supported and enabled in the advanced properties of the NIC driver; typically, layer-2 priority tagging is disabled by default.

From a network stack layering perspective, it's important to understand that `Pacer.sys` is an NDIS Lightweight Filter (LWF) driver and will always be inserted above a miniport driver, which will always be the lowest network software in the stack because it communicates directly with the NIC hardware. Also note that network sniffing applications like Microsoft Network Monitor are also network stack filters, and will always be inserted above the miniport driver. This is important knowledge because it should be clear that taking a network sniff of traffic on the sending computer will never show the tag in a packet (because the tag is added below the sniffing software).

What about trying to do a network sniff on the receiving computer? This is a good question, but it also will not show the layer-2 tag. The reason for this is that NDIS developer documentation clearly states that miniport drivers must strip the tag when received and populate the `NDIS_NET_BUFFER_LIST_8021Q_INFO` `UserPriority` and `VlanId` fields with the values in the tag. This out-of-band structure can then be used by NDIS filter drivers higher up in the stack for implementing these features. The functional reason for stripping the layer-2 tag is because `Tcpip.sys` will drop any received packet that contains this tag. Therefore, if a misbehaving miniport driver does not strip the tag, the packet will never be received by the user-mode application because it will be dropped internally.

In conclusion:

- A network sniffing app on the sending PC will never see a tag.
- A network sniffing app on the receiving PC will never see a tag.
- Monitoring tagged packets from intermediate network elements (such as a switch) is hard if at all possible.

*Gabe Frost, Product Manager*

*Core Windows Networking*

## Verifying DSCP Resilience

If you are not experiencing the performance benefit you expect from a QoS policy, first verify that the QoS policy is being applied correctly. Follow the steps in the section titled “Analyzing QoS Policies” earlier in this chapter to verify that the target computer has the appropriate QoS policies applied and that they match the traffic you are attempting to prioritize.

Next, use Network Monitor to verify that outgoing traffic has the correct DSCP value assigned to it. For more information, see “Network Monitor” earlier in this chapter. If the DSCP value is not assigned, the QoS policies are not being applied correctly. Verify that the GPO is being applied to the computer and that the QoS policy matches the traffic by application, port number, or IP address.

Because it’s possible for network infrastructure to remove the DSCP value from packets, you also must verify that the DSCP value is intact when packets reach the remote host. If the remote host is a computer running Windows, you can use Network Monitor to verify the DSCP value of the packets as they are received. If the remote host is not a computer running Windows, use another protocol analyzer. If the packets do not have the DSCP value intact when they reach the remote host, the network infrastructure is removing the DSCP value. Contact your network administrators for troubleshooting assistance.

If the DSCP value is intact when it reaches the remote host, the network infrastructure might not be correctly configured to prioritize traffic or might not support QoS. For best results, every router between the client and server should support QoS and be configured to prioritize packets based on their DSCP value. From the client, you can use the PathPing tool to determine a likely path between the client and server, as the following example demonstrates. (Code in bold indicates user input.)

### pathping www.contoso.com

Tracing route to contoso.com [10.46.196.103] over a maximum of 30 hops:

```

0  contoso-test [192.168.1.207]
1  10.211.240.1
2  10.128.191.245
3  10.128.191.73
4  10.125.39.213
5  gbr1-p70.cb1ma.ip.contoso.com [10.123.40.98]
6  tbr2-p013501.cb1ma.ip.contoso.com [10.122.11.201]
7  tbr2-p012101.cgci1.ip.contoso.com [10.122.10.106]
8  gbr4-p50.st6wa.ip.contoso.com [10.122.2.54]
9  gar1-p370.stwwa.ip.contoso.com [10.123.203.177]
10 10.127.70.6
11 10.46.33.225
12 10.46.36.210
13 10.46.155.17
14 10.46.129.51
15 10.46.196.103
```

The performance information that PathPing shows isn’t necessarily useful when troubleshooting QoS issues because PathPing uses Internet Control Message Protocol (ICMP) packets that

might be assigned a lower or higher priority than the traffic you are troubleshooting. Less frequently, the route between any two paths can vary depending on network conditions, or QoS settings might actually choose a different route for the traffic you are testing than for ICMP traffic.

Once you have used PathPing to identify a possible route between the client and the server, examine each router configuration to verify that it is not removing DSCP values and that it is correctly prioritizing traffic based on DSCP. If possible, use a protocol analyzer to verify that traffic reaching each router still has the DSCP value intact.

## Isolating Network Performance Problems

The most common concern with QoS is that high-priority traffic has too much latency or is not receiving sufficient bandwidth. First, follow the steps in “Analyzing QoS Policies” and “Verifying DSCP Resilience” earlier in this chapter to ensure that you have correctly configured QoS policies and your network infrastructure. Then, check for the following common problems:

- **Latency is near physical limits.** As discussed in “Latency” earlier in this chapter, increased distance causes increased latency because of the limitation of the physical speed of the signal. To minimize this impact, ensure that your routing is efficient. For example, if you have two offices on the East Coast and one office on the West Coast, routing traffic sent between the two East Coast offices through the West Coast office would incur a significant latency penalty. To rectify this, you could add a link directly between the East Coast offices. Similarly, routing traffic through a VPN almost always makes a route less efficient.
- **Bandwidth is near realistic limits.** If you cannot achieve throughput near your expectations, verify that your expectations are realistic for your network types. Wired Ethernet networks can achieve only 65 to 80 percent of their theoretical limits, whereas wireless networks are typically capable of only 35 to 50 percent of the stated bandwidth. Internet connections, including VPNs that use the Internet, are highly variable and dependent not only on your Internet service provider (ISP) but every ISP that might handle traffic between the source and destination.
- **The computer is busy.** If a computer has high processor utilization, it may not be able to handle incoming traffic efficiently, or it may reduce the responsiveness of the client or server application. You can eliminate this possible source of problems by stopping services or applications during testing.
- **The high-priority queues on routers are overused.** Most routers that support QoS will allow you to monitor the amount of traffic in each priority queue. The more packets in the queue, the higher the latency. To alleviate this, either increase the bandwidth on the destination network, or reduce the amount of high-priority traffic.
- **Drivers may be inefficient.** Verify that computers have updated versions of network interface drivers. Additionally, verify that router firmware is updated.

## Chapter Summary

Used properly, the policy-based QoS built into Windows Vista and Windows Server 2008 can improve efficiency of your network and the quality of network applications such as VoIP. Once you understand the common causes of network performance problems, including latency and jitter, you can create a plan to use QoS to optimize your available bandwidth.

A QoS deployment must include configuring both your network infrastructure and the computers on your network. Fortunately, you can use Group Policy settings to set QoS policies for computers running Windows Vista and computers running Windows Server 2008.

After deployment, you can monitor QoS performance by using Performance Monitor, Network Monitor, or third-party monitoring tools. If necessary, you can edit or remove QoS policies to achieve the QoS goals you set in the planning stage. If you are not achieving your goals, you can troubleshoot the performance problem by analyzing your QoS policies, verifying DSCP resilience, and isolating the specific network links that are introducing the problem.

## Additional Information

For additional information about QoS support in Windows, see the following:

- “Quality of Service” at <http://technet.microsoft.com/en-us/network/bb530836.aspx>
- RFC 2474, “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers,” at <http://www.ietf.org/rfc/rfc2474.txt>
- “The MS QoS Components” at <http://www.microsoft.com/technet/prodtechnol/windows2000serv/maintain/featusability/qoscomp.mspix>
- “Quality of Service in Windows Server ‘Longhorn’ and Windows Vista” at <http://www.microsoft.com/downloads/details.aspx?familyid=0230e025-9549-400b-807e-97e8a0cb9703>
- “Windows Vista Policy-based Quality of Service (QoS)” at <http://www.microsoft.com/downloads/details.aspx?FamilyID=59030735-8fde-47c7-aa96-d4108f779f20>
- “Policy-based QoS Architecture in Windows Server 2008 and Windows Vista: The Cable Guy, March 2006” at <http://www.microsoft.com/technet/community/columns/cableguy/cg0306.mspix>
- Network Quality of Service MSDN community forum at <http://forums.microsoft.com/MSDN/ShowForum.aspx?ForumID=825&SiteID=1>

For additional information about managing Group Policy in Windows, see the following:

- Microsoft Windows Server Group Policy at <http://www.microsoft.com/groupppolicy>
- Enterprise Management with the Group Policy Management Console at <http://go.microsoft.com/fwlink/?LinkID=8630>



## Chapter 6

# Scalable Networking

This chapter provides information about how to design, deploy, maintain, and troubleshoot networking features in the Windows Server 2008 operating system that are designed to support network throughput of over 1 gigabit while minimizing overhead on the computer's main processors. This chapter assumes that you have a solid understanding of Transmission Control Protocol/Internet Protocol (TCP/IP).

## Concepts

As network speeds increase, and applications take advantage of that increased bandwidth, the efficiency of client and server software must also increase. For example, consider a computer running the Windows Server 2003 operating system processing network traffic from several fully utilized gigabit or 10-gigabit Ethernet adapters:

- The large number of interrupts from the network adapters indicating that new packets have arrived can consume a significant amount of processor time.
- Processing of network data is limited to a single CPU core, even though many servers now have eight or more cores, limiting scalability.
- The act of moving data from the network adapter to the operating system requires memory copying, which is performed by the computer's processor and thus increases processor utilization.
- If Internet Protocol security (IPsec) communication is used, even more processing time is required for authentication and encryption.

These technical challenges lead to several real-world problems:

- Storage area networks (SANs) are inefficient because of the high overhead of TCP/IP, which slows storage consolidation efforts.
- Applications that use a significant amount of bandwidth, such as network backups, also incur significant processing overhead, slowing all applications.
- Storage, processing, and bandwidth might allow for server consolidation. However, the increased overhead of the cumulative network utilization, which must be handled by a single processor, would become a bottleneck.
- File and Web servers, which should be able to saturate any speed network, become bottlenecked on the utilization of a single processor. Therefore, multiple servers would be required to work around this performance limitation.

The sections that follow describe important network concepts related to scalable networking.



**More Info** TCP Chimney Offload, Receive-Side Scaling (RSS), and NetDMA were first introduced with the Windows Server 2003 Scalable Networking Pack. For more information, read "Windows Server 2003 Scalable Networking Pack Overview" at <http://www.microsoft.com/technet/community/columns/cableguy/cg0606.msp>. The Microsoft Windows 2000, Windows XP, and Windows Server 2003 operating systems are each capable of supporting IPsec Offload.

## TCP Chimney Offload

One of the reasons processor overhead is so significant when processing network communications is that the computer's processors must assemble the data from multiple TCP packets into a single segment. Figure 6-1 shows the TCP Chimney Offload architecture, which allows the network adapter to handle the task of segmenting TCP data for outgoing packets, reassembling data from incoming packets, and acknowledging sent and received data.

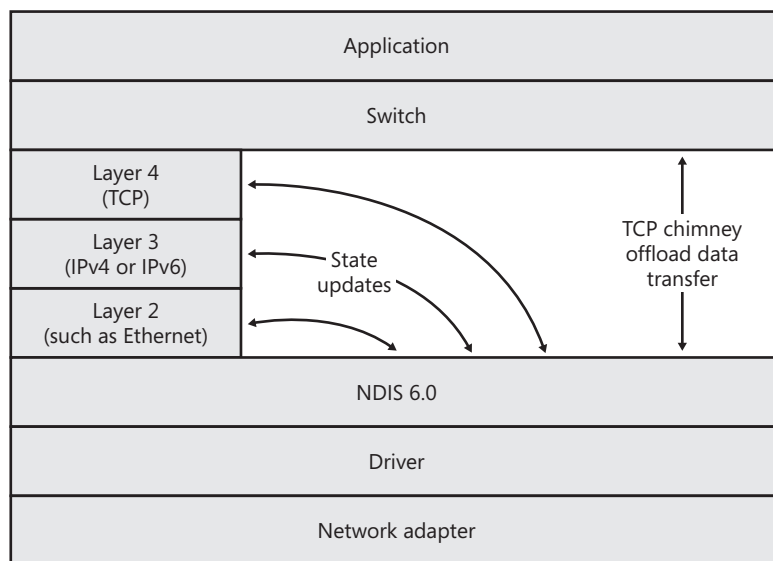


Figure 6-1 TCP Chimney Offload architecture

### How It Works: TCP Chimney Offload

With TCP Chimney Offload, the network adapter hands the data directly to a higher layer switch and communicates state updates only to the intermediate protocol layers, offloading much of the TCP overhead from the computer's processor. The switch layer chooses between the conventional software code path (in which data is passed through intermediate protocol layers) and the more efficient chimney. Without TCP Chimney Offload, all data transfer would need to travel through the Layers 2, 3, and 4 protocols.

TCP Chimney Offload supports both 32-bit and 64-bit versions of the Windows Vista and Windows Server 2008 operating systems and both 32-bit and 64-bit input/output (I/O) buses. TCP Chimney Offload is completely transparent to both systems administrators and application developers. TCP Chimney Offload is not compatible with QoS or adapter teaming drivers developed for earlier versions of Windows.



**Note** As the name suggests, TCP Chimney Offload does not change how non-TCP packets, including Address Resolution Protocol (ARP), Dynamic Host Configuration Protocol (DHCP), Internet Control Message Protocol (ICMP), and User Datagram Protocol (UDP), are handled.

TCP Chimney Offload still requires the operating system to process every application I/O. Therefore, it primarily benefits large transfers, and chatty applications that transmit small amounts of data will see little benefit. For example, file or streaming media servers can benefit significantly. However, a database server that is sending 100–500 bytes of data to and from the database might see little or no benefit.



**More Info** To examine TCP Chimney Offload performance testing data, read “Boosting Data Transfer with TCP Offload Engine Technology” at <http://www.dell.com/downloads/global/power/ps3q06-20060132-broadcom.pdf> and “Enabling Greater Scalability and Improved File Server Performance with the Windows Server 2003 Scalable Networking Pack and Alacritech Dynamic TCP Offload” at [http://www.alacritech.com/Resources/Files/File\\_Serving\\_White\\_Paper.pdf](http://www.alacritech.com/Resources/Files/File_Serving_White_Paper.pdf). For more information about TCP Chimney Offload, read “Scalable Networking: Network Protocol Offload—Introducing TCP Chimney” at [http://www.microsoft.com/whdc/device/network/TCP\\_Chimney.mspx](http://www.microsoft.com/whdc/device/network/TCP_Chimney.mspx).


## Receive-Side Scaling

As 10-gigabit LAN speeds become more common and we look to even higher speeds in the future, software must avoid becoming the performance bottleneck when it processes traffic it receives. One of the most significant bottlenecks is the processing time required for each packet.


Processing capability in computers has continued to increase over the years. However, instead of continuing to increase the clock speed of processors, computer hardware manufacturers have begun relying on multiple processors and multiple cores per processor. To allow Windows networking components to take advantage of this processing power, the software must avoid any process that is single threaded.

Windows Server 2003 supports Network Driver Interface Specification (NDIS) 5.1, which limits processing of incoming traffic to one processor at a time (though the particular processor used could vary depending on which one handled the interrupt), as shown in Figure 6-2.

With NDIS 6.0 and in Windows Vista and Windows Server 2008, the network interrupt service routine (ISR) can parallelize processing by queuing incoming packets received by an RSS-capable network adapter to multiple processors, as shown in Figure 6-3.

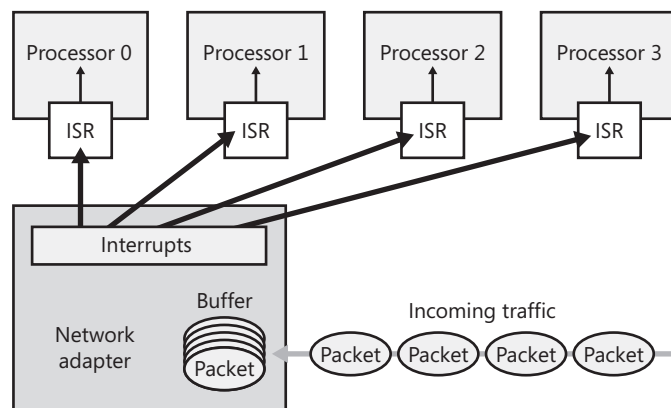


**Figure 6-2** NDIS 5.0 receive processing



**Figure 6-3** NDIS 6.0 receive processing with an RSS-capable network adapter

On PCI-e or PCI-X computers that support MSI or MSI-X, both the queuing and the interrupts can be distributed between multiple processors, as shown in Figure 6-4. Using RSS, applications and services still receive network data in order, but processor utilization in multiprocessor computers is more efficient.



**Figure 6-4** NDIS 6.0 receive processing with an RSS-capable network adapter that supports MSI or MSI-X

### Direct from the Source: MSI and MSI-X Interrupts

There are two methods for a PCI-e/PCI-X device to generate an interrupt:

- Line based
- MSI or MSI-X based

Line-based interrupts are the “old” way of generating interrupts, and most commonly, all line-based interrupts end up being serviced by a single CPU. However, modern systems that support Message Signaled Interrupts (MSI) enable the hardware device to generate an interrupt on any CPU they choose to. Thus RSS-capable NICs that also support MSI-based interrupts bring optimum performance by distributing both the ISRs across multiple CPUs as well as distributing the actual receive packet processing across multiple CPUs. Also note that Windows Vista and Windows Server 2008 are the first Windows operating systems to have software support for MSI/MSI-X systems and devices.

*Rade Trimceski, Program Manager*

*Windows Networking & Devices*

In addition to load balancing incoming traffic across all processors, Windows Server 2008 can also load-balance transmit processing caused by TCP window updates. In summary, RSS can increase transactions per second, connections per second, and network throughput for all multiprocessor computers, especially Web, file, backup, and database servers.



**Note** The default setting for RSS is for RSS to use the first four eligible processors (which is any processor except hyperthreaded virtual processors).



**More Info** For detailed information about RSS, read “Scalable Networking: Eliminating the Receive Processing Bottleneck—Introducing RSS” at [http://download.microsoft.com/download/5/D/6/5D6EAF2B-7DDF-476B-93DC-7CF0072878E6/NDIS\\_RSS.doc](http://download.microsoft.com/download/5/D/6/5D6EAF2B-7DDF-476B-93DC-7CF0072878E6/NDIS_RSS.doc).

## NetDMA

NetDMA, co-designed by Intel and Microsoft, is another technique for reducing the processor overhead associated with processing network traffic and increasing network throughput. NetDMA moves data directly from one location in the computer’s main memory directly to another location without requiring the data to be moved through the processor.

NetDMA requires the underlying hardware platform to support a technology such as Intel I/O Acceleration Technology (Intel I/OAT), a feature that can be used with Intel Xeon processors and Intel 5000 series chipsets. Intel's tests show that Intel I/OAT with NetDMA reduced processor utilization from 36 to 24 percent when four physical gigabit Ethernet network adapters were fully utilized in both directions, producing close to 8 gigabits per second (Gbps) of traffic. With eight-gigabit Ethernet adapters (producing close to 16 Gbps of traffic), Intel I/OAT and NetDMA increased throughput by more than 20 percent. With two or fewer gigabit Ethernet adapters in a computer (producing 4 Gbps or less of traffic), the improvement was minimal.



**More Info** For more information on Intel I/OAT, see <http://www.intel.com/go/ioat>.

NetDMA and TCP Chimney Offload are not compatible. If a network adapter supports both NetDMA and TCP Chimney Offload, Windows Server 2008 will use TCP Chimney Offload.



**More Info** For more information about NetDMA, read "Introduction to Intel I/O Acceleration Technology and the Windows Server 2003 Scalable Networking Pack" at <http://www.intel.com/technology/ioacceleration/317106.pdf>.

## IPsec Offload

IPsec can authenticate and encrypt network traffic without requiring changes to the application. However, authenticating or encrypting each packet requires some processor overhead. On servers that accept a large number of connections and are already processor-limited, the additional processor overhead associated with adding IPsec can cause the processor to become a performance bottleneck.



**Note** Encrypting data within an IPsec session requires processor time because it uses secret key encryption. However, IPsec uses public key encryption when the IPsec session is established to transfer that secret key. It's the public key encryption that takes the most processing time.

IPsec Offload moves IPsec processing to the network adapter, which typically has a processor optimized for handling authentication and encryption tasks. By adding an IPsec Offload card to a server, you can substantially reduce the overhead of using IPsec (which might or might not be significant, depending on the usage and processing capabilities of the server).

For more information about IPsec, see Chapter 4, "Windows Firewall with Advanced Security," and Chapter 16, "IPsec Enforcement."

## Planning and Design Considerations

Scalable networking features typically require the use of supported hardware. Some features require trade-offs, such as disabling software firewalls. Because of these costs, you must evaluate whether the benefits of each scalable networking feature outweigh the costs. The sections that follow guide you through the process of evaluating scalable networking features.

### Evaluating Network Scalability Technologies

When evaluating specific features, consider the following:

- **TCP Chimney Offload** TCP Chimney Offload will work only with NDIS 6.0 drivers on Windows Server 2008, NDIS 5.2 drivers on Windows Server 2003 with SP2, and compatible hardware. Therefore, if you have an NDIS 5.1 or earlier driver, or your network adapter does not support TCP Chimney Offload, it will not work. Because the performance benefits of TCP Chimney Offload are significant only with throughputs of about 2 Gbps or more, there is little benefit to using TCP Chimney Offload at network speeds below gigabit Ethernet, and the benefits will be more pronounced at 10-gigabit and faster speeds.
- **RSS and NetDMA** RSS uses processors more efficiently by distributing load across multiple processors, whereas NetDMA reduces the total amount of processing required for network traffic. In either case, if you need extra budget to purchase hardware that supports RSS or NetDMA, you should use load testing before you purchase the hardware to verify that the processor is limiting the computer's performance and that the server cannot meet your scalability requirements without specialized hardware. If no single processor is fully utilized, RSS and NetDMA will not offer a significant benefit.
- **IPsec Offload** Like RSS and NetDMA, IPsec Offload will improve performance only if the computer is processor-limited. IPsec Offload hardware does reduce the processing overhead associated with cryptographic functions but does not accelerate filter processing time. When testing IPsec Offload hardware, keep in mind that the Offload hardware typically supports a limited number of security associations (SAs). Above that limit, the computer's processors will handle the cryptographic functions as if the IPsec Offload hardware were not present.

During planning, you should also evaluate whether these scalability features are compatible with your server configuration. TCP Chimney Offload and NetDMA will not work with the following features:

- Windows Firewall
- IPsec
- Network Address Translation (NAT)
- Third-party firewalls

Additionally, RSS is not compatible with NAT drivers and is not effective for IPsec traffic unless it was decrypted with IPsec Offload. Table 6-1 illustrates which scalability technologies can benefit performance depending on the network technologies in use.

**Table 6-1 Network Technology Compatibility with Scalability Technologies**

| Technology            | TCP Chimney Offload | RSS                             | NetDMA | IPsec Offload |
|-----------------------|---------------------|---------------------------------|--------|---------------|
| Windows Firewall      | –                   | X                               | X      | X             |
| Third-party firewalls | –                   | X                               | X      | X             |
| IPsec                 | –                   | Only if IPsec Offload is in use | X      | X             |
| NAT                   | –                   | –                               | –      | X             |

Therefore, if you use any of these features and you determine that processing network communications is consuming too much processor time, you will need to rely on RSS and, if you use IPsec, IPsec Offload. Because using TCP Chimney Offload or NetDMA requires you to disable Windows Firewall and IPsec, you should use these features only on servers that have very high scalability requirements and that rely on external security devices, such as a network firewall, to filter traffic.

## Load Testing Servers

Each of the network scalability technologies discussed in this chapter can increase maximum throughput on your servers by decreasing processor utilization. However, if network adapters that support the technology are more costly than standard network adapters, it might not be worthwhile to adopt the technology. Before dedicating part of your hardware budget to these features, you should verify that you require the additional scalability and that network throughput or that the processor is limiting your server’s performance.



**Note** If you determine that network throughput or the processor is already limiting the performance of a production server, load testing might not be worth the effort. Instead, test the new hardware for compatibility, upgrade the server’s network adapter to hardware that supports TCP Chimney Offload, RSS, NetDMA, and, if you use IPsec, IPsec Offload, and monitor the performance in the production environment to determine the benefit.

You can use load testing software to test scalability of servers by simulating a large number of client requests. To avoid impacting your production network, perform the tests in a dedicated lab environment.



Microsoft provides the following tools for different types of servers:

- **Read80Trace and OSTRESS** Allow you to put stress on database servers. You can download these tools at <http://www.microsoft.com/downloads/details.aspx?familyid=5691ab53-893a-4aaf-b4a6-9a8bb9669a8b>.
- **Web Capacity Analysis Tool** Allows you to stress Web servers by submitting a large number of queries. This tool is included with the Internet Information Services (IIS) 6.0 Resource Kit Tools, but they will work with any Web server. You can download the tool at <http://www.microsoft.com/downloads/details.aspx?FamilyID=56fc92ee-a71a-4c73-b628-ade629c89499>.
- **Web Application Stress Tool** Another tool for stressing Web servers, available at <http://www.microsoft.com/downloads/details.aspx?FamilyID=e2c0585a-062a-439e-a67d-75a89aa36495>.
- **Windows Media Load Simulator** Allows you to stress test streaming media servers. For more information, visit <http://www.microsoft.com/windows/windowsmedia/howto/articles/loadsim.aspx>.

Additionally, third-party developers offer stress testing tools for a variety of different server applications. For internally developed applications, talk with your application development team about creating tools that simulate large numbers of client requests. For detailed information about creating custom load test tools by using Microsoft Visual Studio, read “Working with Load Tests” at [http://msdn2.microsoft.com/en-us/library/ms182561\(VS.80\).aspx](http://msdn2.microsoft.com/en-us/library/ms182561(VS.80).aspx).

## Monitoring Server Performance

It’s important that you monitor your server’s performance when using a load testing tool so that you can determine the component that is limiting performance (known as the bottleneck). Using the Performance Monitor snap-in, monitor the following counters to determine the limits of your network performance:

- **Processor\% Processor Time** Add `_Total` and, if you have multiple processors or multiple cores, add `<All Instances>`. `_Total` is useful for measuring the performance benefit of TCP Chimney Offload, NetDMA, and IPsec Offload. `<All Instances>` shows you the utilization of each processor, which is more useful for determining whether a single processor is bottlenecking performance and whether the server is benefitting from RSS.
- **Process\% Processor Time** Monitor the System instance (which will indicate the amount of processor time dedicated to processing network traffic, among other activities) and any other instances that might consume processor time. For example, if you are analyzing the performance of a database server, monitor the database process. When load testing file servers, you can assume that the majority of the System processor utilization can be attributed to processing network traffic.

- **Processor\Interrupts/sec** This number should decrease if you are using TCP Chimney Offload or another form of TCP Offload.
- **Network Interface\Bytes Received/sec and Network Interface\Bytes Sent/sec** These counters will help you understand the server's current load. When you apply sufficient load to reach the server's performance maximum, these numbers should be higher when network scalability features are enabled.
- **Network Interface\Packets Received/sec and Network Interface\Packets Sent/sec** When compared to Bytes Received/sec and Bytes Sent/sec, these counters will allow you to calculate the average number of bytes per packet. NetDMA and TCP Chimney Offload offer more significant benefits with larger packets, whereas RSS is effective with packets of any size.
- **TCPv4\Connections Active and TCPv6\Connections Active** These numbers will show you the current number of active TCP connections, which is helpful for understanding the server's current load.

### To Run Performance Monitor and Gather Data in Real-Time

1. Click Start, click Administrative Tools, and then click Reliability And Performance Monitor.
2. Select the Reliability And Performance\Monitoring Tools\Performance Monitor node.
3. Click the Add button (green plus sign) on the toolbar to add counters.

After adding the counters to Performance Monitor, you can create a data collector set to save data to a file for later analysis. This will allow you to compare the performance before and after implementing a scalability technology.


### To Create a Data Collector Set

1. In Reliability And Performance Monitor, right-click Performance Monitor, click New, and then click Data Collector Set.
2. Type a name for the data collector set, as shown in Figure 6-5. Then click Next.
3. Select a folder to save the data file in, and then click Next.
4. On the final page, click Finish.

After creating the data collector set, it will be available in the Data Collector Sets\User Defined node. Before you begin your load testing, right-click the data collector set, and then click Start. After you have completed the load test, right-click the data collector set, and then click Stop.

After collecting data, you can analyze it by following these steps:

1. In Reliability And Performance Monitor, right-click Performance Monitor, and then click Properties.



**Figure 6-5** The first page of the Create New Data Collector Set Wizard

2. On the Source tab, select Log Files, and then click Add. Select the log file you want to monitor, and then click Open.
3. Click OK to return to Performance Monitor and examine the data.

When examining the data, ask the following questions to evaluate the potential usefulness of scalability features:

- **Was any single processor fully utilized?** If the answer to this question is yes and the server has multiple processors, then RSS, TCP Chimney Offload, NetDMA, or, if you are using IPsec, IPsec Offload could improve performance.
- **Are the Bytes Sent/sec and Bytes Received/sec near the practical limit of the media?** If the answer is yes, scalability features won't improve network performance, but they can reduce processor utilization and provide more processing cycles to applications. If the answer is no and processors are not near full utilization, another network component is limiting your performance. You might need more load testing clients to fully utilize the server, or your network infrastructure might not be able to handle full speed.

## Deployment Steps

Prior to deploying scalable networking features, use load testing software to create a performance baseline of your servers, as discussed in the previous section. After deploying the scalable networking features, rerun the tests and compare the performance to the baseline to verify that you are achieving the expected performance improvements.

Most scalable networking features are enabled by default when compatible network adapters are installed in the computer. Therefore, configuration might not be required. The sections that follow show you how to examine the current configuration and enable or disable each of the scalable networking features.

## Configuring TCP Chimney Offload

TCP Chimney Offload is enabled by default. To view the current status, run the following command and examine the Chimney Offload State row:

```
netsh interface tcp show global
```

Even if TCP Chimney Offload is enabled, it will be active only when there is a compatible network adapter connected. To explicitly enable TCP Chimney Offload, run the following command:

```
netsh interface tcp set global chimney=enabled
```

To disable TCP Chimney Offload, run the following command:

```
netsh interface tcp set global chimney=disabled
```

TCP Chimney Offload will be enabled only if all the following is true:

- No firewall, including Windows Firewall, is enabled.
- No IPsec policies are applied.
- NAT is not enabled.

## Configuring Receive-Side Scaling

Receive-Side Scaling (RSS) is enabled by default. To view the current status, run the following command:

```
netsh interface tcp show global
```

Even if RSS is enabled, it will be active only when you have connected a compatible network adapter. To explicitly enable RSS, run the following command:

```
netsh interface tcp set global rss=enabled
```

To disable RSS, run the following command:

```
netsh interface tcp set global rss=disabled
```

## Configuring NetDMA

Windows does not include tools to configure NetDMA. You should use software provided by the hardware platform provider (such as Intel, in the case of Intel I/OAT) to configure and monitor NetDMA. To download the Intel I/OAT System Check Utility, visit <http://www.intel.com/support/network/adaptor/pro100/sb/CS-023725.htm>.

NetDMA will be enabled only if all the following is true:

- The network adapter does not report that it supports TCP Chimney Offload. (The two technologies are not compatible, and TCP Chimney Offload is preferred when both are available.)
- NAT is not enabled.

## Configuring IPsec Offload

IPsec Offload is enabled by default. To view whether IPsec Offload and all TCP/IP hardware acceleration are enabled, run the following commands at a command prompt and examine the “Task Offload” row:

```
netsh interface ipv4 show global
```

```
netsh interface ipv6 show global
```

Additionally, you can run the following commands to view the offload capabilities of the network adapters in more detail:

```
netsh interface ipv4 show offload
```

```
netsh interface ipv6 show offload
```

To enable or disable IPsec Offload, edit the HKEY\_LOCAL\_MACHINE\System\Current-ControlSet\Services\Ipsec\EnableOffload registry value. Set it to **0** to disable IPsec Offload, or **1** to enable IPsec Offload.

To explicitly enable IPsec Offload and all TCP/IP hardware acceleration, run the following commands:

```
netsh interface ipv4 set global taskoffload=enabled
```

```
netsh interface ipv6 set global taskoffload=enabled
```

To disable IPsec Offload and all TCP/IP hardware acceleration, run the following commands:

```
netsh interface ipv4 set global taskoffload=disabled
```

```
netsh interface ipv6 set global taskoffload=disabled
```

## Ongoing Maintenance

Once you have scalable networking features deployed, you should monitor network throughput and processor utilization on servers to verify that the features remain enabled and are functioning properly. If processor utilization increases or network throughput decreases, the scalable networking features might have been disabled. TCP Chimney Offload and NetDMA, in particular, are incompatible with many common network components and might be automatically disabled as an unwanted side effect of applying updates or configuration changes.

After you verify that scalable networking features provide you with performance benefits and work properly in your environment, you should monitor load on your servers to identify other servers that might benefit from these features. If you identify servers with high network and processor utilization, return to the planning and design phase to determine what hardware upgrades are required and whether enabling scalable networking features would be beneficial.

## Troubleshooting

If you experience poor network throughput, or network performance decreases after enabling TCP Chimney Offload or RSS, disable those features and test performance to determine whether that solves the problem. For instructions on how to disable those features, refer to “Deployment Steps” earlier in this chapter.

You might also be able to enable, disable, or configure scalability features by changing options in your network adapter driver.

### To View and Change the Network Adapter Driver Options

1. Click Start, right-click Computer, and then click Manage.
2. In the Server Manager console, expand Diagnostics, and then click Device Manager.
3. In the Details pane, expand Network Adapters.
4. Right-click your network adapter, and then click Properties.
5. The network adapter properties dialog box appears. Click the Advanced tab.
6. View the advanced properties, and change any settings.
7. Click OK to save your settings.

## Troubleshooting TCP Chimney Offload

To determine whether current connections are being offloaded, run the following command at a command prompt:

```
netstat -t
```

The output will resemble the following:

#### Active Connections

| Proto     | Local Address       | Foreign Address    | State       |
|-----------|---------------------|--------------------|-------------|
| Offload   | State               |                    |             |
| TCP       | 127.0.0.1:27015     | d820:49166         | ESTABLISHED |
| InHost    |                     |                    |             |
| TCP       | 127.0.0.1:49166     | d820:27015         | ESTABLISHED |
| Offloaded |                     |                    |             |
| TCP       | 192.168.1.161:49169 | by1msg3245816:msnp | ESTABLISHED |
| InHost    |                     |                    |             |
| TCP       | 192.168.1.161:50279 | MCE:5900           | ESTABLISHED |
| Offloaded |                     |                    |             |
| TCP       | 192.168.1.161:54109 | beta:5900          | ESTABLISHED |
| Offloaded |                     |                    |             |
| TCP       | 192.168.1.161:54880 | od-in-f103:http    | TIME_WAIT   |
| InHost    |                     |                    |             |
| TCP       | 192.168.1.161:54931 | 76.9.1.18:http     | TIME_WAIT   |
| Offloaded |                     |                    |             |

Netstat displays a list of all connections. The last column shows the current offload status. (You might need to increase the width of the command prompt to view the output easily.) The status will be one of the following:

- **In Host** The network connection is not being offloaded. (The computer's processor is handling it.)
- **Offloaded** The network connection is being handled by the network adapter.
- **Offloading** The network connection is in the process of being transferred to the network adapter.
- **Uploading** The network connection is in the process of being transferred back to the host processor.

To view applications in the TCP Chimney Offload table, run the following command at a command prompt:

```
netsh interface tcp show chimneyapplications
```

To view socket information in the TCP Chimney Offload table, run the following command at a command prompt:

```
netsh interface tcp show chimneyports
```

## Troubleshooting IPsec Offload

If you are using IPsec Offload, Network Monitor will display communications unencrypted, because the IPsec Offload hardware decrypts the data before Network Monitor captures them.

If you experience problems after enabling IPsec Offload, it's possible that the IPsec Offload component is causing compatibility problems. First, verify that you have the latest version of the network adapter driver. If problems persist, disable IPsec Offload by following the steps in "Configuring IPsec Offload" earlier in this chapter. If the problem does not occur with IPsec Offload disabled, you have isolated the cause of the problem as the IPsec Offload capability.

Once you determine that the IPsec Offload adapter is the cause of the problem, collect more information about the problem by doing the following:

- Examine the System event log for IPsec-related events.
- Create a Network Monitor capture, and use IPsec Monitor (Ipsecmon.exe) to analyze each connection attempt. Examine the Confidential Bytes Received counter in Ipsecmon to determine whether packets are being lost on receive.

Contact the IPsec Offload network adapter vendor for additional troubleshooting assistance.

## Chapter Summary

As network speeds increase, many enterprises are discovering that the network throughput of a server can be limited by the server's processors. Although you might expect a database server to dedicate a large amount of processing to the database service, in many cases, the server is spending significant processing time simply processing network communications. Typically, the performance impact becomes noticeable on servers that are transmitting and receiving more than 4 Gbps of sustained bandwidth, and the effect becomes significant above 8 Gbps throughput.

To allow servers to scale to multi-gigabit performance, Windows Server 2008 (when paired with compatible network adapters) supports four significant network scalability technologies:

- **TCP Chimney Offload** TCP data is handed directly to higher layers, bypassing Layer 2, 3, and 4 processing.
- **RSS** In a multi-processor computer, network processing can be handled by multiple processors simultaneously while maintaining in-order delivery.
- **NetDMA** Rather than moving network data through the processor, data is moved directly from the network adapter to the computer's memory.
- **IPsec Offload** Authentication and encryption tasks are handled by a dedicated processor on the network adapter, reducing utilization of the server's main processor.

Each of these technologies has trade-offs, however. First, each requires a network adapter that specifically supports the technology. TCP Chimney Offload and NetDMA cannot be used with Windows Firewall, IPsec, or NAT. NetDMA requires a specialized chipset in addition to a supported network adapter, and it cannot be used with TCP Chimney Offload.



To configure the technologies, use the Netsh command-line tool. Maintenance and troubleshooting requirements should be minimal, because the technologies should function transparently once configured.

## Additional Information

For additional information about scalable networking in Windows, see the following:

- “Scalable Networking” at <http://www.microsoft.com/snp>
- “Scalable Networking: Network Protocol Offload—Introducing TCP Chimney” at [http://www.microsoft.com/whdc/device/network/TCP\\_Chimney.msp](http://www.microsoft.com/whdc/device/network/TCP_Chimney.msp)
- “Scalable Networking: Eliminating the Receive Processing Bottleneck—Introducing RSS” at [http://download.microsoft.com/download/5/D/6/5D6EAF2B-7DDF-476B-93DC-7CF0072878E6/NDIS\\_RSS.doc](http://download.microsoft.com/download/5/D/6/5D6EAF2B-7DDF-476B-93DC-7CF0072878E6/NDIS_RSS.doc)
- “Microsoft Windows Scalable Networking Initiative” at <http://download.microsoft.com/download/5/b/5/5b5bec17-ea71-4653-9539-204a672f11cf/scale.doc>
- “Introduction to Intel I/O Acceleration Technology and the Windows Server 2003 Scalable Networking Pack” at <http://www.intel.com/technology/ioacceleration/317106.pdf>

To examine TCP Chimney Offload performance testing data, see the following:

- “Boosting Data Transfer with TCP Offload Engine Technology” at <http://www.dell.com/downloads/global/power/ps3q06-20060132-broadcom.pdf>
- “Enabling Greater Scalability and Improved File Server Performance with the Windows Server 2003 Scalable Networking Pack and Alacritech Dynamic TCP Offload” at [http://www.alacritech.com/Resources/Files/File\\_Serving\\_White\\_Paper.pdf](http://www.alacritech.com/Resources/Files/File_Serving_White_Paper.pdf)

For additional information about load testing, see the following:

- The Read80Trace and OSTRESS tools, available at <http://www.microsoft.com/downloads/details.aspx?familyid=5691ab53-893a-4aaf-b4a6-9a8bb9669a8b>
- The Web Capacity Analysis Tool, part of the Internet Information Services (IIS) 6.0 Resource Kit Tools, at <http://www.microsoft.com/downloads/details.aspx?FamilyID=56fc92ee-a71a-4c73-b628-ade629c89499>
- The Windows Media Load Simulator, available at <http://www.microsoft.com/windows/windowsmedia/howto/articles/loadsim.aspx>
- “Working with Load Tests” at [http://msdn2.microsoft.com/en-us/library/ms182561\(VS.80\).aspx](http://msdn2.microsoft.com/en-us/library/ms182561(VS.80).aspx)
- The Web Application Stress Tool at <http://www.microsoft.com/downloads/details.aspx?FamilyID=e2c0585a-062a-439e-a67d-75a89aa36495>

# Authentication Infrastructure

To deploy authenticated or protected network access, you must first deploy elements of a Microsoft Windows–based authentication infrastructure consisting of Active Directory, Group Policy, Remote Authentication Dial-In User Service (RADIUS), and a public key infrastructure (PKI). The set of elements you need to deploy depends on the type of network access and the design choices you make with regard to security, central configuration, and other issues. This chapter provides information about how to design and deploy these elements of an authentication infrastructure that can be used for wireless, wired, remote access, and site-to-site connections. Once deployed, elements of this infrastructure can also be used for Network Access Protection (NAP).

## Concepts

The following sections provide technical background on the following technologies that are used in the Windows-based authentication infrastructure:

- Active Directory Domain Services
- Group Policy
- PKI
- RADIUS

## Active Directory Domain Services

Active Directory Domain Services in the Windows Server 2008 operating system stores information about objects on the network and makes this information easy for administrators and users to find and use. Active Directory uses a structured data store as the basis for a logical, hierarchical organization of directory information. Active Directory Domain Services can be installed on servers running Windows Server 2008.

This data store, or directory, contains Active Directory objects. These objects typically include shared resources such as servers, volumes, printers, and the network user and computer accounts.

Security is integrated with Active Directory through logon authentication and through access control to objects in the directory. With a single network logon, administrators can manage and organize directory data throughout their network, and authorized users can access resources anywhere on the network. Policy-based administration eases the management of even the most complex network.

Active Directory also includes the following:

- A set of rules (or schema) that defines the classes of objects and attributes contained in the directory, the constraints and limits on instances of these objects, and the format of their names.
- A global catalog that contains information about every object in the directory. This catalog allows users and administrators to find directory information regardless of which domain in the directory actually contains the data.
- A query and index mechanism, which enables objects and their properties to be published and found by network users or applications.
- A replication service that distributes directory data across a network. All domain controllers in a domain participate in replication and contain a complete copy of all directory information for their domain. Any change to directory data is replicated to all domain controllers in the domain.

## User Accounts

Active Directory user accounts and computer accounts represent a physical entity such as a person, computer, or device. User accounts can also be used as dedicated service accounts for some applications.

User accounts and computer accounts (and groups) are also referred to as security principals. *Security principals* are directory objects that are automatically assigned security identifiers (SIDs), which can be used to access domain resources. A user or computer account is used to do the following:

- **Authenticate the identity of a user or computer.** A user account in Active Directory enables a user to log on to computers and domains with an identity that can be authenticated by the domain. Each user who logs on to the network should have his or her own unique user account and password. To maximize security, you should avoid multiple users sharing one account.
- **Authorize or deny access to domain resources.** When the user is authenticated, the user is authorized or denied access to domain resources based on the explicit permissions assigned to that user on the resource.
- **Administer other security principals.** Active Directory creates a foreign security principal object in the local domain to represent each security principal from a trusted external domain.
- **Audit actions performed using the user or computer account.** Auditing can help you monitor account security.

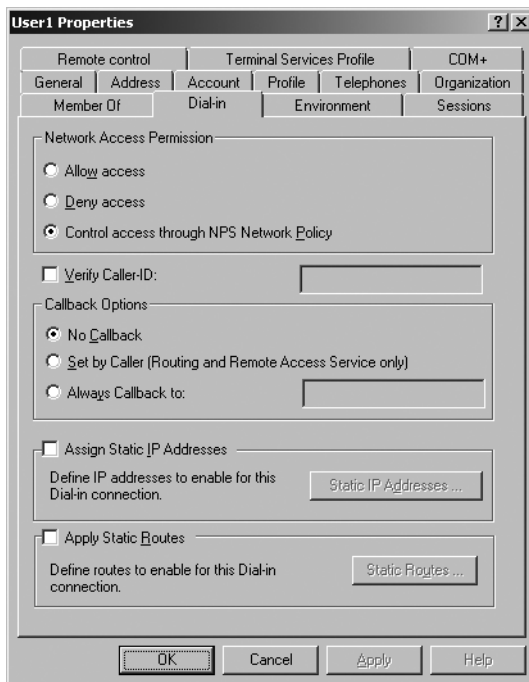
You can manage user or computer accounts by using the Active Directory Users And Computers snap-in.

Each computer that is running the Windows Vista, Windows XP, Windows Server 2008, or Windows Server 2003 operating system and that participates in a domain has an associated computer account. Similar to user accounts, computer accounts provide a means for authenticating and auditing computer access to the network and to domain resources.

User and computer accounts can be added, disabled, reset, and deleted using the Active Directory Users And Computers snap-in. A computer account can also be created when you join a computer to a domain.

## Dial-In Properties of an Account

User and computer accounts in Active Directory contain a set of dial-in properties that can be used when allowing or denying a connection attempt. In an Active Directory–based domain, you can set the dial-in properties on the Dial-In tab of the user and computer account properties dialog box in the Active Directory Users And Computers snap-in. Figure 9-1 shows the Dial-In tab for a user account in a Windows Server 2008 functional level domain.



**Figure 9-1** The Dial-In tab of a user account properties dialog box in a Windows Server 2008 functional level domain

On the Dial-In tab, you can view and configure the following properties:

- **Network Access Permission** You can use this property to set network access permission to be explicitly allowed, denied, or determined through Network Policy Server (NPS) network policies. NPS network policies are also used to authorize the connection attempt. If access is explicitly allowed, NPS network policy conditions and settings and

account properties can still deny the connection attempt. The Control Access Through NPS Network Policy option is available on user and computer accounts in a Windows Server 2008 functional level domain. New accounts that are created for a Windows Server 2008 functional level domain are set to Control Access Through NPS Network Policy.

- **Verify Caller ID** If this property is enabled, the access server verifies the caller's phone number. If the caller's phone number does not match the configured phone number, the connection attempt is denied. This setting is designed for dial-in connections.
- **Callback Options** If this property is enabled, the access server calls the caller back during the connection process. Either the caller or the network administrator sets the phone number that is used by the server. This setting is designed for dial-in connections.
- **Assign Static IP Addresses** You can use this property to assign a specific IP address to a user when a connection is made. This setting is designed for dial-in connections.
- **Apply Static Routes** You can use this property to define a series of static IP routes that are added to the routing table of the server running the Routing and Remote Access service when a connection is made. This setting is designed for demand-dial routing.

## Groups

A *group* is a collection of user and computer accounts and other groups that can be managed as a single unit. Users and computers that belong to a particular group are referred to as group members. Using groups can simplify administration by assigning a common set of permissions and rights to many accounts at once rather than assigning permissions and rights to each account individually.

Groups can be either directory-based or local to a particular computer. Active Directory provides a set of default groups upon installation and also allows you to create groups.

Groups in Active Directory allow you to do the following:

- Simplify administration by assigning permissions on a shared resource to a group rather than to individual users. This assigns the same access on the resource to all members of that group.
- Delegate administration by assigning user rights once to a group through Group Policy and then adding to the group members who require the same rights as the group.

Groups have a scope and type. Group *scope* determines the extent to which the group is applied within a domain or forest. Active Directory defines universal, global, and domain local scopes for groups. Group *type* determines whether a group can be used to assign permissions to a shared resource (for security groups); it also determines whether a group can be used for e-mail distribution lists only (for distribution groups).

Nesting allows you to add a group as a member of another group. You nest groups to consolidate member accounts and reduce replication traffic. Nesting options depend on the functional level of your domain. There are usually multiple domain functional levels, allowing for

a phased upgrade of an environment, enabling additional domain-native functionality at each progressive level.

When you have decided how to nest groups based on your domain functional level, organize your user and computer accounts into the appropriate logical groups for the organization. For a Windows Server 2008 functional level domain, you can use universal and nested global groups. For example, create a universal group named WirelessUsers that contains global groups of wireless user and computer accounts for wireless intranet access. When you configure your NPS network policy for wireless access, you must specify only the WirelessUsers group name.



**More Info** For more information about the types of groups, group scope, and domain functional levels, see the *Windows Server 2008 Active Directory Resource Kit* (Microsoft Press, 2008), which is available both as a stand-alone title and in the *Windows Server 2008 Resource Kit* (Microsoft Press, 2008); Windows Server 2008 Help and Support; or the resources at <http://www.microsoft.com/ad>.

## Public Key Infrastructure

A *public key infrastructure* (PKI) is a system of digital certificates and certification authorities (CAs) that verifies and authenticates the validity of each entity—such as a user, computer, or Windows service—that is participating in secure communications through the use of public key cryptography.

### Certification Authorities

When a certificate is presented to an entity as a means of identifying the certificate holder (the subject of the certificate), it is useful only if the entity being presented the certificate trusts the issuing CA. When you trust an issuing CA, it means that you have confidence that the CA has the proper policies in place when evaluating certificate requests and will deny certificates to any entity that does not meet those policies. In addition, you trust that the issuing CA will revoke certificates that should no longer be considered valid and will publish an up-to-date certificate revocation list (CRL). For more information about CRLs, see “Certificate Revocation” later in this chapter.

For Windows users, computers, and services, trust in a CA is established when you have a copy of the self-signed certificate of the root CA of the issuing CA locally installed and there is a valid certification path to the issuing CA. For a certification path to be valid, there cannot be any certificates in the certification path that have been revoked or whose validity periods have expired. The certification path includes every certificate issued to each CA in the certification hierarchy from a subordinate issuing CA to the root CA. For example, for a root CA, the certification path consists of a single certificate: its own self-signed certificate. For a subordinate CA, just below the root CA in the hierarchy, its certification path consists of two certificates: its own certificate and the root CA certificate.

If your organization is using Active Directory, trust in your organization's certification authorities will typically be established automatically based on decisions and settings made during the PKI deployment. For example, when joining a domain, a computer will automatically receive the organization's root CA through Group Policy settings.

## Certification Hierarchies

A certification hierarchy provides scalability, ease of administration, and consistency with a growing number of commercial and other CA products. In its simplest form, a certification hierarchy consists of a single CA. However, in general, a hierarchy will contain multiple CAs with clearly defined parent-child relationships. In this model, the subordinate certification authorities are certified by their parent CA-issued certificates, which bind a CA's public key to its identity. The CA at the top of a hierarchy is referred to as the *root authority*, or *root CA*. The child CAs of the root CAs are called *subordinate CAs*.

In Windows Server 2008 and Windows Vista, if you trust a root CA (when you have its certificate in your Trusted Root Certification Authorities certificate store), you trust every subordinate CA in the hierarchy unless a subordinate CA has had its certificate revoked by the issuing CA or has an expired certificate. Thus, any root CA is an important point of trust in an organization and should be secured and maintained accordingly.

Verification of certificates thus requires trust in only a small number of root CAs. At the same time, it provides flexibility in the number of certificate-issuing subordinate CAs. There are several practical reasons for supporting multiple subordinate CAs, including the following:

- **Usage** Certificates can be issued for a number of purposes, such as securing e-mail and network authentication. The issuing policy for these uses can be distinct, and separation provides a basis for administering these policies.
- **Organizational divisions** There might be different policies for issuing certificates, depending upon an entity's role in the organization. You can create subordinate CAs for the purpose of separating and administering these policies.
- **Geographic divisions** Organizations might have entities at multiple physical sites. Network connectivity between these sites might dictate a requirement for multiple subordinate CAs to meet usability requirements.
- **Load balancing** If your PKI will support the issuing of a large number of certificates, having only one CA issue and manage all these certificates can result in considerable network load for that single CA. Using multiple subordinate certification authorities to issue the same kind of certificates divides the network load among certification authorities.
- **Backup and fault tolerance** Multiple certification authorities increase the possibility that your network will always have operational certification authorities available to service users.

Such a certificate hierarchy also provides administrative benefits, including the following:

- Flexible configuration of the CA security environment to tailor the balance between security and usability.

For example, you might choose to employ special-purpose cryptographic hardware on a root CA, operate it in a physically secure area, or operate it offline. These security measures might be unacceptable for subordinate CAs because of cost or usability considerations.

- The ability to deactivate a specific portion of the CA hierarchy without affecting the established trust relationships.


For example, you can easily shut down and revoke an issuing CA certificate that is associated with a specific geographic site without affecting other parts of the organization.

By using the Certificates snap-in, you can view the certification path for a certificate on the Certification Path tab of the properties dialog box of a certificate.

For a small business environment, a certificate hierarchy consisting of a single root CA that is also the issuing CA is adequate. For a medium-sized organization, a single root CA with a single level of issuing CAs is adequate. For an enterprise network, you can deploy a three-tiered CA hierarchy, consisting of the following:

- A root CA that is offline (not available on the network)
- A layer of intermediate CAs that are offline
- A layer of issuing CAs that are online

This CA hierarchy provides flexibility and insulates the root CA from attempts by malicious users to compromise its private key. The offline root and intermediate CAs are not required to be Windows Server 2008–based or Windows Server 2003–based CAs. Issuing CAs can be subordinates of a third-party intermediate CA. Figure 9-2 shows a three-level enterprise network certificate hierarchy.



**Figure 9-2** Three-level certificate hierarchy for enterprise networks



## Certificate Revocation

Revocation of a certificate invalidates that certificate as a trusted security credential prior to the natural expiration of its validity period. There are a number of reasons why a certificate, as a security credential, could become untrustworthy prior to its expiration, including the following:

- Compromise or suspected compromise of the certificate subject's private key
- Compromise or suspected compromise of a CA's private key
- Discovery that a certificate was obtained fraudulently
- Change in the status of the certificate subject as a trusted entity
- Change in the name of the certificate subject

A PKI depends on distributed verification of credentials in which there is no need for direct communication with the central trusted entity that vouches for the credentials. This creates a need to distribute certificate revocation information to individuals, computers, and applications attempting to verify the validity of certificates. The need for revocation information and its timeliness will vary according to the application and its implementation of certificate revocation checking. To effectively support certificate revocation, the validating entity must determine whether the certificate is valid or has been revoked.

Certificate revocation lists (CRLs) are digitally signed lists of unexpired certificates that have been revoked. Clients retrieve this list and can then cache it (based on the configured lifetime of the CRL) and use it to verify certificates presented for use. Because CRLs can become large, depending on the size of the CA, delta CRLs can also be published. *Delta CRLs* contain only the certificates revoked since the last base CRL was published, which allows clients to retrieve the smaller delta CRL and quickly build a complete list of revoked certificates. The use of delta CRLs also allows more frequent publishing because the size of the delta CRL usually does not require as much overhead as a full CRL.

Windows Server 2008 supports industry-standard methods of certificate revocation. These methods include publication of CRLs and delta CRLs in several locations for clients to access in Active Directory and on Web servers and network file shares. Certificate revocation also can be checked by using the Online Certificate Status Protocol (OCSP), which uses the Hypertext Transfer Protocol (HTTP) to obtain a definitive digitally signed response indicating a certificate's revocation status.

## Certificate Validation

The certificates that are offered during the negotiation for secure communication must be validated before secure communication can begin. For example, for network access authentication using Extensible Authentication Protocol-Transport Layer Security (EAP-TLS), the authentication server (the RADIUS server) must validate the certificate offered by the IEEE

802.1X or Point-to-Point Protocol (PPP) client. For authentication using either EAP-TLS or Protected EAP (PEAP), the 802.1X or PPP client can be configured to validate the certificate offered by the authentication server.

## Windows Certificate Support

Windows has built-in support for certificates as follows:

- Every computer running Windows Vista, Windows Server 2008, Windows XP, or Windows Server 2003 has the ability, subject to Windows security and permissions, to store computer and user certificates and manage them by using the Certificates snap-in.
- Windows Server 2008 includes Active Directory Certificate Services and Windows Server 2003 includes Certificate Services, both of which allow a Windows server to act as a CA.

Certificate Services provides customizable services for issuing and managing certificates used in software security systems employing public key technologies. You can use Certificate Services in Windows Server 2008 and Windows Server 2003 to create a CA that will receive certificate requests, verify both the information in the request and the identity of the requester, issue certificates, revoke certificates, and publish CRLs.

You can also use Certificate Services to do the following:

- Enroll users for certificates from the CA by using a Web page (known as Web enrollment), through the Certificates snap-in, or transparently through autoenrollment.
- Use certificate templates to help simplify the choices that a certificate requester must make when requesting a certificate, depending upon the policy used by the CA.
- Take advantage of Active Directory for publishing trusted root certificates to domain member computers, publishing issued certificates, and publishing CRLs.
- Implement the ability to log on to a Windows domain by using a smart card.

If your organization is using Certificate Services, the CA is one of two types:

- **Enterprise CA** An enterprise CA depends on Active Directory being present. An enterprise CA offers different types of certificates to a requester based on the certificates it is configured to issue in addition to the security permissions of the requester. An enterprise CA uses information available in Active Directory to help verify the requester's identity. An enterprise CA can publish its CRL to Active Directory, a Web site, or a shared directory. You can use the Certificate Request Wizard within the Certificates snap-in, CA Web pages (Web enrollment), and autoenrollment to request certificates from an enterprise CA.
- **Standalone CA** For a user, a Standalone CA is less automated than an enterprise CA because it does not require or depend on the use of Active Directory. Standalone certification authorities that do not use Active Directory generally must request that the

certificate requester provide more complete identifying information. A Standalone CA makes its CRL available from a shared folder or from Active Directory if it is available. By default, users can request certificates from a Standalone CA only through Web enrollment.



**More Info** For more information about PKI support in Windows, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/pki>.

## Group Policy

The Group Policy management solution in Windows allows administrators to set configurations for both server and client computers. Local policy settings can be applied to all computers, and for those that are part of a domain, an administrator can use Group Policy to set policies that apply across a given site, domain, or organizational unit (OU) in Active Directory or that apply to a security group. Support for Group Policy is available on computers running Windows Vista, Windows Server 2008, Windows XP, and Windows Server 2003.

Through an Active Directory infrastructure and Group Policy, administrators can take advantage of policy-based management to do the following:

- Enable one-to-many management of users and computers throughout the enterprise.
- Automate enforcement of IT policies.
- Simplify administrative tasks such as system updates and application installations.
- Consistently implement security settings across the enterprise.
- Efficiently implement standard computing environments for groups of users.

Group Policy can be used to specify user-related policies and security, networking, and other policies applied at the computer level for management of domain controllers, member servers, and desktop user computers.

The GPMC snap-in provides a unified graphical user interface for deploying and managing Group Policy settings and enables script-based management of Group Policy operations. You can also use the Group Policy Management Editor snap-in.

On Windows Server 2008, you must install the Group Policy Management feature to use the Group Policy management tools such as the GPMC snap-in and Group Policy Management Editor snap-in.

## Group Policy Overview

Administrators can manage computers centrally through Active Directory and Group Policy. Using Group Policy to deliver managed computing environments allows administrators to

work more efficiently because of the centralized, one-to-many management it enables. Measurements of total cost of ownership (TCO) associated with administering distributed personal computer networks reveal lost productivity for users as one of the major costs for corporations. Lost productivity is frequently attributed to user errors—such as modifying system configuration files and thus rendering a computer unusable—or to complexity, such as the availability of nonessential applications and features on the desktop. Because Group Policy defines the settings and allowed actions for users and computers, it can create desktops that are tailored to users' job responsibilities and level of experience with computers.

**Setting Group Policy** By creating Group Policy settings, administrators use Group Policy to specify configurations for groups of users and computers. These settings are specified through the GPMC snap-in or the Group Policy Management Editor snap-in and are contained in a Group Policy Object (GPO), which is in turn linked to Active Directory containers—such as sites, domains, and OUs—and security groups.

In this way, Group Policy settings are applied to the users and computers in those Active Directory containers or security groups. Administrators can configure the users' work environment once and rely on the user's computer to enforce the policies as set.

**Group Policy Capabilities** Through Group Policy, administrators set the policies that determine how applications and operating systems are configured to keep users and systems functional and secure. Group Policies can be used for the following:

- **Registry-based policy** The most common and the easiest way to provide a policy for an application or operating system component is to implement a registry-based policy. By using the GPMC snap-in or the Group Policy Management Editor snap-in, administrators can create registry-based policies for applications, the operating system, and its components. For example, an administrator can enable a policy setting that removes the Run command from the Start menu for all affected users.
- **Security settings** Group Policy provides to administrators options for setting security options for computers and users within the scope of a GPO. Local computer, domain, and network security settings can be specified. For added protection, you can apply software restriction policies that prevent users from running files based on the path, URL zone, hash, or publisher criteria. You can make exceptions to this default security level by creating rules for specific software.

## Using Group Policy

Administrators use Group Policy and Active Directory together to institute policies across domains, sites, and OUs according to the following rules:

- GPOs are stored on a per-domain basis.
- Multiple GPOs can be associated with a single site, domain, or OU.
- Multiple sites, domains, or OUs can use a single GPO.

- Any site, domain, or OU can be associated with any GPO, even across domains (although doing so slows performance).
- The effect of a GPO can be filtered to target particular groups of users or computers based on their membership in a security group.

**Computer and User Configuration** Administrators can configure specific desktop environments and enforce policy settings on groups of computers and users on the network as follows:

- **Computer configuration** Computer-related policies specify operating system behavior, desktop behavior, application settings, security settings, assigned applications options, and computer startup and shutdown scripts. Computer-related policy settings are applied during the computer startup process and during a periodic refresh of Group Policy.
- **User configuration** User-related policies specify operating system behavior, desktop settings, application settings, security settings, assigned and published application options, user logon and logoff scripts, and folder redirection options. User-related policy settings are applied when users log on to the computer and during the periodic refresh of Group Policy.

**Applying Group Policy** Group Policy is applied in an inherited and cumulative fashion and affects all computers and users in an Active Directory container. Group Policy is applied when the computer starts up and when the user logs on. When a user turns on the computer, the system applies computer-based Group Policy settings. When a user logs on interactively, the system loads the user's profile and then applies user-based Group Policy settings. By default, policy settings are reapplied every 90 minutes. (You can set this period between 0 and 45 days.) You can also locally reapply policy settings on demand by running the **gpupdate** command at a Windows command prompt.

When applying policy, the system queries the directory service for a list of GPOs to process. Active Directory resources that are enforced with Group Policy settings will require read access to the GPOs. If a computer or user is not allowed access to a GPO, the system does not apply the specified policy settings. If access is permitted, the system applies the policy settings specified by the GPO.

The scope of Group Policy can extend from a single computer—the local GPO that all computers include—to Active Directory sites, domains, and OUs. For example, a GPO might be linked to an Active Directory site to specify policy settings for proxy settings and network-related settings that are specific to that site. A GPO becomes useful only after it is linked to a container—the settings in the GPO are then applied according to the scope of the container.

GPOs are processed in the order of local, site, domain, and then OU. As a result, a computer or user receives the policy settings of the last Active Directory container processed—that is, a policy applied later overwrites policy applied earlier.



**More Info** For more information about Group Policy in Windows, see the *Microsoft Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/gp>.

## RADIUS

When deploying a network access authentication infrastructure, it is possible to have each network access server store the account information and credentials for authentication and the network access policies for connection authorization. When a connection attempt is made, the access server can authenticate the connection attempt against the locally stored accounts and credentials, evaluate whether the connection attempt is authorized through the local account properties and network access policies, and locally store information about the connection attempt for later analysis. However, this method does not scale, especially in an enterprise environment with a large number of access servers. A scalable and more manageable solution is to offload the authentication and authorization evaluation and the storage of each connection attempt onto a central server that can utilize the existing accounts database.

RADIUS is a widely deployed protocol that allows authentication, authorization, and accounting for network access to be centralized at RADIUS servers. Originally developed for dial-up remote access, RADIUS is now supported by wireless access points (APs), authenticating Ethernet switches, virtual private network (VPN) servers, Digital Subscriber Line (DSL) access servers, and other types of network access servers.




**More Info** RADIUS is described in Request for Comments (RFC) 2865, "Remote Authentication Dial-In User Service (RADIUS)," and RFC 2866, "RADIUS Accounting." The listed RFCs can be viewed at <http://www.ietf.org/rfc.html>.

### Components of a RADIUS Infrastructure

A RADIUS authentication, authorization, and accounting infrastructure consists of the following components:

- Access clients
- Access servers (RADIUS clients)
- RADIUS servers
- User account databases
- RADIUS proxies

Figure 9-3 shows the components of a RADIUS infrastructure.



**Figure 9-3** The components of a RADIUS infrastructure

These components are described in detail in the following sections.

**Access Clients** An access client requires access to a network or another part of the network. Examples of access clients are dial-up or VPN remote access clients, wireless clients, or LAN clients connected to an authenticating switch. Access clients are not RADIUS clients.

**Access Servers (RADIUS Clients)** An access server provides access to a network. An access server using a RADIUS infrastructure is also a RADIUS client, which uses the RADIUS protocol to send connection requests and accounting messages to a RADIUS server. Examples of access servers include:

- Wireless APs that provide physical layer access to an organization's network by using wireless-based transmission and reception technologies.
- Switches that provide physical layer access to an organization's network by using traditional LAN technologies such as Ethernet.
- Network access servers (NASs) that provide remote access connectivity to an organization's network or the Internet. An example is a computer running Windows Server 2008 and Routing and Remote Access and providing either traditional dial-up access or VPN-based remote access to an organization's intranet.

- Network Access Protection (NAP) enforcement points that collect a NAP client's system health status and send it to a Windows Server 2008–based RADIUS server for evaluation. Examples include NAP-enabled Dynamic Host Configuration Protocol (DHCP) servers and Health Registration Authorities (HRAs). For more information about NAP enforcement points, see Chapter 14, “Network Access Protection Overview.”

**RADIUS Servers** A RADIUS server receives and processes connection requests or accounting messages sent by RADIUS clients or RADIUS proxies. During a connection request, the RADIUS server processes the list of RADIUS attributes in the connection request. Based on a set of rules and the information in the user account database, the RADIUS server authenticates and authorizes the connection and sends back either an accept or reject message. The accept message can contain connection restrictions that are enforced by the access server for the duration of the connection.



**Note** The NPS component of Windows Server 2008 is an industry standard–compliant RADIUS server.

**User Account Databases** A user account database is a list of user accounts and their properties that can be checked by a RADIUS server to verify authentication credentials and to obtain user account properties containing authorization and connection setting information.

The two user account databases that NPS can use are the local Security Accounts Manager (SAM) and Active Directory. For Active Directory, NPS can provide authentication and authorization for user or computer accounts in the domain in which the NPS server is a member, two-way trusted domains, and trusted forests with domain controllers running Windows Server 2008 or Windows Server 2003.

If the user accounts for authentication reside in a different type of database, you can use a RADIUS proxy to forward the authentication request to another RADIUS server that has access to the user account database.

**RADIUS Proxies** A RADIUS proxy routes RADIUS connection requests and accounting messages between RADIUS clients and RADIUS servers. The RADIUS proxy uses information within the RADIUS message to route the RADIUS message to the appropriate RADIUS client or server.

A RADIUS proxy can be used as a forwarding point for RADIUS messages when the authentication, authorization, and accounting must occur at multiple RADIUS servers within an organization or in different organizations.

With the RADIUS proxy, the definitions of *RADIUS client* and *RADIUS server* become blurred. A RADIUS client to a RADIUS proxy can be an access server (that originates connection requests or accounting messages) or another RADIUS proxy (in a chained proxy configuration). There can be multiple RADIUS proxies between the originating RADIUS client and the



final RADIUS server using chained RADIUS proxies. In a similar way, a RADIUS server to a RADIUS proxy can be the final RADIUS server (at which the RADIUS message is evaluated for authentication and authorization) or another RADIUS proxy. Therefore, when referring to RADIUS clients and servers from a RADIUS proxy perspective, a RADIUS client is the RADIUS entity that receives RADIUS request messages, and a RADIUS server is the RADIUS entity that forwards RADIUS request messages.



**Note** The NPS component of Windows Server 2008 is an industry standard-compliant RADIUS proxy.

### How It Works: RADIUS Messages and the RADIUS Authentication, Authorization, and Accounting Process

RADIUS messages are sent as User Datagram Protocol (UDP) messages. RADIUS authentication messages are sent to destination UDP port 1812, and RADIUS accounting messages are sent to UDP port 1813. Legacy access servers might use UDP port 1645 for RADIUS authentication messages and UDP port 1646 for RADIUS accounting messages. Only one RADIUS message is included in the UDP payload of a RADIUS packet.

A RADIUS message consists of a RADIUS header and RADIUS attributes. Each RADIUS attribute contains a specific item of information about the connection. For example, there are RADIUS attributes for the user name, the user password, the type of service requested by the user, the type of access server, and the IP address of the access server.

RADIUS attributes are used to convey information between RADIUS clients, RADIUS proxies, and RADIUS servers. For example, the list of attributes in the RADIUS Access-Request message includes information about the user credentials and the parameters of the connection attempt. In contrast, the list of attributes in the Access-Accept message includes information about the type of connection that can be made, connection constraints, and any vendor-specific attributes (VSAs).



**More Info** RADIUS attributes are described in RFCs 2548, 2865, 2866, 2867, 2868, 2869, 3162, and 3579. RFCs and Internet drafts for VSAs define additional RADIUS attributes. The listed RFCs can be viewed at <http://www.ietf.org/rfc.html>.

RFCs 2865 and 2866 define the following RADIUS message types:

- **Access-Request** Sent by a RADIUS client to request authentication and authorization for a network access connection attempt.
- **Access-Challenge** Sent by a RADIUS server in response to an Access-Request message. This message is a challenge to the RADIUS client that requires a

response. The Access-Challenge message is typically used for challenge-response based authentication protocols to verify the identity of the access client.

- **Access-Accept** Sent by a RADIUS server in response to an Access-Request message. This message informs the RADIUS client that the connection attempt is authenticated and authorized.
- **Access-Reject** Sent by a RADIUS server in response to an Access-Request message. This message informs the RADIUS client that the connection attempt is rejected. A RADIUS server sends this message if the credentials are not authentic or if the connection attempt is not authorized.
- **Accounting-Request** Sent by a RADIUS client to specify accounting information for a connection that was accepted.
- **Accounting-Response** Sent by the RADIUS server in response to the Accounting-Request message. This message acknowledges the successful receipt and processing of the Accounting-Request message.

For PPP authentication protocols such as Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), and Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAP v2), the results of the authentication negotiation between the access server and the access client are forwarded to the RADIUS server for verification in the Access-Request message.

For EAP-based authentication, the negotiation occurs between the RADIUS server and the access client. The RADIUS server uses Access-Challenge messages to send EAP messages to the access client. The access server forwards EAP messages sent by the access client to the RADIUS server as Access-Request messages. Within the Access-Challenge and Access-Request messages, EAP messages are encapsulated as the *RADIUS EAP-Message* attribute.

Authentication, authorization, and accounting of network access connections typically use RADIUS messages in the following way. (See Figure 9-3.)

1. Access servers—such as dial-up network access servers, VPN servers, and wireless APs—receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the RADIUS server.
3. The RADIUS server evaluates the Access-Request message.
4. If required (for example, when the authentication protocol is EAP), the RADIUS server sends an Access-Challenge message to the access server. The response to the challenge is sent as a new Access-Request to the RADIUS server. This can occur multiple times during the EAP negotiation.

5. The RADIUS server verifies the user credentials and the authorization of the connection attempt.
6. If the connection attempt is both authenticated and authorized, the RADIUS server sends an Access-Accept message to the access server. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends an Access-Reject message to the access server.
7. Upon receipt of the Access-Accept message, the access server completes the connection process with the access client and sends an Accounting-Request message to the RADIUS server.
8. After the Accounting-Request message is processed, the RADIUS server sends an Accounting-Response message.

## Planning and Design Considerations

The following sections describe key planning and design considerations for the following technologies in a Windows-based network access authentication infrastructure:

- Active Directory
- PKI
- Group Policy
- RADIUS

### Active Directory

It is beyond the scope of this book to describe in detail the planning and design considerations for deploying Active Directory in an organization of arbitrary size. For detailed information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or resources at <http://www.microsoft.com/ad>.

The following sections describe the planning and design considerations for Active Directory that will help you create a manageable Windows-based authentication infrastructure for network access.

### Accounts and Groups

Depending on the type of connection, network access authentication can use the credentials and properties of user or computer accounts. For each type, you must ensure that the Network Access Permission on the Dial-In tab is set to either Allow Access or Control Access Through NPS Network Policy (recommended). By default, new computer and

user accounts have the Network Access Permission set to Control Access Through NPS Network Policy.

Accounts contain the account name and an encrypted form of the account password that can be used for validation of the client's credentials. Additional account properties determine whether the account is enabled or disabled, locked out, or permitted to log on only during specific hours. If an account is disabled, locked out, or not permitted to log on during the time of the connection, the connection attempt is rejected.

When using groups to manage access, you can use your existing groups and create network policies in NPS that either allow access (with or without restrictions) or reject access based on the group name. For example, you can configure an NPS network policy that specifies the Employees group, which has no network access restrictions for VPN connections. You can also configure another network policy that specifies that the accounts in the Contractors group can create VPN connections only during business hours.

NPS can use Active Directory user principal names (UPNs) and universal groups. In a large domain with thousands of users, create a universal group for all of the users for whom you want to allow access, and then create a network policy that grants access for this universal group. To minimize the processing of group membership for a user account, do not put all of your user accounts directly into the universal group, especially if you have a large number of user accounts. Instead, create separate global groups that are members of the universal group, and add user accounts to those global groups.

## Domain and Forest Trust Relationships

The NPS server is an Active Directory domain member and can verify authentication credentials for accounts in the domain of which it is a member and in all other domains that trust the NPS server's domain. Therefore, ensure that all of the domains in your Active Directory infrastructure trust the domain of the NPS server (subject to security restrictions and policies for your organization); otherwise, you must configure the NPS server as a RADIUS proxy to forward the connection request messages to another NPS server that can authenticate the user or computer account that is attempting to connect.

For the NPS server to be able to access the dial-in properties for user and computer accounts, you must add the computer account of the NPS server to the RAS and IAS Servers group for each domain: the domain of the NPS server and all the domains that trust the NPS server's domain.

## PKI

It is beyond the scope of this book to describe in detail the planning and design considerations for deploying a PKI in an organization of arbitrary size. For detailed information, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/pki>.

A PKI is needed for the following purposes in a Windows-based network access infrastructure:

- Autoenrollment of computer certificates on domain member computers for computer-level certificate-based network access
- Autoenrollment of user certificates on domain member computers for user-level certificate-based network access
- Automatic provisioning of computer health certificates on domain member computers for Internet Protocol security (IPsec) enforcement when deploying NAP.

Subsequent chapters in this book describe additional PKI requirements for different types of network access and for NAP.

The following planning and design considerations for your PKI are specific to a Windows-based authentication infrastructure for network access:

- When using certificates for computer-level network access authentication, configure Group Policy for autoenrollment of computer certificates.  
Examples are the use of EAP-TLS or Protected EAP-TLS (PEAP-TLS) for computer-level wireless authentication.
- When using certificates for user-level network access authentication, configure a certificate template for user certificates, and configure Group Policy for autoenrollment of user certificates.  
Examples are the use of EAP-TLS or PEAP-TLS for user-level wireless authentication.
- When using PEAP-MS-CHAP v2 for network access authentication, configure Group Policy for autoenrollment of computer certificates to install computer certificates on the NPS servers. You can use computer certificates when NPS is not installed on an Active Directory domain controller. Alternatively, you can use the RAS and IAS Server certificate template and configure autoenrollment for members of the RAS and IAS Servers security group.  
Examples are the use of PEAP-MS-CHAP v2 for computer-level or user-level wireless authentication.
- When using IPsec enforcement in NAP, you might need to configure a certificate template for health certificates.
- When using certificates for computer-level or user-level network access authentication, ensure that the CRLs are published in a primary location and in at least one secondary location and that these locations are accessible by all computers, especially the RADIUS servers. The RADIUS servers will first attempt to validate the certificate by using OSCP. If the OSCP validation is not successful, the RADIUS server will attempt to perform a CRL validation of the user or computer certificate. By default, the NPS RADIUS servers

will reject all certificate-based connection attempts if they cannot verify the certificate's revocation status.

### Direct from the Source: Modifying CLR Checking Behavior

Performing CRL checking is enabled by default for security reasons. It is possible to modify the behavior of NPS for certificate revocation checking. There are special cases in which you might want or need to make this change; three examples are as follows:

- If your PKI environment has a poor or slow CRL distribution infrastructure
- If you are using third-party certificates that do not or are not able to provide CRL distribution points with the most up-to-date CRLs
- If you rely on an external distribution point and do not have redundant external connections

Any of these conditions could lead to problems with the certificate revocation checking, thus causing delays or intermittent authentication failure. If you must modify NPS for your deployment, you will be making changes to values in the following registry key:

HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\PPP\EAP\13

The two values you will be most concerned with are:

- **IgnoreNoRevocationCheck** This is set to 0 by default. When set to 1, NPS allows the clients to connect even when it does not perform or cannot complete a revocation check.
- **NoRevocationCheck** This is set to 0 by default. When set to 1, NPS does not attempt a revocation check.

If you set either or both of these registry keys to 1, simply revoking someone's certificate won't limit their network access.

*Chris Irwin, Premier Field Engineer*

*Premier Field Engineering Group*

## Group Policy

It is beyond the scope of this book to describe in detail the planning and design consideration for deploying Group Policy in an organization of arbitrary size. For detailed information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/gp>.

Group Policy is used for the following purposes in a Windows-based network access authentication infrastructure:

- To deploy settings to install a root certificate on domain member computers in order to validate the computer certificates of the NPS servers
- To deploy settings to autoenroll computer certificates on domain member computers for computer-level certificate-based network access authentication
- To deploy settings to autoenroll user certificates on domain member computers for user-level certificate-based network access authentication

Additionally, Group Policy allows you to deploy configuration settings for the following:

- IEEE 802.11 wireless network profiles
- Wired (Ethernet with 802.1X authentication) network profiles
- Windows Firewall with Advanced Security connection security rules to protect traffic
- NAP client configuration

When planning your Group Policy infrastructure, adhere to the recommendations and best practices for Group Policy configuration within your Active Directory infrastructure, as described in the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or in the resources at <http://www.microsoft.com/gp>. There are no specific planning and design considerations for Group Policy objects that are specific to a Windows-based authentication infrastructure for network access and for NAP. However, you must ensure that the correct Group Policy Objects are being applied to those containers or security groups that contain user or computer accounts for authenticated access or for configuration of wireless or wired network profiles, Windows Firewall with Advanced Security connection security rules, or NAP client settings.

## RADIUS

NPS can be used as a RADIUS server, a RADIUS proxy, or both. The following sections describe the planning, design, and security considerations when deploying NPS as a RADIUS server or proxy.

### RADIUS Server Planning and Design Considerations

When planning to deploy an NPS-based RADIUS infrastructure for network access authentication or for NAP, consider the following:

- **Domain membership for NPS servers** You must determine the domain in which to make the NPS server a member. For multiple domain environments, an NPS server can authenticate credentials for user accounts in the domain of which it is a member and all domains that trust its domain. To read the dial-in properties for user and computer

accounts, however, you must add the computer account of the NPS server to the RAS and IAS Servers groups for each domain.

- **UDP ports for RADIUS traffic** If needed, you can configure the NPS server to receive RADIUS messages that are sent to UDP ports other than the default ports of 1812 and 1645 (for RADIUS authentication) and ports 1813 and 1646 (for RADIUS accounting).
- **RADIUS clients to configure on the NPS server** A RADIUS client can be an access server—a network access server (for example, a dial-up or VPN server, a wireless AP, or an Ethernet switch) or a NAP enforcement point—or a RADIUS proxy. NPS supports all access servers and RADIUS proxies that comply with RFC 2865. Configure each access server or RADIUS proxy that sends RADIUS request messages to the NPS server as a RADIUS client on the NPS server.

You can specify IP addresses or DNS names for RADIUS clients. In most cases, it is better to specify IPv4 or IPv6 addresses for RADIUS clients. When you use IP addresses, NPS is not required to resolve host names at startup and will start much more quickly. This is beneficial especially if your network contains a large number of RADIUS clients. Use DNS names to specify RADIUS clients when you require something other than administrative flexibility (for example, the ability to map multiple RADIUS client addresses to a single DNS name).

NPS in Windows Server 2008 allows you to specify a RADIUS client by using an address range. The address range for IPv4-based RADIUS clients is expressed in the network prefix length notation *w.x.y.z/p*, where *w.x.y.z* is the dotted decimal notation of the address prefix, and *p* is the prefix length (the number of high order bits that define the network prefix). This is also known as Classless Inter-Domain Routing (CIDR) notation. An example is 192.168.21.0/24. To convert from subnet mask notation to network prefix length notation, *p* is the number of high order bits in the subnet mask that are set to 1. The address range for IPv6-based RADIUS clients is also expressed in network prefix length notation. An example is 2001:db8:27a1:1c5d::/64.

- **Wireless APs, switches, and third-party remote access servers** To determine whether a third-party access server is interoperable with NPS as a RADIUS server, refer to the third-party access server documentation for its RFC 2865 compliance and its use of RADIUS attributes and vendor-specific attributes.
- **Connection request policy configuration** Connection request policies determine whether the NPS server is used as a RADIUS server, a RADIUS proxy, or both, depending on the information in the incoming RADIUS request messages. The Use Windows Authentication For All Users default connection request policy is configured for NPS when it is used as a RADIUS server. Additional connection request policies can be used to specify more specific conditions, manipulate attributes, and specify advanced attributes. Connection request policies are processed in order, so place the more specific policies at the top of the list. You use the Network Policy Server snap-in to manage new connection request policies.



- **Realm replacement to convert user name formats** The *realm* name is the part of the account name that identifies the location of the user account, such as the name of an Active Directory domain. To correctly replace or convert realm names within the user name of a connection request, configure realm name rules for the User-Name RADIUS attribute on the appropriate connection request policy.
- **Network policy configuration** Network policies are used to grant or deny network access and to set specific conditions for allowed network access, such as dial-in constraints, allowed authentication protocols and encryption strength, and additional RADIUS attributes. Use the Network Policy Server snap-in to manage network policies.
- **Network policies and authorization by user or group** In small organizations, you can manage authorization by setting the network access permission on each user account. For a large organization, set the network access permission on each user account to be controlled through the settings of an NPS network policy. Then, configure network policies to grant access by using group membership.
- **Additional RADIUS attributes and vendor-specific attributes** If you plan to return additional RADIUS attributes or vendor-specific attributes (VSAs) with the responses to RADIUS requests, you must add the RADIUS attributes or VSAs to the appropriate network policy.
- **Event logging** Event logging for authentication events, enabled by default, can assist with troubleshooting connection attempts.
- **Access logging** Access logging stores the authentication and accounting request messages received from access servers and collects this information in a central location. You can store the information in local log files or a Microsoft SQL Server database.
- **Interim accounting** Some access servers send interim accounting messages periodically during a connection, in contrast to the accounting message that is sent when the connection attempt is made. To use interim accounting, first verify that your access server supports sending interim accounting messages. Next, add the Acct-Interim-Interval RADIUS attribute as a standard RADIUS attribute from the Settings tab of the appropriate network policy. Configure the Acct-Interim-Interval attribute with the interval (in minutes) to send periodic interim accounting messages.

## RADIUS Server Security Considerations

When using NPS as a RADIUS server, consider the following to ensure a protected RADIUS infrastructure:

- **RADIUS shared secrets** RADIUS shared secrets are used to verify that RADIUS messages, with the exception of the Access-Request message, are sent by a RADIUS-enabled device that is configured with the same shared secret. Shared secrets also verify that the RADIUS message has not been modified in transit (message integrity). The shared secret is also used to encrypt some sensitive RADIUS attributes, such as User-Password and Tunnel-Password. Configure strong shared secrets and change them frequently to prevent dictionary attacks. Strong shared secrets are a long (more than 22 characters)

sequence of random letters, numbers, and punctuation. You can use the Network Policy Server snap-in to generate strong RADIUS shared secrets.

- **Message Authenticator attribute** To ensure that an incoming RADIUS Access-Request message—for connection requests that use the PAP, CHAP, MS-CHAP, and MS-CHAP v2 authentication protocols—was sent from a RADIUS client configured with the correct shared secret, you can use the RADIUS Message Authenticator attribute (also known as a *digital signature* or the *signature attribute*). You must enable the use of the Message Authenticator attribute on both the NPS server (as part of the configuration of the RADIUS client in the Network Policy Server snap-in) and the RADIUS client (the access server or RADIUS proxy). Ensure that the RADIUS client supports the Message Authenticator attribute before enabling it. The Message Authenticator attribute is always used with EAP-based authentication methods.

For information about enabling the RADIUS Message Authenticator attribute for your access server, see your access server documentation.

- **Firewall configuration for RADIUS traffic** If your NPS server is on a perimeter network, configure your Internet firewall (between your perimeter network and the Internet) to allow RADIUS traffic to pass between your NPS server and RADIUS clients on the Internet. You might need to configure an additional firewall that is placed between your perimeter network and your intranet to allow traffic to flow between the NPS server on the perimeter network and domain controllers on the intranet.
- **Network access authentication protocols** NPS includes support for several different authentication protocols. The order of included authentication protocols, from the most secure to the least secure, is: PEAP-TLS, EAP-TLS, PEAP-MS-CHAP v2, MS-CHAP v2, CHAP, and PAP. Microsoft recommends using only the strongest authentication protocols that are required for your configuration. For password-based authentication protocols, strong password policies must be enforced to protect from dictionary attacks. The use of PAP is not recommended unless it is required.

### Direct from the Source: EAP-MD5 Removed

With the release of Windows Vista, the Microsoft EAP-MD5 implementation has been removed. The decision to remove the Microsoft EAP-MD5 implementation was made in the interest of improving security in Windows Vista. The removal of the Microsoft implementation of EAP-MD5 directly affects remote access services, VPN services, and wired 802.1X deployments. By default, these components can no longer use the Microsoft EAP-MD5 implementation for authentication. The server implementation of EAP-MD5 will continue to ship with Windows Server 2008, but it will be disabled by default. Microsoft will continue to terminate EAP-MD5 connections for legacy network devices but will not initiate them from Microsoft's client operating systems.

*Tim Quinn, Support Escalation Engineer*

*Enterprise Platform Support*

- **Remote access account lockout** To provide protection for online dictionary attacks launched against access servers by using known user names, you can enable remote access account lockout. Remote access account lockout disables remote access for user accounts after a configured number of failed connection attempts has been reached. For more information, see Chapter 12, “Remote Access VPN Connections.”

Remote access account lockout can also be used to prevent a malicious user from intentionally locking out a domain account by attempting multiple dial-up or VPN connections with the wrong password. You can set the number of failed attempts for remote access account lockout to a number that is lower than the logon retries for domain account lockout. By doing this, remote access account lockout occurs before domain account lockout, which prevents the domain account from being intentionally locked out.

- **Certificates to install on NPS servers for network access authentication** When you use the included EAP-TLS, PEAP-TLS, or PEAP-MS-CHAP v2 authentication protocols, by default you must install a computer certificate on the NPS server containing the Server Authentication purpose in the Enhanced Key Usage (EKU) extensions. Other authentication protocols provided by independent software or hardware vendors might also require certificates on NPS servers.
- **Using Windows Firewall with Advanced Security connection security rules to protect NPS servers** You can configure Windows Firewall with Advanced Security connection security rules to protect RADIUS traffic sent between RADIUS servers and access servers and between RADIUS servers and RADIUS proxies with IPsec. These rules can be configured as part of Group Policy settings and applied to Active Directory containers or filtered for security groups, or they can be created and applied to individual servers.

## RADIUS Proxy Planning and Design Considerations

When planning to deploy a RADIUS infrastructure for network access authentication or for NAP, consider the following:

- **When to use NPS as a RADIUS proxy** The following uses of NPS as a RADIUS proxy are described in this chapter:
  - When you want to provide authentication and authorization for user accounts that are not members of either the domain in which the NPS server is a member or another domain that has a two-way trust with the domain in which the NPS server is a member. This includes accounts in untrusted domains, one-way trusted domains, and other forests. Instead of configuring your access servers to send their connection requests to an NPS RADIUS server, you can configure them to send their connection requests to an NPS RADIUS proxy. The NPS RADIUS proxy uses the realm name portion of the user name to forward the request to an NPS server in the correct domain or forest. Connection attempts for user accounts in one domain or forest can be authenticated for network access servers that are members of another domain or forest.

- ❑ When you want to process a large number of connection requests. In this case, instead of configuring your RADIUS clients to attempt to balance their connection and accounting requests across multiple RADIUS servers, you can configure them to send their connection and accounting requests to an NPS RADIUS proxy. The NPS RADIUS proxy dynamically balances the load of connection and accounting requests across multiple RADIUS servers and increases the processing of large numbers of RADIUS clients and authentications per second.

For more information about these configurations, see “Using RADIUS Proxies for Cross-Forest Authentication” and “Using RADIUS Proxies to Scale Authentications” later in this chapter.

- **Connection request policy configuration** The Use Windows Authentication For All Users default connection request policy uses NPS as a RADIUS server. To create a connection request policy to use NPS as a RADIUS proxy, you must first create a remote RADIUS server group whose members are the set of RADIUS servers to which a RADIUS message is forwarded. Next, create a connection request policy that forwards authentication requests to a remote RADIUS server group. Finally, either delete the Use Windows Authentication For All Users connection request policy or move the new connection request policy higher in the list so that it is evaluated first.
- **Realm replacement and attribute manipulation** To convert realm names and configure RADIUS message forwarding based on the realm name, you must use realm rules for the User-Name attribute on the appropriate connection request policy. If you are using the MS-CHAP v2 authentication protocol, you cannot manipulate the User Name attribute if the connection request policy is used to forward the RADIUS message. The only exception occurs when a backslash character (\) is used, and the manipulation affects only the information to the left of it. A backslash character is typically used to indicate a domain name (the information to the left of the backslash) and a user account name within the domain (the information to the right of the backslash). In this case, only attribute manipulation rules that modify or replace the domain name are allowed.
- **The use of additional RADIUS attributes and vendor-specific attributes** If you plan to include additional RADIUS attributes and vendor-specific attributes (VSAs) to RADIUS requests that are being forwarded, you must add the RADIUS attributes and VSAs to the appropriate connection request policy.
- **Remote RADIUS server group configuration** A remote RADIUS server group contains the set of RADIUS servers to which RADIUS messages matching a connection request policy are forwarded.
- **Copying logging information at the NPS proxy** The NPS proxy can record all RADIUS accounting information that it receives in the local log file. This creates a central location for all authentication and accounting information for all of the access servers of the NPS proxy.

- **Authentication and accounting ports** When you configure a server in a remote RADIUS server group, you can configure custom UDP ports to which RADIUS authentication and accounting messages are sent. The default UDP port for authentication requests is 1812. The default UDP port for accounting requests is 1813.
- **Load balancing and failure detection** When you configure multiple servers in a remote RADIUS server group, you can configure settings that determine how the NPS proxy balances the load of authentication and accounting requests over the RADIUS servers in the group. By default, the RADIUS traffic is balanced equally across the members of the group. You can use additional settings to configure NPS to detect and recover from the failure of a remote RADIUS server group member.

### Direct from the Source: RADIUS Proxies and Trusts

It is best to avoid creating arbitrary trusts for cross-domain network authentication. If your goal is to allow domain users the ability to log on to networks in different domains, use RADIUS proxies rather than a transitive trust. With a RADIUS proxy, you are passing only the essential data between the two NPS servers necessary for granting user or computer access. Additionally, this requires at most only two UDP ports to be available between the two domains. With a trust, far more traffic, such as resource access validation, is being passed, and many more ports are required to be opened.

*Clay Seymour, Support Escalation Engineer*

*Enterprise Platform Support*

## RADIUS Proxy Security Considerations

When using NPS as a RADIUS proxy, consider the following to ensure a protected RADIUS infrastructure:

- **Shared secrets** Configure strong shared secrets to prevent dictionary attacks, and change them frequently. Strong shared secrets are a long (more than 22 characters) sequence of random letters, numbers, and punctuation.
- **Firewall configuration** If your NPS proxy is on a perimeter network, configure your Internet firewall (between your perimeter network and the Internet) to allow RADIUS messages to pass between your NPS proxy and RADIUS clients on the Internet. You might need to configure an additional firewall that is placed between your perimeter network and your intranet to allow RADIUS traffic to flow between the NPS proxy on the perimeter network and an NPS server on the intranet.
- **Message Authenticator attribute** You can use the RADIUS Message Authenticator attribute (also known as a *digital signature* or the *signature attribute*) to ensure that RADIUS Access-Request messages for connection requests were sent from a RADIUS

client configured with the correct shared secret. The Message Authenticator attribute is always used with EAP, and you don't have to enable it on the NPS server or access server. For the PAP, CHAP, MS-CHAP, and MS-CHAP v2 authentication protocols, you must enable the use of the Message Authenticator attribute on both the NPS server (as part of the configuration of the RADIUS client) and the RADIUS client (the access server or RADIUS proxy). Ensure that the RADIUS client supports the Message Authenticator attribute before enabling it.

- **Using Windows Firewall with Advanced Security connection security rules to protect NPS proxies** You can configure the Windows Firewall with Advanced Security connection security rules to use IPsec to protect RADIUS traffic sent between NPS proxies and access servers and between the NPS proxies and RADIUS servers.
- **Password Authentication Protocol (PAP)** The use of PAP is strongly discouraged, especially when using RADIUS proxies.

## High Availability for RADIUS Authentication

To provide high availability for RADIUS-based authentication and accounting, you should always use at least two NPS servers. One NPS server is used as the primary RADIUS server, and the other is used as a backup. Access servers or other RADIUS proxies are configured for both NPS servers (a primary and a secondary) and automatically switch to the secondary NPS RADIUS server when the primary NPS RADIUS server becomes unavailable. When using multiple RADIUS servers, failover is based on a RADIUS client switching to another RADIUS server and performing a new authentication transaction. Failover within a transaction is not supported.

## High Scalability for RADIUS Authentication

Consider the following for scaling RADIUS authentication to an organization containing a large number of accounts or connection attempt activity:

- **Use universal groups and group-based network policies** If you are using network policies to restrict access for all but certain groups, create a universal group for all of the users or computers for whom you want to allow access, and then create a network policy that grants access for this universal group. Do not put all of your user and computer accounts directly into the universal group, especially if you have a large number of them on your network. Instead, create separate groups that are members of the universal group, and add the user and computer accounts to those groups.
- **Use user principal names** Use user principal names (UPNs), such as user@contoso.com, to refer to users whenever possible. A user can have the same user principal name regardless of domain membership. This practice provides scalability that might be required in organizations with a large number of domains.

- **Install NPS on domain controllers** If possible, install NPS on domain controllers for best authentication and authorization performance. When NPS is running on a domain controller, the traffic and processing delays incurred when an NPS RADIUS server contacts a domain controller over the network to verify account credentials and obtain account properties are eliminated.

If the NPS server is on a computer other than a domain controller, and it is receiving a large number of authentication requests per second, you can improve performance by increasing the number of concurrent authentications between the NPS server and the domain controller. To do this, edit the following registry key:

HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters. Add a new value (REG\_DWORD value type) named MaxConcurrentApi, and although the range can be between 0 and 10, assign it a setting from 2 through 5.

This value specifies the maximum number of simultaneous logon calls that can be transmitted to the domain controller over the secure channel at any given time, and the default is 2 for a member server computer. Increasing the setting will allow additional logon calls to be processed simultaneously to improve performance on the NPS server. Avoid setting the MaxConcurrentApi value to a setting higher than 5 because the additional load might cause depletion of resources on the domain controller.

## Deployment Steps

This section contains the steps or resources for the steps to deploy the following components of a Windows-based network access authentication infrastructure:

- Active Directory
- PKI
- Group Policy
- RADIUS

## Deploying Active Directory

It is beyond the scope of this book to instruct you on the specific steps to deploy Active Directory for an organization of arbitrary size. For additional information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/ad>.

The elements of configuring Active Directory to best support a Windows-based authentication infrastructure for network access are as follows:

- Ensure that all users who are making user-level authenticated connections have a corresponding user account that is enabled.

- Ensure that all computers that are making computer-level authenticated connections have a corresponding computer account that is enabled.
- Set the network access permission on user and computer accounts to the appropriate setting: either Allow Access or Control Access Through NPS Network Policy (recommended). The network access permission setting is on the Dial-In tab on the properties dialog box of a user or computer account in the Active Directory Users And Computers snap-in.
- Organize your network access user and computer accounts into the appropriate groups. Use a Windows 2000, Windows Server 2003, or Windows Server 2008 functional-level domain and universal groups and global groups to organize your accounts for a specific type of access into a single group. For example, for wireless access, create a universal group named WirelessUsers that contains global groups of wireless user and computer accounts for intranet access.

## Deploying PKI

It is beyond the scope of this book to provide the specific steps to deploy a PKI for an organization of arbitrary size. For additional information, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/pki>.

The elements of configuring a certificate services-based PKI to best support a Windows-based authentication infrastructure for network access are as follows:

- When using certificates for user-level network access authentication, configure a certificate template for user certificates. If you are running a Windows enterprise CA, you can make a copy of the standard user template. Standalone CAs do not support certificate templates.
- When using IPsec enforcement in NAP, you might need to configure a certificate template for health certificates.



**Note** A certificate template for a computer certificate is already configured by default with Windows Certificate Services.

After your PKI has been deployed, there are a set of procedures for deploying certificates that are common to wireless, wired, remote access VPN, and site-to-site VPN connections. These procedures are as follows:

- Configuring autoenrollment of computer certificates to computers in an Active Directory domain
- Using the Certificates snap-in to request a computer certificate



- Using the Certificates snap-in to import a computer certificate
- Executing a CAPICOM script that requests a computer or user certificate
- Configuring autoenrollment of user certificates in an Active Directory domain
- Using the Certificates snap-in to request a user certificate
- Using the Certificates snap-in to import a user certificate
- Installing third-party certificate chains by using Group Policy
- Requesting a certificate via the Web

**Configuring the Autoenrollment of Computer Certificates to Computers in an Active Directory Domain** If you are using a Windows Server 2008 enterprise CA as an issuing CA, each computer can automatically request a computer certificate from the issuing CA by using a Computer Configuration group policy setting. This method allows a single point of configuration for an entire domain.

### **To Configure an Active Directory Domain for Automatic Enrollment of Computer Certificates**

1. Open the Group Policy Management snap-in.
2. In the console tree, expand Forest, expand Domains, and then click the name of the domain to which your CA belongs.
3. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
4. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, then Windows Settings, then Security Settings, and then Public Key Policies.
5. Right-click Automatic Certificate Request Settings, point to New, and then click Automatic Certificate Request.
6. The Automatic Certificate Request Setup Wizard appears. Click Next.
7. On the Certificate Template page, click Computer, and then click Next.
8. Click Finish.

To immediately obtain an updated Computer Configuration Group Policy to request a computer certificate for a computer running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP, restart the computer, or type **gpupdate /target:computer** at a command prompt.

**Using the Certificates Snap-In to Request a Computer Certificate** If you are using a Windows Server 2008 enterprise CA as an issuing CA, each computer can separately request a computer certificate from the issuing CA by using the Certificates snap-in.

### To Request a Computer Certificate by Using the Certificates Snap-In

1. Log on to the computer using an account that has administrator privileges for that computer.
2. On the Start menu, click Run, type **mmc**, and then press Enter.
3. On the Console menu, click File, and then click Add/Remove Snap-In.
4. In the Add Or Remove Snap-Ins dialog box, under Available Snap-Ins, double-click Certificates. In the Certificates Snap-In dialog box, click Computer Account, and then click Next.
5. Do one of the following:
  - ❑ To manage certificates for the local computer, click Local Computer.
  - ❑ To manage certificates for a remote computer, click Another Computer and type the name of the computer, or click Browse to select the computer name. Then click OK.
6. Click Finish. Certificates (Local Computer) or Certificates (*computername*) appears on the list of selected snap-ins for the new console. Click OK.
7. In the console tree, expand the Certificates\Personal node.
8. Right-click the Personal node, point to All Tasks, and then click Request New Certificate.

The Certificate Request Wizard guides you through the steps of requesting a certificate. For a Windows-based client computer, the certificate imported into the Local Computer store must have the Client Authentication Enhanced Key Usage (EKU). For the certificate installed on the VPN server or the NPS server, the certificate imported into the Local Computer store must have the Server Authentication EKU.

**Using the Certificates Snap-In to Import a Computer Certificate** If you have a certificate file that contains the computer certificate, you can import the computer certificate by using the Certificates snap-in. This must be done when you purchase individual computer certificates for your VPN or RADIUS servers from a third-party CA for PEAP-MS-CHAP v2 authentication or for Secure Socket Tunneling Protocol (SSTP) connections.

### To Import a Computer Certificate by Using the Certificates Snap-In

1. Open the Certificates (Local Computer)\Personal node.
2. Right-click the Personal node, point to All Tasks, and then click Import.

The Certificate Import Wizard guides you through the steps of importing a certificate from a certificate file. For a Windows-based client computer, the certificate imported into the Local Computer store must have the Client Authentication EKU. For the certificate installed on the VPN or NPS server, the certificate imported into the Local Computer store must have the Server Authentication EKU.



**Note** It is also possible to import a certificate by double-clicking a certificate file that is stored in a folder or sent in an e-mail message. Although this works for certificates created with Windows-based CAs, this method might not work for third-party CAs. The recommended method of importing certificates is to use the Certificates snap-in.

**Executing a CAPICOM Script That Requests a Computer or User Certificate** In this method, each computer must execute a CAPICOM script that requests a computer or user certificate from the issuing CA. CAPICOM is a COM client that performs cryptographic functions (the CryptoAPI) by using Microsoft ActiveX and COM objects. CAPICOM can be used with Microsoft Visual Basic, Visual Basic Scripting Edition, and C++. For more information about CAPICOM, visit <http://msdn2.microsoft.com/en-us/library/ms995332.aspx>.

To perform an enterprise deployment of user and computer certificates, a CAPICOM program or script can be distributed through e-mail for execution, or users can be directed to a Web site containing a link to a CAPICOM program or script. Alternately, the CAPICOM program or script can be placed in the user's logon script file for automatic execution. The storage location of the user or computer certificate can be specified using the CAPICOM application programming interfaces (APIs).

### Configuring Autoenrollment of User Certificates to Users in an Active Directory

**Domain** This method allows a single point of configuration for the entire domain. All members of the domain automatically request the user certificate through a User Configuration group policy setting. If you use as an issuing CA an enterprise CA from Windows Server 2008, Windows Server 2003 Enterprise Edition, or Windows Server 2003 Datacenter Edition, you can install user certificates through autoenrollment.

### To Configure User Certificate Enrollment for an Enterprise CA

1. On the Start menu, click Run, type **mmc**, and then click OK.
2. On the File menu, click Add/Remove Snap-In.
3. Under Available Snap-Ins, double-click Certificate Templates, and then click OK.
4. In the console tree, click Certificate Templates. All certificate templates appear in the details pane.
5. In the details pane, right-click the User template, and then click Duplicate Template. When prompted for the minimum version of the CA to support the certificate template, click Windows Server 2003, Enterprise Edition, and then click OK.
6. In the Template Display Name field, type the name of the new user certificate template (for example, **VPNAccess**).

Make sure that the Publish Certificate In Active Directory check box is selected.

7. Click the Security tab.

8. In the Group Or User Names list, click Domain Users.
9. In the Permissions For Domain Users list, select the Read, Enroll, and Autoenroll permission check boxes, and then click OK.
10. Open the Certification Authority snap-in.
11. In the console tree, expand your CA's name, and then click Certificate Templates.
12. On the Action menu, point to New, and then click Certificate Template To Issue.
13. Click the name of the newly created user certificate template (for example, VPNAccess), and then click OK.
14. Open the Group Policy Management snap-in.
15. In the console tree, expand Forest, expand Domains, and then click the name of your domain to which your CA belongs.
16. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
17. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, then Windows Settings, then Security Settings, and then Public Key Policies.
18. In the details pane, double-click Certificate Services Client – Auto-Enrollment.
19. In Configuration Model, select Enabled from the drop-down list.
20. Select the Renew Expired Certificates, Update Pending Certificates, and Remove Revoked Certificates check box.
21. Select the Update Certificates That Use Certificate Templates check box, and then click OK.

Perform steps 15–21 for each domain container, as appropriate. Ensure that all appropriate domain containers are configured for autoenrollment of user certificates, either through the inheritance of group policy settings of a parent container or through explicit configuration.

To immediately update User Configuration group policy and request a user certificate for a computer that is running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP and is a member of the domain for which autoenrollment is configured, restart the computer, or at a command prompt, type **gpupdate /target:user**.

**Using the Certificates Snap-In to Request a User Certificate** If you are using a Windows Server 2008 enterprise CA as an issuing CA, each computer can separately request a user certificate from the issuing CA by using the Certificates snap-in.

#### **To Request a User Certificate by Using the Certificates Snap-In**

1. Log on to the computer using an account that has administrator privileges for that computer.

2. On the Start menu, click Run, type **mmc**, and then press Enter.
3. On the Console menu, click File, and then click Add/Remove Snap-In.
4. In the Add Or Remove Snap-Ins dialog box, under Available Snap-Ins, double-click Certificates. In the Certificates Snap-In dialog box, click My User Account, click Finish, and then click OK.
5. In the console tree, expand the Certificates\Personal node.
6. Right-click the Personal node, point to All Tasks, and then click Request New Certificate.

The Certificate Request Wizard guides you through the steps of requesting a user certificate. For a Windows-based client computer, the imported certificate must have the Client Authentication EKU.

**Using the Certificates Snap-In to Import a User Certificate** If you have a certificate file that contains the user certificate, you can import the user certificate by using the Certificates snap-in.

#### **To Import a User Certificate by Using the Certificates Snap-In**

1. Open the Certificates (Current User)\Personal node.
2. Right-click the Personal node, point to All Tasks, and then click Import.

The Certificate Import Wizard guides you through the steps of importing a certificate from a certificate file. For a Windows-based client computer, the certificate imported into the Local Computer store must have the Client Authentication EKU.

**Installing Third-Party Certificate Chains by Using Group Policy** When you are using a third-party CA for the computer certificates that are installed on access servers or RADIUS servers, you might need to install the chain of certificates (the root CA certificate to the issuing CA certificate) for the certificate installed on the access or RADIUS server. If the access client does not trust the certificate chain of the certificate submitted by the access or RADIUS server, certificate validation can fail.

A certificate chain consists of the root CA certificate and the certificate of each intermediate CA, including the issuing CA. The following procedures describe how to deploy a root CA certificate and an intermediate CA certificate to access clients by using Group Policy.

#### **To Install a Root CA Certificate by Using Group Policy**

1. In the console tree of the Certificates snap-in for the access or RADIUS server computer account, expand Certificates (Local Computer), expand Trusted Root Certification Authorities, and then click Certificates.
2. In the details pane, right-click the root CA certificate of the issuing CA of the computer certificate on the authentication server, point to All Tasks, and then click Export.
3. In the Certificate Export Wizard, on the Welcome to the Certificate Export Wizard page, click Next.

4. On the Export File Format page, click Cryptographic Message Syntax Standard–PKCS #7 Certificates (.p7b).
5. Click Next. On the File To Export page, type the file name for the exported certificate, or click Browse to specify a location and file name.
6. Click Next. On the Completing The Certificate Export Wizard page, click Finish.
7. Open the Group Policy Management snap-in.
8. In the console tree, expand Forest, expand Domains, and then click the name of your domain to which your CA belongs.
9. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
10. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, Windows Settings, Security Settings, and then Public Key Policies.
11. Right-click Trusted Root Certification Authorities, and then click Import.
12. In the Certificate Import Wizard, specify the file that was saved in step 5.
13. Repeat steps 8–12 for all appropriate domain containers and their Group Policy Objects.

The next time the access client computers update their Computer Configuration group policy, the root CA certificates of the issuing CAs of the authentication server computer certificates are installed in their local certificate store.

### **To Install an Intermediate CA Certificate by Using Group Policy**

1. In the console tree of the Certificates snap-in for the access or RADIUS server computer account, expand Certificates (Local Computer), expand Intermediate Certification Authorities, and then click Certificates.
2. In the details pane, right-click the intermediate CA certificate of the issuing CA of the computer certificate on the authentication server, point to All Tasks, and then click Export.
3. In the Certificate Export Wizard, on the Welcome To The Certificate Export Wizard page, click Next.
4. On the Export File Format page, click Cryptographic Message Syntax Standard–PKCS #7 Certificates (.p7b).
5. Click Next. On the File To Export page, type the file name for the exported certificate, or click Browse to specify a location and file name.
6. Click Next. On the Completing The Certificate Export Wizard page, click Finish.
7. Open the Group Policy Management snap-in.
8. In the console tree, expand Forest, expand Domains, and then click the name of your domain to which your CA belongs.

9. On the Linked Group Policy Objects pane, right-click the appropriate Group Policy Object (the default object is Default Domain Policy), and then click Edit.
10. In the console tree of the Group Policy Management Editor snap-in, expand Computer Configuration, Windows Settings, Security Settings, and then Public Key Policies.
11. Right-click Intermediate Certification Authorities, point to All Tasks, and then click Import.
12. In the Certificate Import Wizard, specify the file that was saved in step 5.

Repeat steps 8–12 for all appropriate domain containers and their Group Policy Objects.

If you cannot use Group Policy, you can manually install root and intermediate certificates on individual access client computers.

### To Manually Install a Root or Intermediate CA Certificate on an Access Client

1. Export the root CA certificate of the access or RADIUS server's computer certificate to a .p7b file.
2. On the access client computer, in the console tree of the Certificates (Local Computer) snap-in, expand Certificates (Local Computer), expand Trusted Root Certification Authorities (for a root CA certificate) or Intermediate Certification Authorities (for an intermediate CA certificate), and then click Certificates.
3. Right-click Certificates, point to All Tasks, and then click Import.
4. The Welcome To The Certificate Import Wizard page of the Certificate Import Wizard appears. Click Next.
5. On the File To Import page, in the File Name box, type the file name of the certificate file saved in step 1, or click Browse and use the Browse dialog box to locate it.
6. Click Next. On the Certificate Store page, click Place All Certificates In The Following Store, and then specify the import location.
7. Click Next. On the Completing The Certificate Import Wizard page, click Finish.

**Requesting a Certificate via the Web** Requesting a certificate via the Web, also known as Web enrollment, is done with Microsoft Windows Internet Explorer. For the address, type **http://servername/certsrv**, where *servername* is the computer name of the Windows Server 2008 or Windows Server 2003 CA that is also running Internet Information Services (IIS). A Web-based wizard takes you through the steps of requesting a certificate. The location where the certificate is stored (whether it is the Current User store or the Local Computer store) is determined by whether the Use Local Machine Store check box was selected when an advanced certificate request was performed. This check box is cleared by default, and certificates are stored in the Current User store. You must have local administrator privileges to store a certificate in the Local Computer store.

You can use Web enrollment with either an enterprise or a Standalone CA.

### **Direct from the Source: Duplicating Default Certificate Templates**

When using certificate templates, you should always make a duplicate of the default template, and if applicable, make your scenario-specific changes to the new template. For example, if you want to change the security groups that can autoenroll for a user certificate, make a duplicate of the user certificate. Then, obtain the properties of the new certificate template, click the Security tab, and add the specific groups that you want to have access to the template.

*Clay Seymour, Support Escalation Engineer*

*Enterprise Platform Support*

## **Group Policy**

It is beyond the scope of this book to provide the specific steps to deploy Group Policy for an organization of arbitrary size. For additional information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or resources at <http://www.microsoft.com/gp>.

The elements of configuring Group Policy to best support a Windows-based authentication infrastructure for network access are as follows:

- When using certificates for computer-level network access authentication, configure Group Policy for autoenrollment of computer certificates. This requires deployment of a Windows enterprise CA. Autoenrollment cannot be configured when using a Stand-alone CA.
- When using certificates for user-level network access authentication, configure a certificate template for user certificates, and configure Group Policy for autoenrollment of user certificates.
- When using PEAP-MS-CHAP v2 for network access authentication, optionally configure Group Policy for autoenrollment of computer certificates to install computer certificates on the NPS servers.
- When you are using PEAP-MS-CHAP v2 for network access authentication and a third-party CA for the computer certificates installed on the NPS RADIUS servers, ensure that the root CA certificate for the NAP RADIUS server's computer certificate is installed on the access clients. If not, configure Group Policy to install the appropriate root CA certificate on domain member computers.



For information about how to configure Group Policy to deploy certificate settings, see “Deploying PKI” earlier in this chapter.

For information about how to configure Group Policy to deploy configuration settings for specific types of network access, see the following:

- Chapter 10, “IEEE 802.11 Wireless Networks”
- Chapter 11, “IEEE 802.1X-Authenticated Wired Networks”
- Chapter 16, “IPsec Enforcement”
- Chapter 17, “802.1X Enforcement”
- Chapter 18, “VPN Enforcement”
- Chapter 19, “DHCP Enforcement”

## RADIUS Servers

Configuring a fault-tolerant RADIUS infrastructure requires at a minimum the configuration of at least two NPS RADIUS servers, a primary NPS RADIUS server, and a secondary RADIUS NPS server. You must do the following:

- Configure the primary NPS server.
- Copy the configuration of the primary NPS server to the secondary NPS server.

Because the configuration of the primary NPS server is being copied to the secondary NPS server, you should always make configuration changes to the primary NPS server.

### Configuring the Primary NPS Server

To configure the primary NPS server on a computer, complete these steps as discussed in the following sections:

1. Obtain and install a computer certificate.
2. Install NPS and configure NPS server properties.
3. Configure NPS with RADIUS clients.
4. Use IPsec to protect RADIUS traffic.
5. Configure the appropriate policies.

**Obtaining and Installing a Computer Certificate** If you have configured computer certificate autoenrollment, force a refresh of computer configuration Group Policy by typing **gpupdate /target:computer** at a command prompt.

If you use a Windows Server 2008 or Windows Server 2003 enterprise CA and you are not using autoenrollment for computer certificates, you can request one, as described in the following procedure.

### To Request a Computer Certificate

1. Click Start, click Run, type **mmc**, and then click OK.
2. On the File menu, click Add/Remove Snap-In.
3. Under Available Snap-Ins, double-click Certificates, click Computer Account, and then click Next.
4. Do one of the following:
  - ❑ To manage certificates for the local computer, click Local Computer, and then click Finish.
  - ❑ To manage certificates for a remote computer, click Another Computer and type the name of the computer, or click Browse to select the computer name. Click Finish.
5. Click OK.
6. In the console tree, expand Certificates (Local Computer or *Computername*), and then click Personal.
7. On the Action menu, point to All Tasks, and then click Request New Certificate to start the Certificate Enrollment Wizard.
8. On the Before You Begin page, click Next.
9. On the Request Certificates page, click Computer, and then click Enroll.
10. Click Finish.

If your PKI does not support autoenrollment of computer certificates, obtain the computer certificate as a saved file, and then use the following procedure to import the computer certificate on the primary NPS server.



**Note** To perform the next procedure, you must be a member of the Administrators group on the local computer, or you must have been delegated the appropriate authority.

### To Import the Computer Certificate on the Primary NPS Server

1. In the console tree of the Certificates snap-in, expand Certificates (Local Computer or *Computername*).
2. Right-click Personal, point to All Tasks, and then click Import.
3. On the Welcome To The Certificate Import Wizard page, click Next.
4. On the File To Import page, in the File Name box, type the file name of the certificate file provided by the commercial CA. Alternatively, you can click Browse and use the Browse dialog box to locate it.

5. Click Next. On the Certificate Store page, click Place All Certificates In The Following Store. By default, the Personal node should appear as the import location. Click Next, and then click Finish.

**Configuring NPS Server Properties** NPS is installed on computers running Windows Server 2008 with the Network Policy and Access Services role through the Initial Configuration Tasks or Server Manager tools. However, the primary NPS server computer must be able to access account properties in the appropriate domains. If NPS is being installed on a domain controller, no additional configuration is required for NPS to access account properties in the domain to which it belongs. If NPS is not installed on a domain controller, you must configure the primary NPS server computer to read the properties of user accounts in the domain, as described in the following procedure:

#### **To Configure the Primary NPS Server Computer to Read the Properties of User Accounts in the Domain**

1. In the console tree of the Network Policy Server snap-in, right-click NPS (Local), and then click Register Server In Active Directory.
2. In the Network Policy Server dialog box, click OK twice.

Alternatively, you can do one of the following:

- Use the **netsh nps add registeredserver** command.
- Use the Active Directory Users And Computers snap-in to add the computer account of the NPS server to the RAS and IAS Servers security group.

If the NPS server authenticates and authorizes network access attempts for user accounts in other domains, verify that the other domains have a two-way trust with the domain in which the NPS server computer is a member. Next, configure the NPS server computer to read the properties of user accounts in other domains by using the **netsh nps add registeredserver** command or by using the Active Directory Users And Computers snap-in.

If there are accounts in other domains, and the domains do not have a two-way trust with the domain in which the NPS server computer is a member, you must configure a RADIUS proxy between the two untrusted domains. If there are accounts in other untrusted Active Directory forests, you must configure a RADIUS proxy between the forests. For more information, see “Using RADIUS Proxies for Cross-Forest Authentication” later in this chapter.

If you want to store authentication and accounting information for connection analysis and security investigation purposes, enable logging for accounting and authentication events. Windows Server 2008 NPS can log information to a local file and to a SQL Server database.

#### **To Enable and Configure Local File Logging for NPS**

1. In the console tree of the Network Policy Server snap-in, click Accounting.
2. In the details pane, click Configure Local File Logging.

3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests in the NPS log files:
  - ❑ To capture accounting requests and responses, select the Accounting Requests check box.
  - ❑ To capture authentication requests, access-accept packets, and access-reject packets, select the Authentication Requests check box.
  - ❑ To capture periodic status updates, such as interim accounting packets, select the Periodic Accounting Status or Periodic Authentication Status check boxes.

All these logging options are enabled by default.

4. On the Log File tab, type the log file directory as needed, and then select the log file format and new log time period. The default log file directory is %SystemRoot%\System32\LogFiles.

### To Enable and Configure SQL Server Database Logging for NPS

1. In the console tree of the Network Policy Server snap-in, click Accounting.
2. In the details pane, click Configure SQL Server Logging.
3. On the Settings tab, select one or more check boxes for recording authentication and accounting requests. All these logging options are enabled by default.
4. In Maximum Number of Concurrent Sessions, type the maximum number of simultaneous sessions that NPS can create with SQL Server.
5. To configure a SQL data source, click Configure.
6. In the Data Link Properties dialog box, configure the appropriate settings for the SQL Server database.

If needed, configure additional UDP ports for authentication and accounting messages that are sent by RADIUS clients (the access servers). By default, NPS uses UDP ports 1812 and 1645 for authentication messages and UDP ports 1813 and 1646 for accounting messages.

### To Configure NPS for Different UDP Ports

1. In the console tree of the Network Policy Server snap-in, right-click NPS, and then click Properties.
2. Click the Ports tab, and then in the Authentication section, type the UDP port numbers for your RADIUS authentication traffic. In the Accounting section, type the UDP port numbers for your RADIUS accounting traffic.

To use multiple port settings for authentication or accounting traffic, separate the port numbers with commas. You can also specify an IP address to which the RADIUS messages must be sent by typing in the following syntax: **IPAddress:UDPPort**. For example, if you have multiple network adapters and you want to receive RADIUS authentication

messages sent only to the IP address of 10.0.0.99 and UDP port 1812, in the Authentication box, type **10.0.0.99:1812**. However, if you specify IP addresses and copy the configuration of the primary NPS server to the secondary NPS server, you must modify the ports on the secondary NPS server to either remove the IP address of the primary NPS server or change the IP address to that of the secondary NPS server.

**Configuring NPS with RADIUS Clients** You must configure the primary NPS server with the access servers or RADIUS proxies as RADIUS clients.

#### To Add a RADIUS Client for NPS

1. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients And Servers, right-click RADIUS Clients, and then click New RADIUS Client.
2. In the New RADIUS Client dialog box, under Name And Address, in the Friendly Name text box, type a name for the RADIUS client (the access server or RADIUS proxy). In the Address (IP Or DNS) text box, type the IP address or DNS domain name of the RADIUS client. If you type a DNS domain name, click Verify to resolve the name to the correct IP address for the access server.
3. Under Shared Secret, in the Shared Secret and Confirm Shared Secret text boxes, type the shared secret for this combination of NPS server and RADIUS client or click Generate to have the NPS service generate a strong RADIUS shared secret.
4. Under Additional Options, specify whether this RADIUS client will always use the Message-Authenticator attribute in RADIUS messages and whether the RADIUS client is a NAP enforcement point that is running Windows Server 2008 (the RADIUS Client Is NAP-Capable check box), and then click OK.

If you have multiple wireless APs on a single subnet, you can simplify RADIUS client administration by specifying an IPv4 or IPv6 address range instead of specifying the address or DNS name of a single RADIUS client. All of the RADIUS clients in the range must be configured to use the same RADIUS server and shared secret. If you are not using this feature, use a different shared secret for each wireless AP.

Use as many RADIUS shared secrets as you can. Each shared secret should be a random sequence of uppercase and lowercase letters, numbers, and punctuation marks that is at least 22 characters long. To create a strong RADIUS shared secret, use the Generate option when configuring a shared secret with the Network Policy Server snap-in.

**Using IPsec to Protect RADIUS Traffic** To ensure maximum security for RADIUS messages, it is recommended that you use IPsec and Encapsulating Security Payload (ESP) to provide data confidentiality, data integrity, and data origin authentication for RADIUS traffic sent between the NPS servers and the RADIUS clients. Computers running Windows Server 2008 and Windows Server 2003 support IPsec. You configure the NPS RADIUS server for IPsec protection of RADIUS traffic through Windows Firewall with Advanced Security connection security rules. To secure RADIUS traffic sent from third-party access servers, the access

servers must also support IPsec. For more information about connection security rules, see Chapter 4, “Windows Firewall with Advanced Security.”

**Configuring the Appropriate Policies** To evaluate authorization and connection constraints for incoming connection requests, you must configure the appropriate policies consisting of connection request policies, network policies, and for NAP, health policies. The Network Policy Server snap-in has a set of wizards to automatically configure a set of policies for common network access and NAP scenarios. The following procedure describes how to run the Network Policy Server wizards.

### To Run the Network Policy Server Wizards

1. In the console tree of the Network Policy Server snap-in, click NPS (Local).
2. In the details pane, in the drop-down list select one of the following:
  - ☐ Network Access Protection (NAP)
  - ☐ RADIUS Server For Dial-up Or VPN Connections
  - ☐ RADIUS Server For 802.1X Wireless or Wired Connections
3. If you selected Network Access Protection (NAP), click Configure NAP and use the pages of the Configure NAP Wizard to specify the set of policies for NAP enforcement.
4. If you selected RADIUS Server For Dial-up Or VPN Connections, click Configure VPN Or Dial-up and use the pages of the Configure VPN Or Dial-up Wizard to specify the set of policies for VPN or dial-up-based network access.
5. If you selected RADIUS Server For 802.1X Wireless or Wired Connections, click Configure 802.1X and use the pages of the Configure 802.1X Wizard to specify the set of policies for VPN or dial-up-based network access.

See the following chapters for information about configuring the appropriate policies with the Network Policy Server wizards:

- Chapter 10
- Chapter 11
- Chapter 12
- Chapter 13, “Site-to-Site Connections”
- Chapter 16
- Chapter 17
- Chapter 18
- Chapter 19

If the access servers require vendor-specific attributes (VSAs), you must add the VSAs to the appropriate network policy.

### To Add a VSA to a Network Policy

1. In the console tree of the Network Policy Server snap-in, expand Policies, and then click Network Policies.
2. Right-click the NPS network policy to which the VSA will be added, and then click Properties.
3. Click the Settings tab, click Vendor Specific, and then click Add. A list of predefined attributes appears in the Add Vendor Specific Attribute dialog box.
4. Look at the list of available RADIUS attributes to determine whether your vendor-specific attribute is already present. If it is, double-click it and configure it as specified in your access server's documentation.
5. If the vendor-specific attribute is not in the list of available RADIUS attributes, double-click Vendor-Specific. The Attribute Information dialog box appears.
6. Click Add. The Vendor-Specific Attribute Information dialog box appears.
7. To specify the network access server vendor for your access server from the list, click Select From List, and then select the network access vendor for which you are configuring the VSA.
8. If the vendor is not listed, click Enter Vendor Code, and then type the vendor code in the space provided.



**More Info** If you do not know the vendor code for your access server, see RFC 1007 for a list of SMI Network Management Private Enterprise Codes. RFC 1007 can be viewed at <http://www.ietf.org/rfc.html>.

9. Specify whether the attribute conforms to the RFC 2865 VSA specification. If you are not sure, see your access server documentation. If your attribute conforms, click Yes. It Conforms, and then click Configure Attribute. The Configure VSA (RFC-Compliant) dialog box appears.
10. In the Vendor-Assigned Attribute Number spin box, type the number that is assigned to the attribute (the numbers available are 0 through 255). In the Attribute Format drop-down list, specify the format for the attribute, and then in the Attribute Value text box, type the value that you are assigning to the attribute. Click OK twice.
11. If the attribute does not conform, click No. It Does Not Conform, and then click Configure Attribute. The Configure VSA (Non-RFC-Compliant) dialog box appears.

12. In the Hexadecimal Attribute Value text box, type the value for the attribute. Click OK twice.

## Configuring the Secondary NPS Server

To configure the secondary NPS server on a computer, do the following:

1. Obtain and install a computer certificate.
2. Configure the secondary NPS server computer to read the properties of user accounts in the domain.
3. Copy the configuration of the primary NPS server to the secondary NPS server.

### Copying the Configuration of the Primary NPS Server to the Secondary

**NPS server** To copy the configuration of the primary NPS server to the secondary NPS server, do the following:

1. On the primary NPS server computer, type **netsh nps export *path\file* exportpsk=yes** at a command prompt, which stores the configuration settings, including RADIUS shared secrets, in a text file at *path\file*. The path can be a relative, an absolute, or a network path.
2. Copy the file created in step 1 to the secondary NPS server.
3. On the secondary NPS server computer, type **netsh nps import *path\file*** at a command prompt, which imports all the settings configured on the primary NPS server into the secondary NPS server.

If you must change the NPS server configuration in any way, use the Network Policy Server snap-in to change the configuration of the NPS server that is designated as the primary configuration server, and then use this procedure to synchronize those changes on the secondary NPS server.

## Using RADIUS Proxies for Cross-Forest Authentication

Because NPS uses Active Directory to validate credentials and obtain user and computer account properties, a RADIUS proxy must be placed between the access servers and the NPS server computers when the user and computer accounts for access client computers and users exist in the following authentication databases:

- Two different Active Directory forests that do not trust each other
- Two different domains that do not trust each other
- Two different domains that have a one-way trust



### Direct from the Source: RADIUS Proxies and EAP-TLS

The use of a RADIUS proxy is required for EAP-TLS because part of the process requires a service principal name (SPN) lookup in Active Directory. However, SPN lookups do not work across trusts. When the NPS server receives the computer identity, it is in the form of an SPN (*host/ComputerName.DNSDomainName*). The NPS server passes the SPN to the local global catalog. If the global catalog is unable to match the SPN to a local domain account, it will fail the request with a No Valid Account Found error condition. SPN requests are not passed to the other domains.

*Clay Seymour, Support Escalation Engineer*

*Enterprise Platform Support*



**Note** You do not need to use a RADIUS proxy if you use PEAP-MS-CHAP v2 and user names like those used prior to Windows 2000 (*microsoft\user1*, for example).

When an access client sends user credentials, a user name is often included, which includes two elements:


- Identification of the user account name
- Identification of the user account location

For example, for the user name *user1@contoso.com*, *user1* is the user account name, and *contoso.com* is the location of the user account. The identification of the location of the user account is known as a *realm*, which has different forms:

- **The realm name can be a prefix.** In *contoso\user1*, *contoso* is the name of a domain like those used prior to Windows 2000.
- **The realm name can be a suffix.** For *user1@contoso.com*, *contoso.com* is either a DNS domain name or the name of an Active Directory-based domain.

The user name is passed from the access client to the access server during the authentication phase of the connection attempt. This user name becomes the User-Name RADIUS attribute in the Access-Request message sent by the access server to its configured RADIUS server, which is a RADIUS proxy in this configuration. When the RADIUS proxy receives the Access-Request message, connection request policies on the RADIUS proxy determine the RADIUS server to which the Access-Request message is forwarded based on the realm name.

Figure 9-4 shows NPS RADIUS proxies forwarding RADIUS messages between access servers and multiple NPS RADIUS servers in two different Active Directory forests.



**Figure 9-4** Using NPS RADIUS proxies for cross-forest authentication

The following configuration is for an organization that uses the following:

- **Active Directory domains** Active Directory domains contain the user accounts, passwords, and dial-in properties that each NPS RADIUS server requires to authenticate user credentials and evaluate authorization.
- **At least two NPS RADIUS servers in each forest** At least two NPS RADIUS servers (one primary and one secondary) can provide fault tolerance for RADIUS-based authentication, authorization, and accounting in each forest. If only one NPS RADIUS server is configured and it becomes unavailable, access clients for that forest cannot be authenticated. By using at least two NPS RADIUS servers and configuring the NPS RADIUS proxies for both the primary and secondary NPS RADIUS servers, the NPS RADIUS proxies can detect when the primary NPS RADIUS server is unavailable and then automatically fail over to the secondary NPS RADIUS server.
- **A network policy for network access** A network policy is configured on the NPS RADIUS servers to authorize network connections based on group membership.
- **At least two NPS RADIUS proxies** At least two NPS RADIUS proxies can provide fault tolerance for RADIUS requests that are sent from the access servers.

To deploy the configuration just described, do the following:

1. Configure the certificate infrastructure.
2. Configure the Active Directory forests for accounts and groups.
3. Configure the primary NPS RADIUS server on a computer in the first forest.
4. Configure the secondary NPS RADIUS server on another computer in the first forest.
5. Configure the primary NPS RADIUS server on a computer in the second forest.
6. Configure the secondary NPS RADIUS server on another computer in the second forest.
7. Configure the primary NPS RADIUS proxy.
8. Configure the secondary NPS RADIUS proxy.
9. Configure RADIUS authentication and accounting on the access servers.

## **Configuring the Certificate Infrastructure**

Follow the instructions in the “Deploying PKI” subsection of “Deployment Steps” earlier in this chapter.

## **Configuring the Active Directory Forests for Accounts and Groups**

Follow the instructions in the “Deploying Active Directory” subsection of “Deployment Steps” earlier in this chapter.

## **Configuring the Primary NPS Server on a Computer in the First Forest**

To configure the primary NPS RADIUS server on a computer in the first forest, perform on a computer in the first forest the steps described in the following subsections of “Configuring the Primary NPS Server” earlier in this chapter:

- “Obtaining and Installing a Computer Certificate”
- “Configuring NPS Server Properties”
- “Configuring Appropriate Policies”

Next, configure the primary NPS RADIUS server in the first forest with the primary and secondary NPS RADIUS proxies as RADIUS clients. To do this, perform the steps in the “Configuring NPS with RADIUS Clients” subsection of “Configuring the Primary NPS Server” earlier in this chapter. (Instead of the access servers, add the primary and secondary NPS RADIUS proxies as RADIUS clients.)

## **Configuring the Secondary NPS Server on Another Computer in the First Forest**

To configure the secondary NPS RADIUS server on another computer in the first forest, follow the instructions in “Configuring the Secondary NPS Server” earlier in this chapter.

## Configuring the Primary NPS Server on a Computer in the Second Forest

To configure the primary NPS RADIUS server on a computer in the second forest, perform the steps in the following subsections of “Configuring the Primary NPS Server” earlier in this chapter on a computer in the second forest:

- “Obtaining and Installing a Computer Certificate”
- “Configuring NPS Server Properties”
- “Configuring Appropriate Policies”

Next, configure the primary NPS RADIUS server in the second forest with the primary and secondary NPS RADIUS proxies as RADIUS clients. To do this, follow the instructions in the “Configuring NPS with RADIUS Clients” subsection of “Configuring the Primary NPS Server” earlier in this chapter (instead of the access servers, add the primary and secondary NPS RADIUS proxies as RADIUS clients).

## Configuring the Secondary NPS Server on Another Computer in the Second Forest

To configure the secondary NPS RADIUS server on another computer in the second forest, perform the steps in “Configuring the Secondary NPS Server” earlier in this chapter.

## Configuring the Primary NPS RADIUS Proxy

The computer acting as the primary NPS RADIUS proxy is not required to be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server. Because the primary NPS RADIUS proxy computer is not performing authentication or authorization of network access connections, it can be a member of a domain of either forest.

### To Configure the Primary NPS RADIUS Proxy for RADIUS Ports and Clients

1. In the Network Policy Server snap-in for the primary NPS RADIUS proxy, configure additional UDP ports for RADIUS messages that are sent by the access servers as needed. By default, NPS uses UDP ports 1812 and 1645 for authentication and UDP ports 1813 and 1646 for accounting.
2. Add the access servers as RADIUS clients by using the instructions in the “Configuring NPS with RADIUS Clients” section of “Configuring the Primary NPS Server” earlier in this chapter.

### To Configure the Primary NPS RADIUS Proxy for a Remote RADIUS Server Group Corresponding to the NPS RADIUS Servers in the First Forest

1. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients And Servers.
2. Right-click Remote RADIUS Server Groups, and then click New.

3. In the New Remote RADIUS Server Group dialog box, in the Group Name field, type the group name for the NPS RADIUS servers in the first forest (for example: RADIUS Servers in Forest1). Click Add.
4. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the primary NPS RADIUS server in the first forest. If you specify a name, click Verify to resolve the name to an IP address.
5. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the primary NPS server in the first forest.
6. Click OK to add the server to the list of servers in the group.
7. In the New Remote RADIUS Server Group dialog box, click Add.
8. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the secondary NPS RADIUS server in the first forest.
9. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the secondary NPS server in the first forest.
10. Click OK to add the server to the list of servers in the group, and then click OK again.

**To Configure the Primary NPS RADIUS Proxy for a Remote RADIUS Server Group Corresponding to the NPS RADIUS Servers in the Second Forest**

1. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients And Servers.
2. Right-click Remote RADIUS Server Groups, and then click New.
3. In the New Remote RADIUS Server Group dialog box, in the Group Name field, type the group name for the NPS RADIUS servers in the second forest (for example: RADIUS Servers in Forest2). Click Add.
4. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the primary NPS RADIUS server in the second forest. If you specify a name, click Verify to resolve the name to an IP address.
5. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the primary NPS RADIUS server in the second forest.
6. Click OK to add the server to the list of servers in the group.
7. In the New Remote RADIUS Server Group dialog box, click Add.
8. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of the secondary NPS RADIUS server in the second forest.
9. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the secondary NPS RADIUS server in the second forest.
10. Click OK to add the server to the list of servers in the group, and then click OK again.

**To Configure the Primary NPS RADIUS Proxy for a Connection Request Policy to Forward RADIUS Request Messages to the NPS RADIUS Servers in the First Forest**

1. In the console tree of the Network Policy Server snap-in, expand Policies, right-click Connection Request Policies, and then click New.
2. On the Specify Connection Request Policy Name And Connection Type page, in the Policy Name box, type the name for the connection request policy (for example: Forward Requests to RADIUS Servers in Forest1). Click Next.
3. On the Specify Conditions page, click Add.
4. In the Select Conditions dialog box, double-click User Name.
5. In the User Name dialog box, type the realm name for all names in the first forest (for example: forest1.example.com), click OK, and then click Next.
6. On the Specify Connection Request Forwarding page, select Forward Requests To The Following Remote RADIUS Server Group For Authentication, and then in the drop-down list, select the remote RADIUS server group for the NPS RADIUS servers in the first forest (for example: RADIUS Servers in Forest1). Click Next.
7. On the Configure Settings page, click Next,
8. On the Completing Connection Request Policy Wizard page, click Finish.

**To Configure the Primary NPS RADIUS Proxy for a Connection Request Policy to Forward RADIUS Request Messages to the NPS RADIUS Servers in the Second Forest**

1. In the console tree of the Network Policy Server snap-in, expand Policies, right-click Connection Request Policies, and then click New.
2. On the Specify Connection Request Policy Name And Connection Type page, in the Policy Name box, type the name for the connection request policy (for example: Forward Requests to RADIUS Servers in Forest2). Click Next.
3. On the Specify Conditions page, click Add.
4. In the Select Conditions dialog box, double-click User Name.
5. In the User Name dialog box, type the realm name for all names in the second forest (for example: forest2.example.com), click OK, and then click Next.
6. On the Specify Connection Request Forwarding page, select Forward Requests To The Following Remote RADIUS Server Group For Authentication, and then, in the drop-down list, select the remote RADIUS server group for the NPS RADIUS servers in the second forest (for example: RADIUS Servers in Forest2). Click Next.
7. On the Configure Settings page, click Next,
8. On the Completing Connection Request Policy Wizard page, click Finish.

## Configuring the Secondary NPS RADIUS Proxy

The computer acting as the secondary NPS RADIUS proxy is not required to be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server. Like the primary NPS RADIUS proxy, the secondary NPS RADIUS proxy computer can be a member of a domain of either forest because it is not performing authentication or authorization of network access connections.

### To Configure the Secondary NPS RADIUS Proxy on Another Computer

1. On the primary NPS RADIUS proxy computer, type **netsh nps export *path*\file exportpsk=yes** at a command prompt.

This command stores the configuration settings, including RADIUS shared secrets, in a text file. The path can be relative, absolute, or a network path.

2. Copy the file created in step 1 to the secondary NPS RADIUS proxy.
3. On the secondary NPS RADIUS proxy computer, type **netsh nps import *path*\file** at a command prompt.

This command imports all the settings configured on the primary NPS RADIUS proxy into the secondary NPS RADIUS proxy.

Based on the default load-balancing settings of the RADIUS servers in the two remote RADIUS server groups, each NPS RADIUS proxy will distribute the authentication request load equally to the two NPS servers in each forest.

## Configuring RADIUS Authentication on the Access Servers

Configure the RADIUS client on your access servers with the following settings:

- The IP address or name of a primary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings.
- The IP address or name of a secondary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings.

To balance the load of RADIUS traffic between the primary and secondary NPS RADIUS proxies, configure half of the access servers with the primary NPS RADIUS proxy as their primary RADIUS server and the secondary NPS RADIUS proxy as their secondary RADIUS server. Configure the other half of the access servers with the secondary NPS RADIUS proxy as their primary RADIUS server and the primary NPS RADIUS proxy as their secondary RADIUS server.

## Using RADIUS Proxies to Scale Authentications

When performing authentication for a large number of access clients by using certificate-based authentication or for a large NAP deployment, the volume of RADIUS authentication traffic necessary to keep access clients connected can be substantial. In a large deployment, it is best to spread the load of authentication traffic among multiple NPS server computers. Because you cannot rely on the access servers to consistently or adequately spread their


authentication traffic among multiple RADIUS servers, intermediate NPS RADIUS proxies can provide this function.

Without the RADIUS proxies, each access server sends its RADIUS requests to one or multiple RADIUS servers and detects unavailable RADIUS servers. The access server might or might not be balancing the load of RADIUS traffic across multiple RADIUS servers. By using NPS RADIUS proxies, consistent load balancing spreads the load of authentication, authorization, and accounting traffic across all the NPS servers in the organization. Additionally, there is a consistent scheme for failure detection and RADIUS server failover (the detection of an unavailable RADIUS server and avoidance of its use for future authentication requests) and failback (the detection that a previously unavailable RADIUS server is available).

The following configuration is for an organization that uses the following:

- **Active Directory domains** Active Directory domains contain the user accounts, passwords, and dial-in properties that each NPS server requires to authenticate user credentials and evaluate authorization.
- **Multiple NPS servers** To balance a large load of RADIUS authentication, authorization, and accounting traffic, there are multiple NPS servers.
- **Network policies** Network policies are configured to authenticate and authorize network access based on group membership.
- **Two NPS RADIUS proxies** Two NPS RADIUS proxies provide fault tolerance for RADIUS requests that are sent from the access servers.

Figure 9-5 shows the use of NPS RADIUS proxies to balance the load of RADIUS traffic from access servers across multiple NPS servers.



**Figure 9-5** Using NPS RADIUS proxies to load-balance RADIUS traffic



To deploy this configuration, do the following:

1. Configure the certificate infrastructure.
2. Configure Active Directory for accounts and groups.
3. Configure NPS as a RADIUS server on multiple computers.
4. Configure the primary NPS RADIUS proxy.
5. Configure the secondary NPS RADIUS proxy.
6. Configure RADIUS authentication and accounting on access servers.

## Configuring the Certificate Infrastructure

Follow the instructions in the “Deploying PKI” subsection of “Deployment Steps” earlier in this chapter.

## Configuring Active Directory for Accounts and Groups

Follow the instructions in the “Deploying Active Directory” subsection of “Deployment Steps” earlier in this chapter.

## Configuring NPS as a RADIUS Server on Multiple Computers

To configure NPS on each NPS server computer, perform on each NPS server computer the steps described in the following subsections of “Configuring the Primary NPS Server” earlier in this chapter:

- “Obtaining and Installing a Computer Certificate”
- “Configuring NPS Server Properties”
- “Configuring Appropriate Policies”

Next, configure each NPS server computer with the primary and secondary NPS RADIUS proxies as RADIUS clients. To do this, perform the steps in the “Configuring NPS with RADIUS Clients” subsection of “Configuring the Primary NPS Server” earlier in this chapter. (Instead of the access servers, add the primary and secondary NPS RADIUS proxies as RADIUS clients.)



**Note** You can configure each NPS RADIUS server separately rather than configuring an initial NPS RADIUS server and copying its configuration to other NPS RADIUS server computers. This is done so that different RADIUS shared secrets can be used between the NPS RADIUS proxies and the NPS RADIUS server.

## Configuring the Primary NPS RADIUS Proxy

The computer acting as the primary NPS RADIUS proxy need not be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server.

### To Configure the Primary NPS RADIUS Proxy

1. In the Network Policy Server snap-in, configure additional UDP ports for RADIUS messages that are sent by the access servers if needed.  
  
By default, NPS uses UDP ports 1812 and 1645 for authentication and UDP ports 1813 and 1646 for accounting.
2. Add the access servers as RADIUS clients of the NPS RADIUS proxy by following the steps in the “Configuring NPS with RADIUS Clients” subsection of “Configuring the Primary NPS Server” earlier in this chapter.
3. In the console tree of the Network Policy Server snap-in, expand RADIUS Clients and Servers.
4. Right-click Remote RADIUS Server Groups, and then click New.
5. In the New Remote RADIUS Server Group box, type the group name for all of the NPS RADIUS servers (for example: RADIUS Servers in the contoso.com Domain).
6. Click Add.
7. On the Address tab, type the DNS name, IPv4 address, or IPv6 address of an NPS RADIUS server. If you specify a name, click Verify to resolve the name to an IP address.
8. On the Authentication/Accounting tab, type the shared secret between the primary and secondary NPS RADIUS proxies and the NPS RADIUS server.
9. Click OK to add the server to the list of servers in the group.
10. Repeat steps 6–9 for each NPS RADIUS server, and then click OK.
11. In the console tree of the Network Policy Server snap-in, expand Policies, right-click Connection Request Policies, and then click New.
12. On the Specify Connection Request Policy Name And Connection Type page, in the Policy Name box, type the name for the connection request policy (for example: Forward Requests to RADIUS Servers in the contoso.com Domain). Click Next.
13. On the Specify Conditions page, click Add.
14. In the Select Conditions dialog box, double-click User Name.
15. In the User Name dialog box, type the realm name for all names in the second forest (for example: forest2.example.com), click OK, and then click Next.
16. On the Specify Connection Request Forwarding page, select Forward Requests To The Following Remote RADIUS Server Group For Authentication, and then in the drop-down list, select the remote RADIUS server group for all of the NPS RADIUS servers in the domain. Click Next.

17. On the Configure Settings page, click Next,
18. On the Completing Connection Request Policy Wizard page, click Finish.

## Configuring the Secondary NPS RADIUS Proxy

The computer acting as the secondary NPS RADIUS proxy need not be dedicated to forwarding RADIUS messages. For example, you can install NPS on a file server.

### To Configure the Secondary NPS RADIUS Proxy on Another Computer

1. On the primary NPS RADIUS proxy computer, type **netsh nps export *path*\file exportpsk=yes** at a command prompt.

This command stores the configuration settings, including RADIUS shared secrets, in a text file. The path can be relative, absolute, or a network path.

2. Copy the file created in step 1 to the secondary NPS RADIUS proxy computer.
3. On the secondary NPS RADIUS proxy computer, type **netsh nps import *path*\file** at a command prompt. This command imports all the settings configured on the primary NPS RADIUS proxy into the secondary NPS RADIUS proxy.

Based on the default load-balancing settings of the RADIUS servers in the remote RADIUS server group, each NPS RADIUS proxy distributes the authentication request load equally to all of the NPS RADIUS servers.

## Configuring RADIUS Authentication on the Access Servers

Configure the RADIUS client on your access servers with the following settings:

- The IP address or name of a primary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings
- The IP address or name of a secondary RADIUS server, the shared secret, UDP ports for authentication and accounting, and failure-detection settings

To balance the load of RADIUS traffic between the primary and secondary NPS RADIUS proxies, configure half of the access servers with the primary NPS RADIUS proxy as their primary RADIUS server and the secondary NPS RADIUS proxy as their secondary RADIUS server. Configure the other half of the access servers with the secondary NPS RADIUS proxy as their primary RADIUS server and the primary NPS RADIUS proxy as their secondary RADIUS server.

## Ongoing Maintenance

This section describes the ongoing maintenance for the following components of a Windows authentication infrastructure for network access:

- Active Directory
- PKI

- Group Policy
- RADIUS

## Active Directory

It is beyond the scope of this book to describe the ongoing maintenance of an Active Directory infrastructure for an organization of an arbitrary size. For detailed information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/ad>.

The elements of maintaining Active Directory to best support a Windows-based authentication infrastructure for network access are as follows:

- When adding user or computer accounts, ensure that the new accounts have the appropriate security group membership to allow network access. For example, if wireless access is being granted through membership in the WirelessUsers group, add new user or computer accounts to this group or to a group that is a member of this group.
- When adding new domains or forests, ensure that the appropriate trust relationships are created to allow NPS RADIUS servers to authenticate account credentials. Additionally, add the computer accounts of the NPS RADIUS servers to the RAS and IAS Servers security groups of the new domains. If the new domains or forests do not have a trust relationship, use NPS RADIUS proxies to provide cross-domain or cross-forest authentication. For more information, see “Using RADIUS Proxies for Cross-Forest Authentication” earlier in this chapter.

## PKI

It is beyond the scope of this book to describe the ongoing maintenance of a PKI for an organization of an arbitrary size. For detailed information, see Windows Server 2008 Help and Support or the resources at <http://www.microsoft.com/pki>.

## Group Policy

It is beyond the scope of this book to describe the ongoing maintenance of Group Policy for an organization of an arbitrary size. For detailed information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista*, Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/gp>.

The elements of maintaining Group Policy to best support a Windows-based authentication infrastructure for network access are as follows:

- When adding new domains or forests, ensure that the appropriate Group Policy objects are applied to the appropriate Active Directory containers to propagate settings for autoenrollment of certificates or configuration settings.

## RADIUS

The following sections describe how to maintain the RADIUS component of the network access infrastructure.

### Adding a New NPS RADIUS Server to the RADIUS Infrastructure

When you add a new NPS RADIUS server to the RADIUS infrastructure, you must do the following:

1. Register the new NPS server in its default domain.
2. Register the new NPS server in other domains.
3. If the new NPS server is a secondary RADIUS server, obtain and install a computer certificate if needed, and copy the configuration of the primary RADIUS server to the new NPS server.
4. If the new NPS server is a primary RADIUS server, do the following:
  - ❑ Obtain and install a computer certificate.
  - ❑ Configure NPS server properties.
  - ❑ Configure NPS with RADIUS clients.
  - ❑ Configure NPS with the appropriate network policies.
5. Configure access servers (RADIUS clients) to use the new NPS server.
6. If IPsec is being used to protect RADIUS traffic, update Windows Firewall with Advanced Security connection security rules to include protection for RADIUS traffic to and from the new NPS server.

Instructions for these procedures can be found in the “RADIUS Servers” subsection of “Deployment Steps” earlier in this chapter.

### Removing an NPS RADIUS Server from the RADIUS Infrastructure

When you remove an NPS RADIUS server from the RADIUS infrastructure, you must do the following:

1. Reconfigure your access servers to remove references to the NPS server that is being removed.
2. Remove the computer account of the NPS server that is being removed from the RAS and IAS Servers security group of its default domain.
3. Remove the computer account of the NPS server that is being removed from the RAS and IAS Servers security group of other domains.
4. If IPsec is being used to protect RADIUS traffic to and from the NPS server that is being removed, update Windows Firewall with Advanced Security connection security rules to remove protection for the NPS server.

## Maintaining RADIUS Clients

When you deploy a new access server, such as a new wireless AP for your wireless network, you must do the following:

1. Add the access server as a RADIUS client to either your NPS RADIUS servers or your NPS RADIUS proxies.
2. Configure the access server to use your NPS RADIUS servers or your NPS RADIUS proxies.
3. If IPsec is being used to protect traffic between your RADIUS servers or proxies and the access server, update Windows Firewall with Advanced Security connection security rules to include protection for RADIUS traffic to and from the new access server.

When you remove an access server, you must do the following:

1. Delete the access server as a RADIUS client on either your NPS RADIUS servers or your NPS RADIUS proxies.
2. If IPsec is being used to protect traffic between your RADIUS servers and the access server, update Windows Firewall with Advanced Security connection security rules to remove protection for RADIUS traffic between the access server and the NPS RADIUS servers or proxies.

## Troubleshooting Tools

This section describes the troubleshooting tools or the resources that describe troubleshooting tools for the following components of a Windows authentication infrastructure for network access:

- Active Directory
- PKI
- Group Policy
- RADIUS

### Active Directory

It is beyond the scope of this book to describe in detail the troubleshooting tools for Active Directory. For additional information, see the *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit*, Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/ad>.

Active Directory–specific troubleshooting issues are described as needed in subsequent chapters to troubleshoot network access or NAP.

## PKI

It is beyond the scope of this book to describe in detail the troubleshooting tools for a Windows-based PKI. For additional information, see *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008), Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/pki>.

Digital certificate and PKI-specific troubleshooting issues are described as needed in subsequent chapters to troubleshoot network access or NAP.

## Group Policy

It is beyond the scope of this book to describe in detail the troubleshooting tools for Group Policy. For additional information, see the *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* by Derek Melber, Group Policy MVP, with the Windows Group Policy Team (Microsoft Press, 2008) Windows Server 2008 Help and Support, or the resources at <http://www.microsoft.com/gp>.

Group Policy-specific troubleshooting issues are described as needed in subsequent chapters to troubleshoot network access or NAP.

## RADIUS

To help you gather information to troubleshoot problems with NPS, Microsoft provides the following troubleshooting tools:

- NPS event logging and Windows Event Viewer
- Network Monitor 3.1
- Performance Monitor counters
- SNMP Service

### NPS Event Logging and Windows Event Viewer

Use Event Viewer, available from the Administrative Tools program group, to obtain information about hardware and software problems and to monitor all security events, including informational, warning, and error events.

To troubleshoot NPS authentication attempts, view the NPS events in Windows Logs\Security. Viewing the authentication attempts in this log is useful in troubleshooting network policies. When you have multiple network policies configured, you can use the security event log to determine the name of the network policy that either accepted or rejected the connection attempt. Enabling NPS event logging and reading the text of NPS authentication events in the security event log is the most useful tool for troubleshooting failed NPS authentications.

To view the NPS events, configure a filter with the Event Sources option set to Microsoft Windows Security Auditing and the Task Category option set to Network Policy Server.

Both types of logging (rejected authentication requests and successful authentication requests) are enabled by default.

### To Configure NPS for Event Logging

1. In the console tree of the Network Policy Server snap-in, right-click NPS, and then click Properties.
2. On the General tab, select each required check box, and then click OK.

## Network Monitor 3.1

You can use Network Monitor 3.1 (or later) or a commercial packet analyzer (also known as a *network sniffer*), to capture and view RADIUS authentication and accounting messages that are sent to and from the NPS server. Network Monitor 3.1 includes a RADIUS parser, which you can use to view the attributes of a RADIUS message and troubleshoot network access or NAP issues.



**On the Disc** You can link to the download site for Network Monitor from the companion CD-ROM.

## Reliability and Performance Counters

You can use the Reliability and Performance snap-in to monitor the resource use of specific components and program processes. With Performance Monitor, which is in the Reliability and Performance snap-in, you can use charts and reports to determine how efficiently your server uses NPS and both identify and troubleshoot potential problems.

You can use Performance Monitor to monitor the following NPS-related performance objects:

- NPS Accounting Clients
- NPS Accounting Server
- NPS Authentication Clients
- NPS Authentication Server

## SNMP Service

You can use the Simple Network Management Protocol (SNMP) service to monitor status information for your NPS server. NPS supports the RADIUS Authentication Server Management Information Base (MIB), as specified in RFC 2619, and the RADIUS Accounting Server MIB, as specified in RFC 2621.



## Chapter Summary

A Windows-based network access infrastructure consists of Active Directory, PKI, Group Policy, and RADIUS components. Active Directory stores user and computer account credentials and properties and provides an infrastructure to deploy centrally configured user and computer configuration Group Policy settings. A PKI issues and validates digital certificates used in different types of network access scenarios or NAP enforcement methods. Group Policy settings can instruct computers to automatically request specific types of certificates or configure network access and protection settings. RADIUS provides a standard protocol and centralized management of network access authorization, authentication, and accounting.

The combination of Active Directory, PKI, Group Policy, and RADIUS creates a Windows-based infrastructure that provides centralized authentication for 802.11 wireless access, 802.1X wired access, dial-up or VPN-based remote access connections, and dial-up or VPN-based site-to-site connections. The combination of PKI, Group Policy, and RADIUS creates a Windows-based infrastructure that provides centralized configuration and validation of system health status for NAP.

## Additional Information

For additional information about Active Directory, see the following:

- *Windows Server 2008 Active Directory Resource Kit* in the *Windows Server 2008 Resource Kit* (both from Microsoft Press, 2008)
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- Microsoft Windows Server Active Directory (<http://www.microsoft.com/ad>)

For additional information about PKI, see the following:

- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- Microsoft Public Key Infrastructure for Windows Server (<http://www.microsoft.com/pki>)
- *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008)

For additional information about Group Policy, see the following:

- *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* (Microsoft Press, 2008)
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- Microsoft Windows Server Group Policy (<http://www.microsoft.com/gp>)

For additional information about RADIUS and NPS, see the following:

- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- Network Policy Server (<http://www.microsoft.com/nps>)
- RFC 2548, “Microsoft Vendor-Specific RADIUS Attributes”
- RFC 2619, “RADIUS Authentication Server MIB”
- RFC 2621, “RADIUS Accounting Server MIB”
- RFC 2865, “Remote Authentication Dial-In User Service (RADIUS)”
- RFC 2866, “RADIUS Accounting”
- RFC 2867, “RADIUS Accounting Modifications for Tunnel Protocol Support”
- RFC 2868, “RADIUS Attributes for Tunnel Protocol Support”
- RFC 2869, “RADIUS Extensions”
- RFC 3162, “RADIUS and IPv6”
- RFC 3579, “RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)”

For additional information about Windows-based network access, see the following:

- Chapter 10, “IEEE 802.11 Wireless Networks”
- Chapter 11, “IEEE 802.1X-Authenticated Wired Networks”
- Chapter 12, “Remote Access VPN Connections”
- Chapter 13, “Site-to-Site VPN Connections”

For additional information about NAP, see the following:

- Chapter 14, “Network Access Protection Overview”
- Chapter 15, “Preparing for Network Access Protection”
- Chapter 16, “IPsec Enforcement”
- Chapter 17, “802.1X Enforcement”
- Chapter 18, “VPN Enforcement”
- Chapter 19, “DHCP Enforcement”
- The Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- Microsoft Network Access Protection (<http://www.microsoft.com/nap>)

## Chapter 12

# Remote Access VPN Connections

This chapter provides information about how to design, deploy, maintain, and troubleshoot remote access virtual private network (VPN) connections. Once deployed, the remote access VPN solution can be modified for the VPN Enforcement method of Network Access Protection (NAP), as described in Chapter 18, “VPN Enforcement.”

This chapter assumes that you understand the role of Active Directory, public key infrastructure (PKI), Group Policy, and Remote Authentication Dial-up User Service (RADIUS) elements of a Windows-based authentication infrastructure for network access, as described in Chapter 9, “Authentication Infrastructure.”



**More Info** This chapter does not describe the deployment planning and steps for dial-up remote access. For more information on those topics, see Windows Server 2008 Help and Support or the Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>.

## Concepts

A VPN is the extension of a private network that encompasses links across shared or public networks such as the Internet. With a VPN, you can send data between two computers across a shared or public network in a manner that emulates a point-to-point private link, such as a long-haul T-Carrier-based wide area network (WAN) link. Virtual private networking is the act of creating and configuring a virtual private network.

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information, which allows the data to traverse the shared or public network to reach its endpoint. To emulate a private link, the data is encrypted for confidentiality. The link in which the private data is encapsulated and encrypted is the VPN connection.

Users working at home or on the road can use VPN connections to establish a remote access connection to an organization’s server by using the infrastructure provided by a public network such as the Internet. From the user’s perspective, the VPN connection is a point-to-point connection between the computer (the VPN client) and an organization server (the VPN server). The exact infrastructure of the shared or public network is irrelevant because it appears logically as if the data is sent over a dedicated private link.

Organizations can also use VPN connections to establish routed connections with geographically separate offices or with other organizations over a public network such as the Internet while maintaining secure communications. A routed VPN connection across the Internet logically operates as a dedicated WAN link. For more information about routed VPN connections, see Chapter 13, “Site-to-Site VPN Connections.”

With both remote access and routed connections, an organization can use VPN connections in place of long-distance dial-up or leased lines for connecting to an Internet service provider (ISP).

There are three types of remote access VPN technologies in the Windows Server 2008 and Windows Vista operating systems:

- **Point-to-Point Tunneling Protocol (PPTP)** PPTP uses Point-to-Point Protocol (PPP) authentication methods for user-level authentication and Microsoft Point-to-Point Encryption (MPPE) for data encryption.
- **Layer Two Tunneling Protocol with Internet Protocol security (L2TP/IPsec)** L2TP/IPsec uses PPP authentication methods for user-level authentication and IPsec for computer-level peer authentication, data authentication, data integrity, and data encryption.
- **Secure Socket Tunneling Protocol (SSTP)** SSTP uses PPP authentication methods for user-level authentication and Hypertext Transfer Protocol (HTTP) encapsulation over a Secure Sockets Layer (SSL) channel (also known as a Transport Layer Security or TLS channel) for data authentication, data integrity, and data encryption.

A remote access client (a single user computer) makes a remote access VPN connection to a private network through a VPN server. The VPN server can provide access to the entire network to which the VPN server is attached. The packets sent from the remote client across the VPN connection originate at the remote access client computer.

During the connection process, the remote access client (the VPN client) authenticates itself to the remote access server (the VPN server), and for authentication methods that support mutual authentication, the server authenticates itself to the client.



**Note** Using IPsec tunnel mode as a remote access VPN technology is not supported by Windows-based VPN clients or servers because of the lack of an industry standard method of performing user authentication and IP address configuration over an IPsec tunnel. IPsec tunnel mode is described in Requests for Comments (RFCs) 2401, 2402, and 2406.

### Direct from the Source: Enhancements to PPTP and L2TP/IPsec

In Windows Server 2008 and Windows Vista, VPN security has been enhanced for the following:

- **PPTP** MPPE encryption with a 40-bit and 56-bit key has been disabled by default in Windows Server 2008 and Windows Vista. PPTP connections now support only 128-bit MPPE keys by default. If a Windows Vista-based VPN client is connecting to a Windows Server 2003-based VPN server, or if a Windows XP-based VPN client is connecting to a Windows Server 2008-based VPN server, connections will be successful only if both the VPN client and VPN server are configured to use 128-bit MPPE encryption.

You can configure Windows Server 2008 and Windows Vista to use 40-bit and 56-bit MPPE keys for PPTP connections by setting the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Rasman\Parameters\AllowPPTPWeakCrypto` registry value to **1** and then restarting the computer. However, this is not recommended.

- **L2TP/IPsec** In L2TP connections, use of IPsec with the Data Encryption Standard (DES) and the Message Digest 5 (MD5) hashed message authentication code (HMAC) in Windows Server 2008 and Windows Vista has been disabled by default. L2TP/IPsec connections now support only 3DES encryption and the Secure Hash Algorithm-1 (SHA1) HMAC by default. If a Windows Vista-based VPN client is connecting to a Windows Server 2003-based VPN server, or if a Windows XP-based VPN client is connecting to a Windows Server 2008-based VPN server, connections will be successful only if both the VPN client and VPN server are configured to use 3DES encryption and the SHA1 HMAC. However, support for the Advanced Encryption Standard (AES) using 128-bit or 256-bit keys has been added.

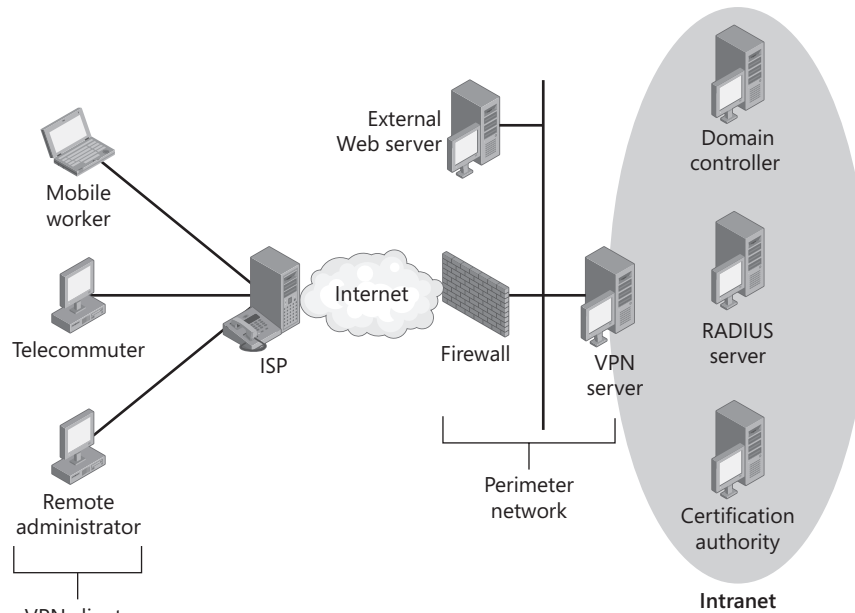
You can configure Windows Server 2008 and Windows Vista to use DES encryption and the MD5 HMAC for L2TP/IPsec connections by setting the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Rasman\Parameters\AllowL2TPWeakCrypto` registry value to **1** and then restarting the computer. However, this is not recommended.

*Samir Jain, Lead Program Manager*

*India Development Center*

## Components of Windows Remote Access VPNs

Figure 12-1 shows the components of Windows-based remote access VPNs.



**Figure 12-1** Components of Windows-based remote access VPNs

The components are:

- **VPN clients** VPN clients initiate remote access VPN connections to VPN servers and communicate with intranet resources once connected.
- **VPN servers** VPN servers listen for remote access VPN connection attempts, enforce authentication and connection requirements, and route packets between VPN clients and intranet resources.
- **RADIUS servers** RADIUS servers provide centralized authentication and authorization processing and accounting for network access attempts from multiple VPN servers (and other types of access servers).
- **Active Directory domain controllers** Active Directory domain controllers validate user credentials for authentication and provide user account information to evaluate authorization.
- **Certification authorities (CAs)** CAs are part of the PKI, and they issue computer or user certificates to VPN clients and computer certificates to VPN servers and RADIUS servers for computer-level and user-level authentication of VPN connections.



**More Info** *Computer certificates* are certificates that are stored in the local computer certificate store and have the appropriate properties to perform PPP-based, SSL-based, or IPsec-based authentication. For more details about certificate requirements for PPP-based or SSL-based authentication, see “Network Access Authentication and Certificates” at <http://go.microsoft.com/fwlink/?LinkID=20016>. For more details about certificate requirements for IPsec-based authentication, see “How IPsec Works” at <http://go.microsoft.com/fwlink/?LinkID=67907>.

Typical users of remote access VPN connections are:

- Laptop users who connect to an intranet to access e-mail and other resources while traveling
- Telecommuters who use the Internet to access intranet resources from home
- Remote administrators who use the Internet to connect to a private network and configure network or application services

## Planning and Design Considerations

When deploying a remote access VPN solution, you must consider the following planning and design issues:

- VPN protocols
- Authentication methods
- VPN servers
- Internet infrastructure
- Intranet infrastructure
- Concurrent intranet and Internet access for VPN clients
- Authentication infrastructure
- VPN clients
- PKI
- VPN enforcement with NAP

## VPN Protocols

Windows Server 2008 includes support for the following remote access VPN protocols:

- **PPTP** PPTP uses PPP user authentication and MPPE encryption. When Microsoft Challenge Handshake Authentication Protocol (MS-CHAP v2) or Protected EAP (PEAP)-MS-CHAP v2 is used with strong passwords, PPTP is a secure VPN technology. For certificate-based authentication, Extensible Authentication Protocol-Transport Layer

Security (EAP-TLS) can be used with registry-based certificates or smart cards. PPTP is widely supported, easily deployed, and can be used across most network address translators (NATs). PPTP is supported by the Windows Server 2008, Windows Vista, Windows Server 2003, and Windows XP operating systems.

- **L2TP/IPsec** L2TP utilizes PPP user authentication and IPsec packet protection. L2TP/IPsec uses certificates (by default) and the IPsec computer-level authentication process to negotiate the protected IPsec session and then PPP-based user authentication to authenticate the user of the VPN client computer. By using IPsec, L2TP/IPsec provides data confidentiality (encryption), data integrity (proof that the data was not modified in transit), and data origin authentication (proof that the data was sent by the authorized user) for each packet. However, L2TP/IPsec requires a PKI to allocate computer certificates to each L2TP/IPsec-based VPN client. L2TP/IPsec is supported by Windows Server 2008, Windows Vista, Windows Server 2003, and Windows XP.
- **SSTP** SSTP utilizes PPP user authentication and an HTTP-over-SSL channel for encapsulation and encryption. Because SSTP uses SSL traffic (using TCP port 443), SSTP can be used in many different network configurations, such as when VPN clients or servers are behind network address translations (NATs), firewalls, or proxy servers that can block or are not designed to forward PPTP or L2TP/IPsec traffic. SSTP is supported only by Windows Server 2008 and Windows Vista SP1.

## Design Choices for VPN Protocols

- When using PEAP-MS-CHAP v2, EAP-MS-CHAP v2, or MS-CHAP v2 for authentication, PPTP does not require a certificate infrastructure to issue certificates to each VPN client.
- PPTP-based VPN connections provide data confidentiality (encryption) for packets. PPTP-based VPN connections do not provide data integrity or data origin authentication.
- By using IPsec, L2TP/IPsec-based VPN connections provide data confidentiality, data integrity, and data origin authentication.
- SSTP-based VPN clients and servers can be behind NATs, firewalls, or Web proxies. However, SSTP does not support VPN clients or servers that are located behind authenticating Web proxies.
- By default, a VPN server running Windows Server 2008 supports all three types of VPN connections simultaneously. You can use PPTP for some remote access VPN connections (for example, from VPN clients that do not have an installed computer certificate), L2TP/IPsec for other remote access VPN connections (for example, from VPN clients that have an installed computer certificate), and SSTP for VPN clients running Windows Vista SP1.
- If you are using a combination of VPN protocols, you can create separate network policies that define different connection settings for PPTP, L2TP/IPsec, or SSTP-based connections.



- In Windows Server 2008 and Windows Vista, IPv6 traffic can be sent over a PPTP-based VPN connection as IPv4-tunneled traffic or as native IPv6 traffic inside the VPN tunnel. For more information, see “How It Works: IPv6 and VPN Connections” later in this chapter.
- In Windows Server 2008 and Windows Vista, L2TP/IPsec and SSTP-based VPN connections support IPv6 traffic as IPv4-tunneled traffic, as native IPv6 traffic inside the VPN tunnel, and for VPN connections over IPv6. For more information, see “How It Works: IPv6 and VPN Connections” later in this chapter.

## Requirements for VPN Protocols

- PPTP-based VPN clients can be located behind a NAT if the NAT includes a NAT editor that knows how to properly translate PPTP tunneled data. For example, both the Internet Connection Sharing (ICS) feature of the Network Connections folder and the NAT routing protocol component of Routing and Remote Access include a NAT editor that translates PPTP traffic to and from PPTP clients located behind the NAT. VPN servers cannot be behind a NAT unless there are multiple public IP addresses and there is a one-to-one mapping of a public IP address to the private IP address of the VPN server. If there is only one public address, the NAT must be configured to translate and forward the PPTP tunneled data to the VPN server. Most NATs using a single public IPv4 address, including ICS and the NAT routing protocol component, can be configured to allow inbound traffic based on IPv4 addresses and TCP and UDP ports. However, PPTP tunneled data does not use TCP or UDP headers. Therefore, a VPN server cannot be located behind a computer using ICS or the NAT routing protocol component when using a single public IPv4 address.
- L2TP/IPsec-based VPN clients or servers cannot be behind a NAT unless both the client and server support IPsec NAT Traversal (NAT-T). IPsec NAT-T is supported by Windows Server 2008, Windows Vista, Windows Server 2003, and Windows XP SP2.
- L2TP/IPsec supports computer certificates as the default and recommended authentication method for IPsec. Although you can configure a preshared key to authenticate L2TP/IPsec connections, this is not recommended except as a transition authentication method when deploying a PKI. Computer certificate authentication requires a PKI to issue computer certificates to the VPN server computer and all VPN client computers.
- SSTP is supported only by Windows Server 2008 (as a VPN server or client) and Windows Vista SP1 (as a VPN client).
- SSTP uses an encrypted SSL channel to protect all data sent across the VPN connection. To create this encrypted channel, the VPN server must have a computer certificate, and the VPN client computer must be able to validate the computer certificate of the VPN server. This means that the VPN clients must have the root CA certificate of the issuing CA of the VPN server’s computer certificate installed.

- If you want to send native IPv6 traffic inside the VPN tunnel across a demand-dial VPN connection, you must use L2TP/IPsec. For more information, see “How It Works: IPv6 and VPN Connections.”
- If you want to use demand-dial VPN connections across the IPv6 Internet, you must use L2TP/IPsec.

## Best Practices for VPN Protocols

- If you already have a PKI in place, use L2TP/IPsec instead of PPTP.
- If you are not using all the VPN protocols, configure the Ports node in the Routing and Remote Access snap-in to set the number of ports for unused VPN protocols to 0.

### How It Works: IPv6 and VPN Connections

For VPN connections, Windows Server 2008 and Windows Vista support IPv6 in the following ways:

- IPv4-tunneled IPv6 traffic
- Native IPv6 traffic inside the VPN tunnel
- VPN connections over IPv6

#### IPv4-Tunneled IPv6 Traffic

In Windows XP and Windows Server 2003, you could send IPv6 traffic over a VPN connection, but only if it was already wrapped in an IPv4 header (IPv4 tunneling). With IPv4-tunneled IPv6 traffic support, a remote access client can create a VPN connection across the IPv4 Internet and then use IPv4-tunneled IPv6 traffic to communicate with IPv6/IPv4 nodes or IPv6 nodes on the intranet.

IPv4-tunneled IPv6 traffic sent over a VPN connection consist of IPv6 packets that are wrapped with an IPv4 header (this is the IPv4 tunneling), which are wrapped with a PPP header and a VPN protocol header (such as PPTP or L2TP/IPsec), which are wrapped with a final IPv4 header to allow the packet to traverse the IPv4 Internet.

PPTP, L2TP/IPsec, and SSTP in Windows Server 2008 and Windows Vista support IPv4-tunneled IPv6 traffic. IPv4-tunneled IPv6 traffic sent over a VPN connection requires Internet Protocol Control Protocol (IPCP) support on the VPN client and VPN server, IPv6 transition technology support on the VPN client, and an IPv6 transition technology infrastructure (such as ISATAP) on the intranet. IPCP is a PPP network control protocol that allows PPP hosts to configure settings for using IPv4 over a PPP link.

#### Native IPv6 Traffic Inside the VPN Tunnel

Windows Server 2008 and Windows Vista support VPN connections with native IPv6 traffic inside the VPN tunnel. The VPN client creates a VPN connection with a VPN

server over the IPv4 Internet and then negotiates the use of IPv6 over the PPP link. IPv6 packets are encapsulated by the VPN protocol inside of the VPN tunnel. With native support for IPv6 traffic inside the VPN tunnel, a remote access client can create a VPN connection across the IPv4 Internet and then use native IPv6 traffic to communicate with IPv6 nodes on the intranet.

Native IPv6 traffic inside the VPN tunnel consists of IPv6 packets that are wrapped with a PPP header and a VPN protocol header, which are wrapped with a final IPv4 header to allow the packet to traverse the IPv4 Internet.

Native IPv6 traffic inside the VPN tunnel requires IPv6 Control Protocol (IPV6CP) support on the VPN client and VPN server, IPv6 routing support on the VPN server, and a native IPv6 routing infrastructure on the intranet. PPTP, L2TP/IPsec, and SSTP in Windows Server 2008 and Windows Vista support native IPv6 traffic inside the VPN tunnel. IPV6CP is a PPP network control protocol that allows PPP hosts to configure settings for using IPv6 over a PPP link.



**Note** Windows XP and Windows Server 2003 do not support native IPv6 traffic inside the VPN tunnel.

### VPN Connections Over IPv6

Windows Server 2008 and Windows Vista also support VPN connections over IPv6. The VPN client creates a VPN connection with a VPN server over the IPv6 Internet and then negotiates the use of either IPv6 or IPv4 over the PPP link. With VPN connections over IPv6 support, a remote access client can create a VPN connection across the IPv6 Internet and then use either native IPv6 or IPv4 traffic to communicate with nodes on the intranet.

Traffic for VPN connections over IPv6 consist of IPv6 or IPv4 packets that are wrapped with a PPP header and a VPN protocol header, which are wrapped with a final IPv6 header to allow the packet to traverse the IPv6 Internet.

SSTP and L2TP/IPsec in Windows Server 2008 and Window Vista support VPN connections over IPv6. VPN connections over IPv6 require native IPv6 support for VPN protocols on the VPN client and VPN server, IPv6 routing support on the VPN server, and connections to the IPv6 Internet.

Native IPv6 capability for VPN connections, which is the ability to send native IPv6 packets over a VPN connection, is possible with native IPv6 traffic inside the VPN tunnel and with VPN connections over IPv6.



**Note** Windows XP and Windows Server 2003 do not support VPN connections over IPv6 or native IPv6 capability for VPN connections.

## Authentication Methods

To authenticate the user who is attempting a VPN connection, Windows Server 2008 supports a wide variety of authentication protocols, including the following:

- MS-CHAP v2
- EAP-MS-CHAP v2
- EAP-TLS
- PEAP-MS-CHAP v2
- PEAP-TLS



**Note** In Windows Server 2008 and Windows Vista, support for the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP and also known as MS-CHAP v1), Shiva Password Authentication Protocol (SPAP), and EAP-Message Digest 5 (EAP-MD5) protocols has been removed because of security considerations.

EAP-TLS and PEAP-TLS are used in conjunction with a PKI and either user certificates or smart cards. With EAP-TLS, the VPN client sends its user certificate for authentication, and the authentication server sends a computer certificate for authentication. By default, the VPN client validates the VPN server's certificate. With PEAP-TLS, the VPN client and authentication server create an encrypted TLS channel, and then the VPN client and authentication server exchange certificates. Both EAP-TLS and PEAP-TLS are much stronger than either PEAP-MS-CHAP v2 or MS-CHAP v2 because they do not rely on passwords. PEAP-TLS is the strongest authentication method because the certificate exchange between the VPN client and the authentication server is encrypted.

In the absence of user certificates or smart cards, use PEAP-MS-CHAP v2, EAP-MS-CHAP v2, or MS-CHAP v2. PEAP-MS-CHAP v2 is recommended over either MS-CHAP v2 or EAP-MS-CHAP v2 because the MS-CHAP v2 exchange of messages is protected with an encrypted TLS channel, making it much more difficult for a malicious user to capture the message exchange and determine the user's password through an offline dictionary attack.

## Design Choices for Authentication Protocols

- MS-CHAP v2, EAP-MS-CHAP v2, and PEAP-MS-CHAP v2 are password-based authentication protocols.
- EAP-TLS and PEAP-TLS are certificate-based authentication protocols.
- For L2TP/IPsec-based connections, any user-level authentication protocol can be used because the authentication occurs after the VPN client and VPN server have established an IPsec-protected channel. However, the use of PEAP-MS-CHAP v2, MS-CHAP v2,

EAP-MS-CHAP v2, EAP-TLS, or PEAP-TLS is recommended to provide strong user authentication and mutual authentication with the authentication server.

## Requirements for Authentication Protocols

- For encrypted PPTP-based connections, you must use MS-CHAP v2, EAP-MS-CHAP v2, PEAP-MS-CHAP v2, EAP-TLS, or PEAP-TLS. Only these authentication protocols provide a mechanism to generate a per-session initial encryption key that is used by both the VPN client and the VPN server to encrypt PPTP data sent on the VPN connection.
- PEAP-MS-CHAP v2 and EAP-MS-CHAP v2 are supported by VPN clients running Windows Server 2008 and Windows Vista. MS-CHAP v2 is supported by VPN clients running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP.
- PEAP-MS-CHAP v2 requires the installation of a computer certificate on the authentication server (either the VPN server or, more typically, a RADIUS server) and the root CA certificate of the issuing CA of the computer certificate on each of the VPN client computers. PEAP-MS-CHAP v2 is supported only by VPN clients running Windows Server 2008 or Windows Vista.
- For SSTP-based connections, you must use MS-CHAP v2, EAP-MS-CHAP v2, PEAP-MS-CHAP v2, EAP-TLS, or PEAP-TLS. Only these authentication protocols provide a mechanism to generate a per-session initial encryption key that is used by both the VPN client and the VPN server to avoid attacks on the SSTP-based VPN connection by malicious users between the VPN client and server.
- To deploy VPN enforcement with NAP, you must use a PEAP-based authentication method.

## Best Practices for Authentication Protocols

- Use the strongest authentication scheme that is possible for your remote access VPN configuration. The strongest authentication scheme is the use of PEAP-TLS or EAP-TLS with certificates. Otherwise, use PEAP-MS-CHAP v2, MS-CHAP v2, or EAP-MS-CHAP v2 authentication.
- If you are using smart cards or have a PKI that issues user certificates, use PEAP-TLS or EAP-TLS for your VPN connections. PEAP-TLS is supported by VPN clients running Windows Server 2008 or Windows Vista. EAP-TLS is supported by VPN clients running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP.
- If you must use a password-based authentication protocol such as PEAP-MS-CHAP v2, MS-CHAP v2, or EAP-MS-CHAP v2, require the use of strong passwords on your network. Strong passwords are long (longer than eight characters) and contain a mixture of uppercase and lowercase letters, numbers, and punctuation. In an Active Directory domain, use the Computer Configuration\Windows Settings\Security Settings\Account Policies>Password Policy node in Group Policy settings to enforce strong user password requirements.

## VPN Servers

A VPN server is a computer running Windows Server 2008 and Routing and Remote Access that does the following:

- Listens for PPTP connection attempts, IPsec negotiations for L2TP connection attempts, and SSL negotiations for SSTP connection attempts
- Requires authentication and authorization of VPN connections before allowing intranet data to flow to and from VPN clients
- Acts as a router forwarding packets between VPN clients and resources on the intranet

A VPN server typically has two or more installed network adapters: one or more network adapters connected to the Internet and one or more network adapters connected to the intranet.

### Configuring Routing and Remote Access

When you configure and enable Routing and Remote Access, the Routing and Remote Access Server Setup Wizard prompts you to select the role that the computer will fill. For VPN servers, you should select the Remote Access (Dial-Up Or VPN) configuration option. For more information about the Routing and Remote Access Server Setup Wizard, see “Deploying VPN Servers” later in this chapter. With the Remote Access (Dial-Up Or VPN) option, the Routing and Remote Access server operates in the role of a dial-up or VPN server that supports remote access VPN connections.

When you select the Remote Access (Dial-Up Or VPN) option in the Routing and Remote Access Server Setup Wizard, the following occurs:

1. You are first prompted to specify whether VPN, dial-up, or both types of access are needed.
2. Next, you are prompted to select the network interface that is connected to the Internet. By default, the interface that you select will be automatically configured with IPv4 and IPv6 packet filters that allow only VPN-related traffic. All other traffic is silently discarded.

For example, you will no longer be able to ping the Internet interface of the VPN server. If you are running other services on the VPN server (such as Internet Information Services), you must manually add packet filters and exceptions for Windows Firewall to allow the traffic to and from the other services.

3. Next, if you have multiple network adapters that are connected to the intranet, you are prompted to select an interface over which DHCP, DNS, and WINS configuration is obtained.

4. Next, you are prompted to decide whether you want to obtain IPv4 addresses to assign to remote access clients by using either DHCP or a specified range of addresses. If you select a specified range of addresses, you are prompted to add the address ranges.
5. Next, you are prompted to specify whether you want to use RADIUS for authentication and accounting of VPN connections. If you select RADIUS, you are prompted to configure primary and alternate RADIUS servers and the RADIUS shared secret.

When you select and configure the Remote Access (Dial-Up Or VPN) option in the Routing and Remote Access Server Setup Wizard, the configuration results are as follows:

- The Routing and Remote Access service is enabled as an IPv4-based remote access server and LAN and demand-dial router, which performs authentication and accounting either locally or through RADIUS. If there is only one network adapter connected to the intranet, that network adapter is automatically selected as the interface from which to obtain DHCP, DNS, and WINS configuration. Otherwise, the network adapter specified in the wizard is selected to obtain DHCP, DNS, and WINS configuration. If specified, the static IPv4 address ranges are configured.
- Depending on the version of Windows Server 2008, up to 128 PPTP ports, 128 L2TP ports, and 128 SSTP ports are created. Each port represents a possible remote access VPN connection. All of them are enabled for both inbound remote access connections and inbound and outbound demand-dial connections (used for site-to-site VPN connections).
- The selected Internet interface is configured with input and output IPv4 and IPv6 packet filters that allow only VPN traffic.
- The DHCP Relay Agent component is added with the Internal interface. The Internal interface is a logical interface that represents the connection to all other authenticated remote access clients. If the VPN server is a DHCP client at the time the wizard is run, the DHCP Relay Agent is automatically configured with the IPv4 address of a DHCP server. Otherwise, you must manually configure the properties of the DHCP Relay Agent with an IPv4 address of a DHCP server on your intranet. IPv4-based remote access clients send a DHCPInform message to obtain additional configuration settings, such as DNS settings and static routes. The DHCP Relay Agent forwards DHCPInform messages between VPN remote access clients and an intranet DHCP server.
- The IGMP component is added and the Internal interface is configured for Internet Group Management Protocol (IGMP) router mode. All other LAN interfaces are configured for IGMP proxy mode. If your intranet is IPv4 multicast-enabled, this allows VPN remote access clients to send and receive IPv4 multicast traffic.

## Design Choices for VPN Servers

- The VPN server can be configured to obtain IPv4 addresses from DHCP or from a manually configured set of address ranges (known as *static pools* of addresses). Using

DHCP to obtain IPv4 addresses simplifies configuration; however, you must ensure that the DHCP scope for the subnet to which the intranet connection of the VPN server is attached has enough addresses for all the computers physically connected to the subnet and the maximum number of remote access clients.

If you are configuring a static pool of addresses, there might be additional routing considerations. For more information, see “Configuring Intranet Network Infrastructure” later in this chapter.

- The VPN server can either evaluate authentication and authorization for VPN connections itself or rely on a RADIUS server. When configuring the VPN server, you can choose to use Windows or RADIUS for authentication or accounting.

When configured to use Windows for authentication and accounting, the VPN server is a member of an Active Directory domain and communicates with an Active Directory domain controller to validate the credentials of the VPN client and obtain the VPN client’s user-account dial-in properties. The VPN server uses the user-account properties and locally configured network policies to authorize the VPN connection. The VPN server by default logs VPN connection accounting information in local accounting log files.

When configured to use RADIUS for authentication and accounting, the VPN server uses a configured RADIUS server to validate the credentials of the VPN client, authorize the connection attempt, and log VPN connection accounting information. In this configuration, the VPN server need not be a member of an Active Directory domain. If the RADIUS server is a computer running Windows Server 2008 and Network Policy Server (NPS), it must be a member of an Active Directory domain.

- The Routing and Remote Access Server Setup Wizard does not automatically enable IPv6 support for remote access VPN connections. For more information, see “Deploying VPN Servers” later in this chapter.

## Requirements for VPN Servers

- The VPN server must have a manual TCP/IP (IPv4) configuration for its Internet interface and intranet interfaces. Because of possible default route conflicts, you should manually configure your intranet interfaces with an IPv4 address, subnet mask, DNS servers, and WINS servers. However, do not configure a default gateway on the VPN server’s intranet interfaces. It is possible for the VPN server to have a manual TCP/IP configuration and use DHCP to obtain IPv4 addresses that are assigned to VPN clients.
- For VPN connections that use the PEAP-MS-CHAP v2, EAP-TLS, or PEAP-TLS authentication protocols, you must install on the authentication server (either the VPN server or the RADIUS server) a computer certificate that can be validated by the VPN client. You might also need to install the root CA certificate of the issuing CA of the authentication server’s computer certificate on your VPN client.



- For SSTP-based VPN connections, you must install on the VPN server a computer certificate that can be validated by the VPN client. You might also need to install the root CA certificate of the issuing CA of the VPN server's computer certificate on your VPN client.
- For L2TP/IPsec-based VPN connections, you must install on the VPN server a computer certificate that can be validated by the VPN client.
- If you configure the VPN server for local authentication or for RADIUS authentication, and the RADIUS server is a computer running NPS, the default network policy named Connections to Microsoft Routing and Remote Access server rejects all types of connection attempts unless the remote access permission of the user account's dial-in properties is set to Allow Access. If you want to use this network policy for your VPN connections, set the policy type to Allow Access. If you want to manage authorization and connection settings for VPN connections by group or by type of connection, you must configure additional NPS policies. For more information, see "Configuring RADIUS Servers" later in this chapter.

## Best Practices for VPN Servers

- Determine the connection of the VPN server that will be connected to the Internet. Typical Internet-connected VPN servers have at least two LAN connections: one connected to the Internet (either directly or connected to a perimeter network) and one connected to the organization intranet. To make this distinction easier to see when using the Routing and Remote Access Server Setup Wizard, in the Network Connections folder, rename the connections to a name that describes their purpose or role. For example, rename the connection named "Local Area Connection 2" that is connected to the Internet with the name "Internet."

## Internet Infrastructure

For a VPN client to successfully exchange traffic with a VPN server over the Internet, the following must be true:

- The VPN server's DNS name or IP address is reachable.
- The VPN server is reachable.
- VPN traffic is allowed to and from the VPN server.

## VPN Server Name Resolvability

In most cases, you will refer to the VPN server by its fully qualified domain name (FQDN) rather than its IPv4 or IPv6 address. You can use an FQDN (for example, vpn.example.microsoft.com) as long as the name can be resolved to an IPv4 or IPv6 address. Therefore, you must ensure that whatever name you are using for your VPN servers when configuring a VPN connection is resolvable to an IPv4 or IPv6 address using Internet-based DNS servers.

When you use names rather than addresses, you can also take advantage of DNS round-robin load balancing if you have multiple VPN servers with the same DNS host name. Within DNS, you can create multiple records that resolve a specific host name to different IPv4 addresses. In this situation, DNS servers send back all the addresses in response to a DNS name query and typically randomize the order of the addresses for successive queries. Because most DNS clients use the first address in the DNS query response, the result is that VPN client connections are, on average, spread across the VPN servers, as long as both VPN servers are available. To ensure availability of the VPN servers, you can use Network Load Balancing.

## VPN Server Reachability

To be reachable, the VPN server must be assigned a public IPv4 address or a global IPv6 address to which packets are forwarded by the routing infrastructure of the IPv6 or IPv6 Internet. If you have been assigned a static public IPv4 address or global IPv6 address prefix from an ISP or an Internet registry, this is typically not an issue. In some IPv4 configurations, the VPN server is actually configured with a private IPv4 address and has a published static IPv4 address by which it is known on the Internet. A device between the Internet and the VPN server translates the published and actual IPv4 addresses of the VPN server in packets to and from the VPN server.

Although the routing infrastructure might provide reachability, the VPN server might be unreachable because of the placement of firewalls, packet filtering routers, NATs, security gateways, or other types of devices that prevent packets from either being sent to or received from the VPN server computer.

## VPN Servers and Firewall Configuration

There are two approaches to using a firewall with a VPN server:

- **The VPN server is attached directly to the Internet, and the firewall is between the VPN server and the intranet.** In this configuration, the VPN server must be configured with packet filters that allow VPN traffic in and out of its Internet interface only. The firewall can be configured to allow specific types of remote access traffic.
- **The firewall is attached to the Internet, and the VPN server is between the firewall and the intranet.** In this configuration, both the firewall and the VPN server are attached to a subnet known as the *perimeter network* (also known as a *screened subnet*). Both the firewall and the VPN server must be configured with packet filters that allow only VPN traffic to and from the Internet.

For the details of configuring packet filters for the VPN server and the firewall for both of these configurations, see “Firewall Packet Filtering for VPN Traffic” later in this chapter.

## Requirements for Internet Infrastructure

- Ensure that the FQDNs of your VPN servers are resolvable from the Internet by placing either appropriate DNS address (A) or IPv6 address (AAAA) records in your Internet DNS server or the DNS server of your ISP. Test the resolvability by using the Ping tool to ping the name of each of your VPN servers when directly connected to the IPv4 or IPv6 Internet.

Because of packet filtering, the ping command might display the result “Request timed out,” but check to ensure that the name specified was resolved by the Ping tool to the proper address. To force the Ping tool to use an IPv4 address, use the **-4** command-line switch. To force the Ping tool to use an IPv6 address, use the **-6** command-line switch. You can also use the Nslookup tool to test name resolution.

- Ensure that the IPv4 or IPv6 addresses of your VPN servers are reachable from the Internet by using the Ping tool to ping the FQDN or address of your VPN server with a 5-second timeout (using the **-w 5** command-line switch) when directly connected to the Internet. If you see a “Destination unreachable” error message, the VPN server is not reachable.

## Best Practices for Internet Infrastructure

Configure packet filtering for PPTP traffic, L2TP traffic, SSTP traffic, or all types of traffic on the appropriate firewall and VPN server interfaces connecting to the Internet and the perimeter network. For more information, see “Firewall Packet Filtering for VPN Traffic” later in this chapter.

## Intranet Infrastructure

The intranet infrastructure ensures that the VPN client can exchange packets with nodes on the intranet using the VPN server as an IPv4 or IPv6 router. Without proper intranet infrastructure design, VPN clients might be unable to do the following:

- Resolve intranet names
- Obtain an IPv4 address or IPv6 subnet prefix that is reachable from the intranet
- Reach intranet locations

## Intranet Name Resolution

Ensure that each VPN server is configured with the IPv4 or IPv6 addresses of your intranet DNS servers and, if you are using WINS to resolve intranet NetBIOS names, the IPv4 addresses of your intranet WINS servers. The VPN server should be manually configured with DNS and WINS servers.

As part of the PPP connection negotiation process for IPv4, the VPN clients receive IPv4 addresses of DNS and WINS servers. By default, the VPN clients inherit the DNS and WINS server addresses configured on the VPN server. After the PPP connection negotiation is complete, a VPN client running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP sends a DHCPInform message to the VPN server. If properly configured, the VPN server forwards the DHCPInform message to a DHCP server, which responds with a DHCPAck message. The VPN server sends the DHCPAck message to the VPN client, which can contain a DNS domain name, additional DNS server addresses for DNS servers (which are queried before the DNS servers configured through the PPP negotiation), and WINS server addresses (which replace the WINS server addresses configured through the PPP negotiation). The relaying of DHCP messages is facilitated by the DHCP Relay Agent routing protocol component of Routing and Remote Access, which is automatically added by the Routing and Remote Access Server Setup Wizard.

If the VPN server is using DHCP to configure its intranet interfaces (not recommended), the VPN server relays the DHCPInform messages to the DHCP server that was in use when the Routing and Remote Access Server Wizard was run. If the VPN server has a static TCP/IP configuration on its intranet interface (recommended), the DHCP Relay Agent routing protocol component must be configured with the IPv4 address of at least one DHCP server on your intranet. You can add DHCP server IPv4 addresses to the DHCP Relay Agent routing protocol component in Routing and Remote Access snap-in from the General tab for the properties of the DHCP Relay Agent item under IPv4 Routing.

To dynamically configure the IPv6 addresses of DNS servers for VPN connections that support native IPv6 traffic, the Windows Vista-based or Windows Server 2008-based VPN client relies on the Router Advertisement message sent by the VPN server after IPV6CP negotiation completes. If the Router Advertisement message has the Other Stateful Configuration flag (the O flag) set, the VPN client sends a DHCPv6 Information-Request message to the VPN server. If the Windows Server 2008 VPN server is properly configured with the DHCPv6 Relay Agent, it will forward the Information-Request message to a DHCPv6 server. The DHCPv6 Reply message is forwarded back to the VPN client and can contain the IPv6 addresses of DNS servers on the intranet.

### Requirements for Intranet Name Resolution

- Using the Ping and Net tools, test DNS and WINS name resolution for intranet resources from the VPN server computer. If name resolution does not work from the VPN server, it might not work for VPN clients. Troubleshoot and fix all name resolution problems of the VPN server before testing VPN connections.
- Because the intranet interfaces of the VPN server are manually configured with a TCP/IP configuration, the Routing and Remote Access Server Setup Wizard cannot automatically configure the DHCP Relay Agent routing protocol component. You must manually add the IPv4 address of at least one DHCP server on your intranet to the DHCP Relay

Agent component. If you do not, the VPN server discards DHCPInform messages sent by VPN clients, and the VPN clients do not receive the updated DNS and WINS server addresses or the DNS domain name.

- If you have a single-subnet small office/home office (SOHO) with no DHCP, DNS, or WINS server, you must either configure a DNS server or WINS server to resolve names for computers on the SOHO subnet and VPN clients or enable NetBIOS broadcast name resolution, which enables NetBIOS over TCP/IP name resolution between connected VPN clients and computers on the SOHO network. To enable NetBIOS broadcast name resolution, in the Routing and Remote Access snap-in, in the properties dialog box of a VPN server, on the IPv4 tab, select the Enable Broadcast Name Resolution check box.
- To forward DHCPv6 messages between IPv6-capable VPN clients and a DHCPv6 intranet server, you must add and configure the DHCPv6 Relay Agent routing protocol component. For more information, see “Deploying VPN Servers” later in this chapter.

**Best Practices for Intranet Name Resolution** To ensure that VPN clients obtain the most current list of DNS and WINS server IPv4 addresses, manually configure the DHCP Relay Agent component of Routing and Remote Access rather than relying on the VPN server to configure VPN clients with its own DNS and WINS server IPv4 addresses.

## VPN Server Routing to the Internet and the Intranet

The VPN server is an IPv4 and IPv6 router that forwards packets between VPN clients and nodes on the intranet. Therefore, it must be configured with the proper set of routes to be able to reach any Internet location (because a VPN client can connect from anywhere on the Internet) and any intranet location. For both IPv4 and IPv6 traffic, the VPN server needs the following:

- A default route that points to a firewall or router directly connected to the Internet. This route makes all the locations on the Internet reachable.
- One or more routes that summarize the addresses used on your intranet and point to a neighboring intranet router. These routes make all the locations on your intranet reachable from the VPN server. Without these routes, all intranet hosts not connected to the same intranet subnet as the VPN server are unreachable.

For a default route that points to the Internet, configure the Internet interface of the VPN server with a default gateway, but do not configure the intranet interfaces with a default gateway. If you configure your intranet interfaces with default gateways, you will have multiple default routes in the IPv4 and IPv6 routing tables of the VPN server. Because of the way that the TCP/IP protocol selects the default route for forwarding default route traffic, having multiple default routes can result in default route traffic being forwarded to the intranet, rather than the Internet, making Internet locations unreachable.

To add intranet routes to the routing table of the VPN server, you can:

- Add IPv4 and IPv6 static routes by using the Routing and Remote Access snap-in. You do not necessarily need to add a route for each subnet in your intranet. At a minimum, you need to add the routes that summarize the IPv4 or IPv6 address space used on your intranet.

For example, if your intranet uses the private IPv4 address space 10.0.0.0/8 to number its subnets and hosts, you do not need to add a route for each subnet. Just add a route for 10.0.0.0 with the subnet mask 255.0.0.0 that points to a neighboring router on the intranet subnet to which your VPN server is attached.

- If you are using Routing Information Protocol (RIP) in your intranet, you can add and configure the RIP component of the Routing and Remote Access service so that the VPN server participates in the propagation of intranet routing information as a RIP-based router.

If your intranet has only a single subnet, no further configuration is required.

## VPN Client Routing to the Intranet

The reachability of VPN clients from the intranet for IPv4 traffic depends on how you configure the VPN server to obtain IPv4 addresses to assign to VPN clients. The IPv4 addresses assigned to VPN clients as they connect can be from either of the following:

- An *on-subnet address range*, which is an address range of the intranet subnet to which the VPN server is attached.

An on-subnet address range is used whenever the VPN server is configured to use DHCP to obtain IP addresses for VPN clients or when the manually configured pools of IPv4 addresses are within the range of addresses of the attached subnet.

- An *off-subnet address range*, which is an address range that represents a different subnet that is logically attached to the VPN server.

An off-subnet address range is used whenever the VPN server is manually configured with pools of IPv4 addresses for a separate subnet.

If you are using an on-subnet address range, no additional routing configuration is required because the VPN server acts as an Address Resolution Protocol (ARP) proxy for all packets destined for VPN clients. Routers and hosts on the VPN server subnet forward packets destined to VPN clients to the VPN server, which sends them the appropriate VPN client.

If you are using an off-subnet address range, you must add the route(s) that summarize the off-subnet address range to the intranet routing infrastructure so that traffic destined for VPN clients is forwarded to the VPN server and then sent by the VPN server to the appropriate VPN client. To provide the best summarization of address ranges for routes, choose address ranges that can be expressed using a single prefix and subnet mask.

To add the routes that summarize the off-subnet address range to the routing infrastructure of the intranet, add static routes to a neighboring router of the VPN server for the off-subnet address range that point to the VPN server's intranet interface. Configure the neighboring router to propagate this static route to other routers in the intranet using the dynamic routing protocol used in your intranet.

If your intranet consists of a single subnet, you must either configure each intranet host for persistent route(s) of the off-subnet address range that point to the VPN server's intranet interface or configure each intranet host with the VPN server as its default gateway. Therefore, it is recommended that you use an on-subnet address pool for a SOHO network consisting of a single subnet.

For IPv6-based VPN connections, the subnet prefix assigned to VPN clients in the Router Advertisement message is always for a subnet separate from the subnet to which the VPN server is connected. All the VPN clients are assigned the same subnet prefix, which is always an off-subnet prefix. To make the VPN clients reachable from the intranet, you must add the subnet prefix as a route pointing to the VPN server to your IPv6 routing infrastructure.

## Requirements for Intranet Routing Infrastructure

- Configure the Internet interface of the VPN server with a default gateway, but do *not* configure the intranet interfaces of the VPN server with a default gateway.
- Add static IPv4 and IPv6 routes that summarize the addresses used in your intranet to the VPN server. Alternatively, if you use RIP for your IPv4 dynamic routing protocol, configure and enable RIP on the VPN server. If you use a routing protocol other than RIP, you might be able to use route redistribution. For example, if you use Interior Gateway Routing Protocol (IGRP), you might configure the VPN server's neighboring intranet router to use RIP on the interface connected to the subnet to which the VPN server is attached and IGRP on all other interfaces.
- Add the IPv6 subnet prefix for IPv6-capable VPN clients as a route pointing to the VPN server to your IPv6 routing infrastructure.

## Best Practices for Intranet Routing Infrastructure

If possible, configure the VPN server with an on-subnet address range either by obtaining IPv4 addresses through DHCP or by manually configuring on-subnet address pools.

## Concurrent Intranet and Internet Access for VPN Clients

By default, when a Windows-based VPN client makes a VPN connection, it automatically adds a new default route for the VPN connection and modifies the existing default route to have a higher metric. Adding the new default route means that all Internet locations except the IPv4 address of the VPN server and locations based on other routes are not reachable for the duration of the VPN connection.

To prevent the default route from being created, you can configure the VPN connection to not use the default gateway of the remote network. For VPN connections in the Network Connections folder, do the following:

1. Obtain properties of the Internet Protocol (TCP/IP) or Internet Protocol Version 4 (TCP/IPv4) component from the Networking tab for the properties of the VPN connection.
2. Click Advanced.
3. In the Advanced TCP/IP Settings dialog box, on the General tab, clear the Use Default Gateway On Remote Network check box.

When the Use Default Gateway On Remote Network check box is cleared, a default route is not created when the connection is made. However, a route corresponding to the Internet address class of the assigned IPv4 address is created. For example, if the address assigned during the connection process is 10.0.12.119, the Windows-based VPN client creates a route for the class-based address prefix 10.0.0.0 with the subnet mask 255.0.0.0.

Based on the Use Default Gateway On Remote Network setting, one of the following occurs when the VPN connection is active:

- Internet locations are reachable and intranet locations are not reachable except those matching the address class of the assigned IP address. (The Use Default Gateway On Remote Network check box is cleared.)
- All intranet locations are reachable and Internet locations are not reachable except the address of the VPN server and locations available through other routes. (The Use Default Gateway On Remote Network check box is selected.)

For most Internet-connected VPN clients, this behavior does not represent a problem because they are typically engaged in either intranet or Internet communication, not both.

For VPN clients who want concurrent access to intranet and Internet resources when the VPN connection is active (also known as *split tunneling*), you can do one of the following:

- Select the Use Default Gateway On Remote Network check box (the default setting), and allow Internet access through the organization intranet. Internet traffic between the VPN client and Internet hosts would pass through firewalls or proxy servers as if the VPN client is physically connected to the organization intranet. Although there is an impact on performance, this method allows Internet access to be filtered and monitored according to the organization's network policies while the VPN client is connected to the organization network.
- If the IPv4 addressing within your intranet is based on a single class-based address prefix, clear the Use Default Gateway On Remote Network check box. An example is when your intranet is using the private IPv4 address prefix 10.0.0.0/8.



- If the IPv4 addressing within your intranet is not based on a single class-based address prefix, you can use the following solutions:
  - ❑ The Classless Static Routes DHCP option
  - ❑ The Connection Manager Administration Kit
  - ❑ A command (.cmd) file on the VPN client

For more information about these methods, see the section “Configuring Concurrent Access to the Internet and Intranet” later in this chapter.



**Note** For native IPv6-based VPN clients, the default route is added based on the receipt of the Router Advertisement from the VPN server. If the Use Default Gateway On Remote Network check box is selected for the TCP/IPv6 protocol, the interface metric of the VPN connection—which becomes the metric of the default IPv6 route that uses the VPN connection—is set to a low value so that the default route over the VPN connection has the lowest metric. If the Use Default Gateway On Remote Network check box is cleared, the interface metric of the VPN connection is set to a static value or to an automatic metric, but it is not guaranteed to be lowest. Therefore, the default route over the VPN connection might not have the lowest metric.

## Authentication Infrastructure

The authentication infrastructure exists to:

- Authenticate the credentials of VPN clients
- Authorize the VPN connection
- Record the VPN connection creation and termination for accounting purposes

The authentication infrastructure for remote access VPN connections consists of:

- The VPN server computer
- A RADIUS server computer
- A domain controller
- An issuing CA of a PKI (optional)

## Using Windows or RADIUS for Authentication

A Windows Server 2008-based VPN server can be configured to use either Windows or RADIUS for authentication or accounting. RADIUS provides a centralized authentication, authorization, and accounting service when you have multiple VPN servers or a mix of heterogeneous dial-up and VPN equipment or other types of access servers such as wireless access points.

When the VPN server uses Windows for authentication, it performs the authentication of the VPN connection by communicating with a domain controller using a protected remote procedure call (RPC) channel and performs authorization of the connection attempt through the dial-in properties of the user account and locally configured network policies. When the VPN server uses RADIUS for authentication, it relies on a RADIUS server to perform both the authentication and authorization.

When the VPN server uses Windows for authentication, it logs VPN connection information in a local log file (`%SystemRoot%\System32\Logfiles\Logfile.log` by default) based on settings configured in the Accounting node in the Network Policy Server snap-in. When the VPN server uses Windows for authentication, it relies on the RADIUS server to record the accounting information.

If you are using RADIUS and a Windows domain as the user account database from which to verify user credentials and obtain dial-in properties, you should use NPS in Windows Server 2008. NPS is a full-featured RADIUS server and proxy that is tightly integrated with Active Directory and Routing and Remote Access.

When NPS is used as the RADIUS server, it does the following:

- NPS performs the authentication of the VPN connection by communicating with a domain controller over a protected RPC channel. NPS performs authorization of the connection attempt through the dial-in properties of the user account and network policies configured on the NPS server.
- NPS by default logs all RADIUS accounting information in a local log file (`%SystemRoot%\System32\Logfiles\Logfile.log` by default) based on settings configured in the Accounting node in the Network Policy Server snap-in.

## Best Practices for Authentication Infrastructure

- If you have multiple VPN servers and you want to centralize authentication, authorization, and accounting services, or you have a heterogeneous mixture of network access equipment, use a RADIUS server and configure the VPN server to use RADIUS for authentication and accounting.
- If your user account database is the Active Directory domain service, use NPS as your RADIUS server. See Chapter 9 for additional design and planning considerations for NPS-based RADIUS servers.
- To better manage authorization for remote access VPN connections, create a universal group in Active Directory for VPN access that contains global groups for the user accounts that are allowed to make remote access VPN connections. For example, create a universal group named `VPNUsers` that contains the global groups based on your organization's regions or departments. Each global group contains allowed user accounts for VPN remote access. When you configure your NPS policies for VPN connections, you specify the `VPNUsers` group name.

- Whether the VPN server is configured for local or RADIUS-based authentication, use a VPN-specific network policy to authorize VPN connections and specify connection constraints and requirements. For example, use network policies to grant access based on group membership, to require strong encryption, to require the use of specific authentication methods (such as PEAP-MS-CHAP v2 or EAP-TLS), or to limit traffic by using IP packet filtering.

## VPN Clients

The VPN client can be any computer that is capable of creating a PPTP connection using MPPE encryption, an L2TP connection using IPsec encryption, or an SSTP connection using SSL encryption. A Windows-based VPN client running Windows Vista, Windows Server 2008, Windows Server 2003, or Windows XP can create PPTP or L2TP/IPsec-based VPN connections. A Windows-based VPN client running Windows Vista SP1 or Windows Server 2008 can create SSTP-based VPN connections.

You can configure VPN connections on the Windows-based VPN client either manually or by using the Connection Manager components available in Windows Server 2008. The exact method for manually configuring VPN connections varies in Windows. The methods for different versions are as follows:

- For a Windows Vista-based or Windows Server 2008-based VPN client, in the Network and Sharing Center, click Connect to a Network. To create a VPN connection, you must specify the IP address or DNS name of the VPN server on the Internet.
- For a Windows XP-based or Windows Server 2003-based VPN client, use the New Connection Wizard in the Network Connections folder.

## Connection Manager

When scaling the configuration of VPN connections for an enterprise network, you might encounter the following issues:

- The exact procedure to configure a VPN connection varies depending on the version of Windows running on the client computer.
- To prevent configuration errors, it is preferable to have the IT staff, rather than users, configure the VPN connection.
- A configuration method must be able to scale to hundreds or thousands of client computers in a large organization.
- A VPN connection might need a double-dial configuration, in which a user must obtain a dial-up connection to the Internet before creating a VPN connection with the organization intranet.

The solution to these issues of configuring VPN connections across an enterprise is Connection Manager. Connection Manager consists of the following components:

- Connection Manager client dialer
- Connection Manager Administration Kit
- Connection Point Services

**Connection Manager Client Dialer** The Connection Manager (CM) client dialer is software that is installed on each VPN client. It includes advanced features that make it a superset of basic remote access networking. At the same time, the CM client dialer presents a simplified connection experience to the user. It limits the number of configuration options that a user can change, ensuring that the user can always connect successfully. For example, a CM client dialer can:

- Use customized graphics, icons, messages, and help
- Automatically create a dial-up connection before the VPN connection is made
- Run custom actions during various parts of the connection process, such as pre-connect and post-connect actions (executed before or after the dial-up or VPN connection is completed)
- For dial-up connections, select from a list of phone numbers to use, based on physical location

A *customized CM client dialer profile*, also known as a *package*, is a self-extracting executable file that is created by a network administrator using the Connection Manager Administration Kit (CMAK). The CM profile is distributed to VPN users via CD-ROM, e-mail, Web site, or file share. When the user runs the CM profile, it automatically configures the customized dial-up or VPN connection. The CM profile does not require a specific version of Windows; it configures connections for computers running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP.

**Connection Manager Administration Kit** You create a customized CM profile by using the CMAK. With the CMAK, you can develop client dialer and connection software that allows your users to connect to the network by using only the connection features that you specify for them. The CM profile supports a variety of features that both simplify and enhance implementation of connection support for you and your users, most of which can be incorporated using the CMAK. The CMAK allows you to build CM profiles customizing the CM client dialer so that connection reflects the identity of your organization. It allows you to determine which functions and features you want to include and how the dial-up or VPN connection appears to your users.

**Connection Point Services** For dial-up CM profiles, Connection Point Services (CPS) allows you to automatically distribute and update custom phone books. These phone books

contain one or more Point of Presence (POP) entries, with each POP supplying a telephone number that provides dial-up access to an intranet, or more commonly, to an Internet access point. The phone books give users complete POP information so that when they travel, they can connect to different Internet access points rather than being restricted to a single POP.

Without the ability to update phone books (a task CPS handles automatically), users would need to contact their organization's technical support staff to be informed of changes in POP information and to reconfigure their client dialer software.

CPS has two components:

- **Phone Book Administrator** A tool to create and maintain the phone book database and to publish new phone book information to the Phone Book Service.
- **Phone Book Service** A Microsoft Internet Information Services (IIS) 7.0 extension. A CM profile can be configured to check the Phone Book Service running on a specified IIS server to ensure that it is using the latest phone book. If not, the remote access client automatically downloads a phone book update.

## Design Choices for VPN Clients

- If you have a small number of VPN clients, you can perform manual configuration of VPN connections on each computer.
- If you have a large number of VPN clients or they are running different versions of Windows, use the Connection Manager components of Windows Server 2008 to create a CM profile containing customized VPN configuration settings, and for dial-up connections, to maintain the phone book database.

## Requirements for VPN Clients

- For L2TP/IPsec connections, you must install a computer certificate on the VPN client computer.
- For the PEAP-TLS or EAP-TLS authentication methods, you must either install a user certificate on the VPN client computer or issue smart cards to your users.
- For SSTP connections, you must ensure that the VPN clients have the root CA certificate of the issuing CA of the VPN server's computer certificate installed.
- For the PEAP-MS-CHAP v2 or PEAP-TLS authentication methods, if your VPN clients are validating the certificate of the authentication server (recommended), you must ensure that the VPN clients have the root CA certificate of the issuing CA of the authentication server's computer certificate installed.

## Design Choices for Connection Manager Profiles

- The name of the CM profile should reflect its purpose and use because it will be the name of the connection in the Network Connections folder after it is installed on the VPN client.
- You can merge settings from existing profiles into new profiles. The new profiles inherit the settings of the merged profiles.
- The VPN client computer might need to make a dial-up connection to obtain Internet access before attempting the VPN connection.
- If VPN clients need concurrent access to the Internet and the intranet, you can configure the CM profile to add static routes to the VPN client's routing table for intranet locations. For more information, see "Configuring Concurrent Access to the Internet and Intranet" later in this chapter.
- The CM profile can be configured to automatically configure the VPN client with the Internet Explorer proxy settings for the intranet proxy servers.
- If you need to run programs such as disabling certain services or launching a Windows program (known as *actions*) during various phases of the VPN connection setup, you can configure custom actions. For example, you can configure actions that run before the connection is made (a pre-connect action) or after the connection is made (a post-connect action).
- If you want to specify a custom logon picture for your organization's logo that appears when your users activate the VPN connection, create a bitmap file that is 330 by 140 pixels.
- If you want to specify a custom picture for your organization's logo that appears when your users access the phone book, create a bitmap file that is 114 by 309 pixels.
- If you want to specify a custom program and title bar icons for the VPN connection in the Network and Sharing Center or Network Connections folder, create bitmap files that are 32 by 32 pixels and 16 by 16 pixels, respectively.
- If you want to provide your users with customized help for the VPN connection, create a custom help file in compiled help module (CHM) format.
- You can add files that are installed with the CM profile, such as organization information, support, or troubleshooting tools.
- You must determine how to distribute the CM profile to your users. For more information, see "Distributing Your CM Profiles" later in this chapter.

## Requirements for Connection Manager Profiles

- If you have a mixture of VPN clients running Windows Vista, Windows Server 2008, Windows Server 2003, and/or Windows XP, you will need to create separate CM profiles for VPN clients running Windows Vista or Windows Server 2008 and for VPN clients running Windows Server 2003 and/or Windows XP.

## PKI

To perform certificate-based authentication for L2TP connections and smart card or user certificate-based authentication for VPN connections using PEAP-TLS or EAP-TLS, a PKI must be in place to issue the proper certificates to VPN clients, VPN servers, and RADIUS servers to submit during the authentication process and to validate the certificate being submitted.

For PEAP-MS-CHAP v2-based authentication and SSTP-based VPN connections, a PKI is not required. It is possible to purchase certificates from a third-party certification authority (CA) to install on your authentication server (for PEAP-MS-CHAP v2) or VPN server (for SSTP). You might also need to distribute the root CA and intermediate CA certificates of the third-party computer certificates to your VPN client computers.

### Computer Certificates for L2TP/IPsec Connections

When you are using the certificate authentication method for L2TP/IPsec connections, the list of certification authorities (CAs) is not configurable. Instead, each IPsec peer sends a list of root CAs from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A is issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies its IPsec peer that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPsec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPsec negotiation fails.

The VPN client must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

For example, if the VPN client is issued computer certificates by root CAs CertAuth1 and CertAuth2, it notifies the VPN server during IPsec security negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the VPN server does not have a valid computer certificate issued from a CA that follows a certificate chain to either CertAuth1 or CertAuth2, IPsec negotiation fails.

An organization typically has a single root CA and one or multiple issuing CAs of the root CA that issue computer certificates. Because of this, all computers within the organization have both computer certificates from an issuing CA of the single root CA *and* request certificates for authentication from issuing CAs of the same single root CA.

To deploy computer certificates for L2TP/IPsec connections in your organization, perform the following actions:

1. Deploy a PKI.

2. Install a computer certificate on each computer. This is most easily accomplished with Windows Active Directory Certificate Services or Certificate Services installed as an enterprise CA and by configuring Group Policy settings for computer certificate auto-enrollment. For more information, see “Deploying Certificates” later in this chapter.

## PKI for Smart Cards

The use of smart cards is the strongest form of user authentication in Windows Server 2008. For remote access VPN connections, you can use smart cards with the EAP-TLS or PEAP-TLS authentication method.

The individual smart cards are distributed to users who have a computer with a smart card reader. To log on to the computer, the user must insert the smart card into the smart card reader and type the smart card personal identification number (PIN). When the user attempts to make a VPN connection, the smart card certificate is sent during the connection negotiation process.

To manually configure EAP-TLS for smart cards on the VPN client:

- The VPN connection must be configured to use EAP with the Smart Card Or Other Certificate EAP type. In the properties dialog box of the Smart Card Or Other Certificate EAP type, select Use My Smart Card.
- For VPN clients running Windows Server 2008, Windows Vista, Windows Server 2003, Windows XP SP2, or Windows XP SP1, if you want to validate the computer certificate of the authentication server, select Validate Server Certificate (enabled by default). If you want to configure the names of the authentication servers, select Connect To These Servers, and then type the server names. To require the server's computer certificate to have been issued a certificate from a specific set of trusted root CAs, in the Trusted Root Certification Authorities section, select the appropriate CAs.

For instructions on configuring the Connection Manager Administration Kit so that EAP-TLS uses smart cards, see “Configuring and Deploying CM Profiles by Using the CMAK” later in this chapter.

EAP is enabled as an authentication type by default, but if it isn't enabled, in the Routing and Remote Access snap-in, in the properties dialog box of the VPN server, on the Security tab, click Authentication Methods to open its dialog box, and then enable EAP.

To configure EAP-TLS authentication in the NPS network policy for remote access VPN connections, in the properties dialog box of the network policy, ensure that EAP is enabled. On the Constraints tab, add the Smart Card Or Other Certificate EAP type to the list of EAP types for the Authentication Methods constraint. If the authentication server has multiple computer certificates installed, configure the properties of the Smart Card Or Other Certificate EAP type, and then select the appropriate computer certificate to submit during EAP-TLS authentication.



## PKI for User Certificates

User certificates that are stored in the Windows registry for user authentication can be used in place of smart cards. However, it is not as strong a form of authentication. With smart cards, the user certificate issued during the authentication process is made available only when the user possesses the smart card and has knowledge of the PIN to log on to the computer. With user certificates, the user certificate issued during the authentication process is made available when the user logs on to the computer using a domain-based user name and password.

Just as with smart cards, authentication using user certificates for remote access VPN connections use the EAP-TLS or PEAP-TLS authentication methods.

To deploy user certificates in your organization, perform the following steps:

1. Deploy a PKI.
2. Install a user certificate for each user. This is most easily accomplished with Windows Certificate Services installed as an enterprise CA and configuring Group Policy settings for user certificate autoenrollment. For more information, see “Deploying Certificates” later in this chapter.

When the user attempts a VPN connection, the VPN client computer sends the user certificate during the authentication process.

## Requirements for PKI

- For L2TP/IPsec remote access VPN connections using computer certificate authentication for IPsec, you must install computer certificates, also known as *machine certificates*, on each VPN client and VPN server.

The computer certificate of the VPN client must be valid and verifiable by the VPN server; the VPN server must have a root CA certificate for the CA that issued the computer certificate of the VPN client.

The computer certificate of the VPN server must be valid and verifiable by the VPN client; the VPN client must have a root CA certificate for the CA that issued the computer certificate of the VPN server.

- To authenticate VPN connections using a smart card or user certificate with EAP-TLS or PEAP-TLS, the VPN client must have a smart card or registry-based user certificate installed, and the authentication server must have a computer certificate installed.

The smart card or user certificate of the VPN client must be valid and verifiable by the authentication server; the authentication server must have the root CA certificate of the issuing CA of the certificate of the VPN client.

The computer certificate of the authentication server must be verifiable by the VPN client; the VPN client must have the root CA certificate of the issuing CA of the computer certificate of the authentication server.

- To authenticate VPN connections using PEAP-TLS, the authentication server must have a computer certificate installed.

The computer certificate of the authentication server must be verifiable by the VPN client; the VPN client must have the root CA certificate of the issuing CA for the computer certificate of the authentication server.

- For SSTP-based VPN connections, you must install a computer certificate on the VPN server.

The computer certificate of the VPN server must be valid and verifiable by the VPN client. The VPN client must have the root CA certificate of the issuing CA of the computer certificate that is installed on the VPN server.

## Best Practices for PKI

- For computer certificates for L2TP/IPsec, if you are using a Windows Server 2008 enterprise CA as an issuing CA, use Computer Configuration Group Policy to configure your Active Directory domain for autoenrollment of computer certificates. Each computer that is a member of the domain automatically requests a computer certificate when Computer Configuration Group Policy is updated.
- For registry-based user certificates for EAP-TLS or PEAP-TLS, if you are using a Windows Server 2008 enterprise CA as an issuing CA, use Computer Configuration Group Policy to configure your Active Directory domain for autoenrollment of computer certificates. Each user that successfully logs on to the domain automatically requests a user certificate when User Configuration Group Policy is updated.

## VPN Enforcement with NAP

Network Access Protection (NAP) for Windows Server 2008, Windows Vista, and Windows XP with Service Pack 3 provides components and an application programming interface (API) set that help you enforce compliance with health policies for network access or communication. Developers and network administrators can create solutions for validating computers that connect to their networks, can provide needed updates or access to needed resources, and can limit the access of noncompliant computers.

VPN Enforcement is one of the NAP enforcement methods included with Windows Server 2008, Windows Vista, and Windows XP with Service Pack 3. With VPN Enforcement, a VPN-based remote access client must prove that it is compliant with system health requirements before being allowed full access to the intranet. If the VPN client is not compliant with system health requirements, the VPN server places the VPN client on a restricted network containing servers that have resources to bring the VPN client back into compliance. The VPN

server enforces the restricted access through IP packet filters that are placed on the VPN connection. After correcting its health state, the VPN client validates its health state again, and if compliant, the IP packet filters on the VPN connection that confine the access to the restricted network are removed.

For VPN Enforcement to work, you must already have a working VPN deployment with Windows Server 2008–based VPN servers that uses a PEAP-based authentication method. For the details of deploying VPN Enforcement after successfully deploying a remote access VPN solution, see Chapter 18, “VPN Enforcement.”

## Additional Security Considerations

When deploying a remote access VPN solution, you must consider the following additional security considerations:

- Strong link encryption
- Packet filtering on the VPN server
- Firewall packet filtering for VPN traffic
- Multi-use VPN servers
- Preventing traffic routed from VPN clients
- Concurrent access
- Unused VPN protocols

## Strong Link Encryption

For encryption, you can use link encryption or both end-to-end encryption and link encryption, described as follows:

- *Link encryption* encrypts the data only on the link between the VPN client and the VPN server across the Internet. For PPTP connections, you must use MPPE in conjunction with MS-CHAP v2, PEAP-MS-CHAP v2, EAP-TLS, or PEAP-TLS authentication. For L2TP/IPsec connections, IPsec provides encryption. For SSTP connections, SSL provides encryption.
- *End-to-end encryption* encrypts the data between the source host and its final destination. You can use IPsec after the VPN connection is made to encrypt data between the VPN client on the Internet and the node on the intranet. For more information about IPsec and end-to-end protection of IP traffic, see Chapter 4, “Windows Firewall with Advanced Security.”

To configure the VPN server to require link encryption, select the appropriate encryption strengths for the Encryption settings on the network policy that is used for remote access VPN connections. Do not select the No Encryption check box.

## VPN Traffic Packet Filtering on the VPN Server

To prevent the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic (assuming that the VPN server is not hosting other services accessible from the Internet), you must ensure that IPv4 and IPv6 input and output packet filters for PPTP, L2TP/IPsec, and SSTP traffic are configured on the Internet interface of the VPN server. Because a VPN server is an IPv4 and IPv6 router, if the VPN traffic filters are not configured on the Internet interface, the VPN server might forward unwanted Internet traffic to your intranet. These filters are automatically added when you run the Routing and Remote Access Server Setup Wizard using the options described in “Deploying VPN Servers” later in this chapter.

## Firewall Packet Filtering for VPN Traffic

It is a common practice to use a firewall to provide protection for intranet hosts, such as a VPN server, from Internet hosts. If you have a firewall, you must configure packet filters on the firewall to allow traffic to and from VPN clients on the Internet and the VPN server.

The following are common configurations of firewalls with a VPN server:


- The VPN server is directly attached to the Internet, and the firewall is between the VPN server and the intranet.
- The firewall is directly attached to the Internet, and the VPN server is between the firewall and the intranet.
- Two firewalls are used—one between the VPN server and the intranet and one between the VPN server and the Internet.

## VPN Server in Front of the Firewall

To prevent the VPN server from sending or receiving any traffic on its Internet interface except VPN traffic, you must configure PPTP, L2TP/IPsec, and SSTP input and output filters on the interface that corresponds to the connection to the Internet. Because IPv4 and IPv6 routing is enabled by default on the Internet interface by the Routing and Remote Access Server Setup Wizard, if VPN packet filters are not configured on the Internet interface, traffic received on the Internet interface is forwarded.

When the VPN server is in front of the firewall attached to the Internet, you must add to the Internet interface packet filters that allow only VPN traffic to and from the IPv4 or IPv6 address of the VPN server's Internet interface.

For inbound traffic, the VPN server decrypts the tunneled data and forwards it to the firewall. The firewall in this configuration is acting as a filter for intranet traffic and can prevent specific resources from being accessed, scan data for viruses, and perform intrusion detection, among other functions. Figure 12-2 shows the VPN server in front of the firewall.



**Figure 12-2** The VPN server in front of the firewall

The firewall is configured for the appropriate rules for intranet traffic to and from VPN clients according to your network security policies.

For the Internet interface on the VPN server, you can configure the VPN traffic input and output filters for both IPv4 and IPv6 using the Routing and Remote Access snap-in. These filters are automatically configured when you run the Routing and Remote Access Server Setup Wizard and choose the Remote Access (Dial-up or VPN) configuration, the VPN remote access type, select the correct Internet interface, and leave the Enable Security On The Selected Interface By Setting Up Packet Filters check box on the VPN Connection page selected (enabled by default). Additionally, the Routing and Remote Access Server Setup Wizard will automatically add and enable the same ports in the Windows Firewall.

The following sections describe these filters in detail in case you must manually configure them.

**PPTP Traffic Filters** The following are IPv4 input filters (also known as *inbound filters*) for PPTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 1723

This filter allows PPTP tunnel management traffic to the VPN server.

- Destination IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47

This filter allows PPTP tunneled data to the VPN server.

- Destination IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [established] source port of 1723

This filter is required only when the VPN server is acting as a calling router in a site-to-site (also known as router-to-router) VPN connection. TCP [established] traffic is accepted only when the VPN server initiated the TCP connection.

The following are IPv4 output filters (also known as *outbound filters*) for PPTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP source port of 1723

This filter allows PPTP tunnel management traffic from the VPN server.

- Source IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and IP Protocol ID of 47

This filter allows PPTP tunneled data from the VPN server.

- Source IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP [established] destination port of 1723

This filter is required only when the VPN server is acting as a VPN client (a calling router) in a site-to-site VPN connection. TCP [established] traffic is sent only when the VPN server initiated the TCP connection.

The following are IPv6 input filters for PPTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IPv6 address of the VPN server's Internet interface, prefix length of 128, and TCP destination port of 1723
- Destination IPv6 address of the VPN server's Internet interface, prefix length of 128, and IP Protocol ID of 47
- Destination IPv6 address of the VPN server's Internet interface, prefix length of 128, and TCP [established] source port of 1723

The following are IPv6 output filters for PPTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IPv6 address of the VPN server's Internet interface, prefix length of 128, and TCP source port of 1723
- Source IPv6 address of the VPN server's Internet interface, prefix length of 128, and IP Protocol ID of 47
- Source IPv6 address of the VPN server's Internet interface, prefix length of 128, and TCP [established] destination port of 1723

**L2TP/IPsec Traffic Filters** The following are IPv4 input filters for L2TP/IPsec traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 500

This filter allows Internet Key Exchange (IKE) traffic to the VPN server.

- Destination IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 4500

This filter allows IPsec NAT-T traffic to the VPN server.

- Destination IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 1701

This filter allows L2TP traffic to the VPN server.

The following are IPv4 output filters for L2TP/IPsec traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 500

This filter allows IKE traffic from the VPN server.

- Source IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 4500.

This filter allows IPsec NAT-T traffic from the VPN server.

- Source IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 1701

This filter allows L2TP traffic from the VPN server.

There are no filters required for IPsec Encapsulating Security Protocol (ESP) traffic for the IP protocol of 50. The Routing and Remote Access service filters are applied after the IPsec components remove the ESP header.

The following are IPv6 input filters for L2TP/IPsec traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IPv6 address of the VPN server's Internet interface, prefix length of 128, and UDP destination port of 500
- Destination IPv6 address of the VPN server's Internet interface, prefix length of 128, and UDP destination port of 4500
- Destination IPv6 address of the VPN server's Internet interface, prefix length of 128, and UDP destination port of 1701

The following are IPv6 output filters for L2TP/IPsec traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IPv6 address of the VPN server's Internet interface, prefix length of 128, and UDP source port of 500
- Source IPv6 address of the VPN server's Internet interface, prefix length of 128, and UDP source port of 4500
- Source IPv6 address of the VPN server's Internet interface, prefix length of 128, and UDP source port of 1701

**SSTP Traffic Filters** The following is an IPv4 input filter for SSTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP destination port of 443

This filter allows SSTP traffic to the VPN server.

The following is an IPv4 output filter for SSTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IPv4 address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and TCP source port of 443

This filter allows SSTP traffic from the VPN server.

The following is an IPv6 input filter for SSTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Destination IPv6 address of the VPN server's Internet interface, prefix length of 128, and TCP destination port of 443

The following is an IPv6 output filter for SSTP traffic with the filter action set to Drop All Packets Except Those That Meet The Criteria Below:

- Source IPv6 address of the VPN server's Internet interface, prefix length of 128, and TCP source port of 443

## VPN Server Behind the Firewall


In the more common configuration, the firewall is connected to the Internet, and the VPN server is an intranet resource that is connected to the perimeter network, also known as a screened subnet. The perimeter network is a subnet that contains resources that are available to Internet users, such as Web and FTP servers. The VPN server has an interface on both the perimeter network and the intranet. In this approach, the firewall must be configured with



input and output filters on its Internet interface that allow the passing of tunnel maintenance traffic and tunneled data to the VPN server. Additional filters can allow the passing of traffic to Web, FTP, and other types of servers on the perimeter network. For an added layer of security, the VPN server should also be configured with VPN traffic packet filters on its perimeter network interface.

The firewall in this configuration is acting as a filter for Internet traffic and can confine the incoming and outgoing traffic to the specific resources on the perimeter network, perform intrusion attempt detection, prevent denial of service attacks, and perform other functions.

Because the firewall does not have the encryption keys for each VPN connection, it can filter only on the plaintext headers of the tunneled data. In other words, all tunneled data passes through the firewall. This is not a security concern, however, because the VPN connection requires an authentication process that prevents unauthorized access beyond the VPN server. Figure 12-3 shows the VPN server behind the firewall on the perimeter network.



**Figure 12-3** The VPN server behind the firewall on the perimeter network

For both the Internet and network perimeter interfaces on the firewall, configure VPN traffic input and output filters by using the firewall's configuration software. Separate input and output packet filters can be configured on the Internet interface and the perimeter network interface.

Tables 12-1 and 12-2 summarize the packet filters that should be configured on the Internet and perimeter network interfaces of the firewall.

**Table 12-1 Packet Filters on the Internet Interface**

| Filter Type | IP Version | VPN Protocol | Traffic  |
|-------------|------------|--------------|--|
| Input       | IPv4       | PPTP         | Destination IPv4 address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB)  |
| Input       | IPv4       | PPTP         | Destination IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F)           |
| Input       | IPv4       | PPTP         | Destination IPv4 address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB)*      |
| Input       | IPv4       | L2TP/IPsec   | Destination IPv4 address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4)   |
| Input       | IPv4       | L2TP/IPsec   | Destination IPv4 address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194) |
| Input       | IPv4       | L2TP/IPsec   | Destination IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)           |
| Input       | IPv6       | L2TP/IPsec   | Destination IPv6 address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4)   |
| Input       | IPv6       | L2TP/IPsec   | Destination IPv6 address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194) |
| Input       | IPv6       | L2TP/IPsec   | Destination IPv6 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)           |
| Input       | IPv4       | SSTP         | Destination IPv4 address of the VPN server's perimeter network interface and TCP destination port of 443 (0x1BB)   |
| Input       | IPv6       | SSTP         | Destination IPv6 address of the VPN server's perimeter network interface and TCP destination port of 443 (0x1BB)   |
| Output      | IPv4       | PPTP         | Source IPv4 address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB)            |
| Output      | IPv4       | PPTP         | Source IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F)                |
| Output      | IPv4       | PPTP         | Source IPv4 address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB)*           |

**Table 12-1 Packet Filters on the Internet Interface**

| Filter Type | IP Version | VPN Protocol | Traffic  |
|-------------|------------|--------------|--|
| Output      | IPv4       | L2TP/IPsec   | Source IPv4 address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4)   |
| Output      | IPv4       | L2TP/IPsec   | Source IPv4 address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194) |
| Output      | IPv4       | L2TP/IPsec   | Source IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)      |
| Output      | IPv6       | L2TP/IPsec   | Source IPv6 address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4)   |
| Output      | IPv6       | L2TP/IPsec   | Source IPv6 address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194) |
| Output      | IPv6       | L2TP/IPsec   | Source IPv6 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)      |
| Output      | IPv4       | SSTP         | Source IPv4 address of the VPN server's perimeter network interface and TCP source port of 443 (0x1BB)   |
| Output      | IPv6       | SSTP         | Source IPv6 address of the VPN server's perimeter network interface and TCP source port of 443 (0x1BB)   |

\*These filters are required only when the VPN server is acting as a calling router in a site-to-site VPN connection. These filters should be used only in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" earlier in this chapter and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet that use this port.

**Table 12-2 Packet Filters on the Perimeter Network Interface**

| Filter Type | IP Version | VPN Protocol | Traffic  |
|-------------|------------|--------------|--|
| Input       | IPv4       | PPTP         | Source IPv4 address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB)  |
| Input       | IPv4       | PPTP         | Source IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F)      |
| Input       | IPv4       | PPTP         | Source IPv4 address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB)* |
| Input       | IPv6       | PPTP         | Source IPv6 address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB)  |
| Input       | IPv6       | PPTP         | Source IPv6 address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F)      |

**Table 12-2 Packet Filters on the Perimeter Network Interface**

| Filter Type | IP Version | VPN Protocol | Traffic  |
|-------------|------------|--------------|--|
| Input       | IPv6       | PPTP         | Source IPv6 address of the VPN server's perimeter network interface and TCP source port of 1723 (0x6BB)*           |
| Input       | IPv4       | L2TP/IPsec   | Source IPv4 address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4)             |
| Input       | IPv4       | L2TP/IPsec   | Source IPv4 address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194)           |
| Input       | IPv4       | L2TP/IPsec   | Source IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)                |
| Input       | IPv6       | L2TP/IPsec   | Source IPv6 address of the VPN server's perimeter network interface and UDP source port of 500 (0x1F4)             |
| Input       | IPv6       | L2TP/IPsec   | Source IPv6 address of the VPN server's perimeter network interface and UDP source port of 4500 (0x1194)           |
| Input       | IPv6       | L2TP/IPsec   | Source IPv6 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)                |
| Input       | IPv4       | SSTP         | Source IPv4 address of the VPN server's perimeter network interface and TCP source port of 443 (0x1BB)             |
| Input       | IPv6       | SSTP         | Source IPv6 address of the VPN server's perimeter network interface and TCP source port of 443 (0x1BB)             |
| Output      | IPv4       | PPTP         | Destination IPv4 address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB)  |
| Output      | IPv4       | PPTP         | Destination IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F)           |
| Output      | IPv4       | PPTP         | Destination IPv4 address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB)* |
| Output      | IPv6       | PPTP         | Destination IPv6 address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB)  |
| Output      | IPv6       | PPTP         | Destination IPv6 address of the VPN server's perimeter network interface and IP Protocol ID of 47 (0x2F)           |
| Output      | IPv6       | PPTP         | Destination IPv6 address of the VPN server's perimeter network interface and TCP destination port of 1723 (0x6BB)* |

**Table 12-2 Packet Filters on the Perimeter Network Interface**

| Filter Type | IP Version | VPN Protocol | Traffic  |
|-------------|------------|--------------|--|
| Output      | IPv4       | L2TP/IPsec   | Destination IPv4 address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4)   |
| Output      | IPv4       | L2TP/IPsec   | Destination IPv4 address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194) |
| Output      | IPv4       | L2TP/IPsec   | Destination IPv4 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)           |
| Output      | IPv6       | L2TP/IPsec   | Destination IPv6 address of the VPN server's perimeter network interface and UDP destination port of 500 (0x1F4)   |
| Output      | IPv6       | L2TP/IPsec   | Destination IPv6 address of the VPN server's perimeter network interface and UDP destination port of 4500 (0x1194) |
| Output      | IPv6       | L2TP/IPsec   | Destination IPv6 address of the VPN server's perimeter network interface and IP Protocol ID of 50 (0x32)           |
| Output      | IPv4       | SSTP         | Destination IPv4 address of the VPN server's perimeter network interface and TCP destination port of 443 (0x1BB)   |
| Output      | IPv6       | SSTP         | Destination IPv6 address of the VPN server's perimeter network interface and TCP destination port of 443 (0x1BB)   |


\*These filters are required only when the VPN server is acting as a calling router in a site-to-site VPN connection. These filters should be used only in conjunction with PPTP packet filters described in "VPN Server in Front of the Firewall" earlier in this chapter and configured on the VPN server's network perimeter interface. By allowing all traffic to the VPN server from TCP port 1723, there exists the possibility of network attacks from sources on the Internet that use this port.

There are no filters required for L2TP traffic at the UDP port of 1701. All L2TP traffic at the firewall, including tunnel maintenance and tunneled data, is encrypted as an IPsec ESP payload.

There are no IPv6 filters for PPTP traffic because Routing and Remote Access does not support IPv6 over PPTP connections.

## VPN Server Between Two Firewalls

Another configuration is when the VPN server computer is placed on the perimeter network between two firewalls. The Internet firewall, which is the firewall between the Internet and the VPN server, filters all Internet traffic from all Internet clients. The intranet firewall, which is the firewall between the VPN server and the intranet, filters intranet traffic from VPN clients. Figure 12-4 shows the VPN server between two firewalls on the perimeter network.



**Figure 12-4** The VPN server between two firewalls on the perimeter network

In this configuration, you should do the following:

- Configure your Internet firewall and VPN server with the packet filters as described in “VPN Server Behind the Firewall” earlier in this chapter.
- Configure your intranet firewall for the appropriate rules for intranet traffic to and from VPN clients according to your network security policies.

## Multi-Use VPN Servers

Because of the routes that are automatically created on VPN remote access clients, it is possible for the VPN client to send some unencrypted traffic to the VPN server rather than through the encrypted tunnel of the VPN connection. For example, the VPN client might connect to other services running on the VPN server without sending the traffic across the VPN connection. Only traffic that is destined to the public IP address of the VPN server is sent in the clear. If traffic from the client uses the IPv4 address of the Internal interface of the VPN server, however, it will be encrypted.

When a remote access VPN client creates a VPN connection with a VPN server, it creates a series of routes in the IPv4 routing table on the VPN client, including the following:

- **A default route that uses the VPN connection** The new default route for the VPN connection effectively replaces the existing default route for the duration of the connection. After the connection is made, all traffic that does not match an address on the directly connected network or the address of the VPN server is sent encrypted over the VPN connection.
- **A host route to the VPN server’s Internet IPv4 address** The host route for the address of the VPN server is created so that the VPN server is reachable. If the host route is not present, VPN traffic to the VPN server cannot be sent.

The result of having the host route for the VPN server is that all traffic that is sent to applications or services running on the VPN server to the VPN server's Internet IPv4 address is not sent across the VPN connection but is instead sent unencrypted across the Internet.

For example, when a remote access VPN client creates a VPN connection with a VPN server and then accesses a shared file on the VPN server computer using the VPN server's Internet address, that traffic is not sent using the VPN connection. The file sharing traffic is sent in plaintext over the Internet.

Additionally, if packet filters that allow VPN connection traffic only on the Internet interface are configured on the VPN server, all other traffic sent to the VPN server is discarded. All attempts to connect to applications or services running on the VPN server will fail because traffic attempting to connect to those services is not sent over the VPN connection.

The IPv4 address that is used by the VPN client to access services running on the VPN server depends on the way that the name of the VPN server is resolved. Typical users and applications refer to network resources using names rather than IPv4 addresses. The name must be resolved to an IPv4 address using either DNS or WINS. If the intranet DNS and WINS infrastructures do not contain a record mapping the name of the VPN server to the public IPv4 address of the VPN server's interface on the Internet, traffic to services running on the VPN server will always be sent across the VPN connection.

To prevent the VPN server from registering the public IPv4 address of its Internet interface in the intranet DNS, do the following:

1. In the Network Connections folder, obtain properties of the Internet Protocol Version 4 (TCP/IPv4) component of the Internet connection.
2. On the General tab, click Advanced.
3. In the Advanced TCP/IP Settings dialog box, on the DNS tab, clear the Register This Connection's Addresses In DNS check box, and then click OK three times.

If you are using NetBIOS over TCP/IP on your intranet, to prevent the VPN server from registering the public IPv4 address of its Internet interface with intranet WINS servers, do the following:

1. In the Network Connections folder, obtain properties of the Internet Protocol Version 4 (TCP/IPv4) component of the Internet connection.
2. On the General tab, click Advanced.
3. In the Advanced TCP/IP Settings dialog box, on the WINS tab, click Disable NetBIOS Over TCP/IP, and then click OK three times.

Before the VPN connection is made, the VPN client uses the Internet DNS infrastructure to resolve the name of the VPN server computer to its public IPv4 addresses. After the VPN connection is made, assuming that intranet DNS and WINS servers are configured either

during the PPP connection process or through the relaying of the DHCPInform message, the VPN client uses the intranet DNS and WINS infrastructures to resolve the name of the VPN server computer to its intranet IPv4 addresses.

## Blocking Traffic Routed from VPN Clients

After a VPN client successfully establishes a VPN connection, by default, any packet sent over the connection is received by the VPN server and forwarded. Packets sent over the connection can include:

- Packets originated by the VPN client computer
- Packets forwarded by the VPN client computer that are received from other computers

When the client computer makes the VPN connection, it creates by default a default route so that all traffic that matches the default route is sent over the VPN connection. If other computers are forwarding traffic to the VPN client, treating the VPN client computer as a router, that traffic is also forwarded across the VPN connection. This is a security problem because the VPN server has not authenticated the computer that is forwarding its traffic to the VPN client computer. The computer forwarding traffic to the VPN client computer has the same ability to send packets to the intranet as the authenticated VPN client computer.

To prevent the VPN server from forwarding traffic across the VPN connection for computers other than authenticated VPN client computers, configure IPv4 input packet filters on the network policy that is used for your VPN connections to discard all traffic except that originating from VPN clients. The default network policy named Connections to Microsoft Routing and Remote Access Server has a single IPv4 input filter with the filter action to Permit Only The Packets Listed Below and with the settings listed in Table 12-3.

**Table 12-3** Input Filter Settings

| IP Packet Filter Field | Setting        |
|------------------------|----------------|
| Source Address         | User's Address |
| Source Network Mask    | User's Mask    |
| Destination Address    | Any            |
| Destination Mask       | Any            |
| Protocol               | Any            |



**Note** Although the Routing and Remote Access snap-in displays User's Address and User's Mask, the actual filter that is created for each remote access client is for the client's assigned IPv4 address and a subnet mask of 255.255.255.255.

With this IPv4 input packet filter, the VPN server discards all traffic sent across the VPN connection except traffic that originates from VPN clients.



## Concurrent Access

When a VPN client computer has concurrent access to both the Internet and your intranet and has routes that allow reachability to both networks, the possibility exists that a malicious Internet user might use the connected VPN client computer to reach the private intranet through the authenticated VPN connection. This is possible if the VPN client computer has IPv4 routing enabled. IPv4 routing can be manually enabled on Windows-based computers by setting the HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\IPEnableRouter registry entry to **1** (data type is REG\_DWORD).

If your VPN clients must use concurrent access, you can help block unwanted traffic from the Internet by doing the following:

- Use an IPv4 packet filter on your network policies for VPN connections to discard inbound traffic on the VPN connection that has not been sent from the VPN client. The default network policy named Connections To Microsoft Routing And Remote Access Server has this IPv4 packet filter configured by default.
- Use the Network Access Protection feature in Windows Server 2008, Windows Vista, and Windows XP with Service Pack 3 to check whether connecting VPN clients have IPv4 routing enabled. If they do, do not allow unlimited remote access until it has been disabled.

## Unused VPN Protocols

If you not using all the VPN protocols, configure the Ports node in the Routing and Remote Access snap-in to set the number of ports for unused VPN protocols to 0. This prevents connections to the VPN server through protocols other than those being used for remote access VPN connections.

## Deploying VPN-Based Remote Access

To deploy VPN-based remote access by using Windows Server 2008, take the following steps:

- Deploy certificates.
- Configure Internet infrastructure.
- Configure RADIUS servers.
- Deploy VPN servers.
- Configure intranet infrastructure.
- Deploy VPN clients.

## Deploying Certificates

You must deploy certificates if you are using the following:

- **L2TP/IPsec connections with certificate authentication** Each VPN client computer and VPN server requires a computer certificate.

The Routing and Remote Access service supports the configuration of a preshared key for IPsec authentication of L2TP/IPsec connections. In the Routing and Remote Access snap-in, in the properties dialog box of a VPN server, on the Security tab, you can enable a custom IPsec policy and type the preshared key. VPN clients running Windows Server 2008, Windows Vista, Windows XP, or Windows Server 2003 also support the configuration of an IPsec preshared key. (In the properties dialog box of a VPN connection, on the Security tab, click IPsec Settings.) However, preshared key authentication for L2TP/IPsec connections is a weak form of authentication and is not recommended.

- **EAP-TLS or PEAP-TLS authentication with either smart cards or registry-based user certificates** Each VPN client computer needs either a smart card or a user certificate, and each authentication server needs a computer certificate.

It is possible to configure the VPN clients so that they do not validate the certificate of the authentication server, in which case computer certificates would not be required on the authentication servers. However, having the VPN clients validate the certificate of the authentication server is recommended for mutual authentication of the VPN client and authentication server, which helps prevent the VPN client from authenticating with an impersonating authentication server.

- **PEAP-MS-CHAP v2 authentication** Each authentication server needs a computer certificate, and each VPN client needs the certificate chain of the authentication server's computer certificate installed.

It is possible to configure the VPN clients so that they do not validate the certificate of the authentication server, in which case computer certificates on the authentication servers and the root CA certificate of the issuing CA on the VPN client is not required. However, having the VPN clients validate the certificate of the authentication server is recommended for mutual authentication of the VPN client and authentication server.

- **SSTP connections** Each VPN server needs a computer certificate, and each VPN client needs the root CA certificate of the issuing CA of the VPN server's computer certificate.

The VPN server computer certificate can have the Server Authentication or All Purpose usage in the Enhanced Key Usage (EKU) property of the certificate. The computer certificate should be valid, not expired, and have a certification revocation list (CRL) distribution point that is accessible from the Internet. The VPN client verifies that the computer certificate has not been revoked during the SSL authentication by checking the CRL at the distribution point stored in the computer certificate. Certificate revocation can also be checked by using the Online Certificate Status Protocol (OCSP), which uses HTTP to return a definitive digitally signed response of a certificate's status.

Additionally, the name of the Subject property of the VPN server's computer certificate must match the name of the VPN server in the properties dialog box of the VPN connection in the Network Connections folder on the VPN client. This name must match whether you are using DNS host names, IPv4 addresses, or IPv6 addresses for the VPN server.

## Deploying Computer Certificates

To install a computer certificate, a PKI must be present to issue certificates. Once the PKI is in place, you can install a computer certificate on VPN clients, VPN servers, or authentication servers in the following ways:

- By configuring autoenrollment of computer certificates to computers in an Active Directory domain
- By using the Certificates snap-in to request a computer certificate
- By using the Certificates snap-in to import a computer certificate
- By requesting a certificate over the Web
- By executing a CAPICOM script that requests a computer certificate

For more information, see “Deploying PKI” in Chapter 9.

## Deploying Root CA Certificates

You might need to deploy root CA certificates under the following circumstances:

- You are using PEAP-MS-CHAP v2 authentication.
- You are using SSTP connections.

**Root CA Certificates for PEAP-MS-CHAP v2** If you use PEAP-MS-CHAP v2 authentication, you might need to install the root CA certificates on your VPN clients for the computer certificate that your authentication servers (the VPN servers or the RADIUS servers) have been configured to use. If the root CA certificate of the issuer of the computer certificates that are installed on the authentication servers is already installed as a root CA certificate on your VPN clients, no other configuration is necessary. For example, if your root CA is a Windows Server 2008–based or Windows Server 2003–based online root enterprise CA, the root CA certificate is automatically installed on each domain member computer through Group Policy.

To verify whether the correct root CA certificate is installed on your VPN clients, you must:

1. Determine the root CA from the computer certificates installed on the authentication servers
2. Determine whether a certificate for the root CA is installed on your VPN clients

### To Determine the Root CA from the Computer Certificates Installed on the Authentication Servers

1. In the console tree of the Certificates snap-in for the authentication server computer account, expand Certificates (Local Computer or *Computername*), expand Personal, and then click Certificates.
2. In the details pane, double-click the computer certificate used for PEAP-MS-CHAP v2 authentication.
3. On the Certification Path tab for the properties of the certificate, note the name at the top of the certification path. This is the name of the root CA.

### To Determine Whether a Certificate for the Root CA Is Installed on Your VPN Client

1. In the console tree of the Certificates snap-in for the VPN client computer account, expand Certificates (Local Computer or *Computername*), expand Trusted Root Certification Authorities, and then click Certificates.
2. Examine the list of certificates in the details pane for a name matching the root CA for the computer certificates issued to the authentication servers.

You must install the root CA certificates of the issuers of the computer certificates of the authentication servers on each VPN client that does not contain them. The easiest way to install a root CA certificate on all your VPN clients is through Group Policy. For more information, see “Deploying PKI” in Chapter 9.

**Root CA Certificates for SSTP Connections** If you use SSTP connections, you might need to install the root CA certificate of the issuing CA of the computer certificates that are installed on your VPN servers. If the root CA certificate of the issuer of the computer certificates installed on the VPN servers is already installed as a root CA certificate on your VPN clients, no other configuration is necessary. If your root CA is a Windows Server 2008–based or Windows Server 2003–based online root enterprise CA, the root CA certificate is automatically installed on each domain member computer through Group Policy.

To verify whether the correct root CA certificate is installed on your VPN clients, you must:

1. Determine the root CA from the computer certificates installed on the VPN servers
2. Determine whether a certificate for the root CA is installed on your VPN clients

### To Determine the Root CA from the Computer Certificates Installed on the VPN Servers

1. In the console tree of the Certificates snap-in for the VPN server computer account, expand Certificates (Local Computer or *Computername*), expand Personal, and then click Certificates.
2. In the details pane, double-click the computer certificate used for SSL authentication.
3. On the Certification Path tab, note the name at the top of the certification path. This is the name of the root CA.

### To Determine Whether a Certificate for the Root CA Is Installed on Your VPN Client

1. In the console tree of the Certificates snap-in for the VPN client computer account, expand Certificates (Local Computer or *Computername*), expand Trusted Root Certification Authorities, and then click Certificates.
2. Examine the list of certificates in the details pane for a name or names matching the root CA for the computer certificate issued to the VPN servers.

You must install the root CA certificates of the issuers of the computer certificates of the VPN servers on each Windows Server 2008 or Windows Vista SP1-based VPN client that does not contain them. The easiest way to install a root CA certificate on all your VPN clients is through Group Policy. For more information, see “Deploying PKI” in Chapter 9.

### Deploying User Certificates

You can deploy user certificates to VPN client computers in the following ways:

- By configuring autoenrollment of user certificates to users in an Active Directory domain
- By using the Certificates snap-in to request a user certificate
- By using the Certificates snap-in to import a user certificate
- By requesting a certificate over the Web
- By executing a CAPICOM script that requests a user certificate

For more information, see “Deploying PKI” in Chapter 9.

## Configuring Internet Infrastructure

To configuring the Internet infrastructure for remote access VPN connections, perform the following:

- Place VPN servers in the perimeter network or on the Internet.
- Install Windows Server 2008 on VPN servers and configure Internet interfaces.
- Add address records to Internet DNS servers.

### Placing VPN Servers in the Perimeter Network or on the Internet

Decide where to place the VPN servers in relation to your Internet firewall. In the most common configuration, the VPN servers are placed behind the firewall on the perimeter network between the Internet and your intranet. If so, configure packet filters on the firewall to allow VPN traffic to and from the IPv4 or IPv6 address of the VPN servers’ perimeter network interfaces. For more information, see “Firewall Packet Filtering for VPN Traffic” earlier in this chapter.

## Installing Windows Server 2008 on VPN Servers and Configuring Internet Interfaces

Install Windows Server 2008 on the VPN server computer. Name the interfaces in the Network Connections folder with names that identify the network to which they are connecting. Connect the VPN server to either the Internet or to the perimeter network with one network adapter, and connect it to the intranet with another network adapter. Prior to running the Routing and Remote Access Server Setup Wizard, the VPN server computer will not forward IPv4 or IPv6 packets between the Internet and the intranet.

For the connection attached to the IPv4 Internet or the perimeter network, configure the TCP/IP (IPv4) protocol with a public IPv4 address, a subnet mask, and the default gateway of either the firewall (if the VPN server is connected to a perimeter network) or an ISP router (if the VPN server is directly connected to the Internet). Do not configure the connection with DNS server or WINS server IPv4 addresses.

For the connection attached to the IPv6 Internet or the perimeter network, configure the TCP/IP (IPv6) protocol with a global IPv6 address, a 64-bit prefix length, and the default gateway of either the firewall (if the VPN server is connected to a perimeter network) or an ISP router (if the VPN server is directly connected to the IPv6 Internet). Do not configure the connection with DNS server IPv6 addresses.

### Adding Address Records to Internet DNS Servers

To ensure that the name of the VPN server (for example, vpn.example.microsoft.com) can be resolved to its public IPv4 address or global IPv6 address, either add DNS address (A) or IPv6 address (AAAA) records to your Internet DNS server (if you are providing DNS name resolution for Internet users) or have your ISP add A or AAAA records to their DNS server(s) (if your ISP is providing DNS name resolution for Internet users). Verify that the name of the VPN server can be resolved to its public IPv4 address or global IPv6 address when connected to the Internet.

## Configuring Active Directory for User Accounts and Groups

To configure Active Directory for user accounts and groups, do the following:

1. Ensure that all users of VPN client computers have a corresponding user account.
2. Set the remote access permission on VPN client user accounts to Allow Access or Deny Access to manage remote access by user. Or, to manage access by group, set the remote access permission on user accounts to Control Access Through NPS Network Policy.
3. Organize VPN client user accounts into the appropriate universal and nested groups to take advantage of group-based network policies.

## Configuring RADIUS Servers

If you are using RADIUS for authentication, authorization, and accounting of VPN connections, configure and deploy your NPS-based RADIUS servers as described in Chapter 9, including the following steps:

1. Install a computer certificate on the NPS servers (for EAP-TLS, PEAP-TLS, or PEAP-MS-CHAP v2 authentication).
2. Configure logging.
3. Add RADIUS clients to the NPS server corresponding to each VPN server.

The NPS server uses a network policy to authorize remote access VPN connections. The default network policy named Connections To Microsoft Routing And Remote Access Server can be used for remote access VPN connections. However, by default, this network policy has its policy type set to Deny Access.

To use this network policy to accept remote access VPN connections, do the following:

1. In the console tree of the Network Policy Server snap-in, under Policies, click Network Policies.
2. Double-click the network policy named Connections To Microsoft Routing And Remote Access Server.
3. On the Overview tab, under Access Permission, click Grant Access, and then click OK.

You can also use the Configure VPN Or Dial-Up Wizard to create a set of policies that are customized for remote access VPN connections.

### To Create a Set of Policies for Remote Access VPN Connections

1. In the console tree of the Network Policy Server snap-in, click NPS.
2. In the details pane, under Standard Configuration, select RADIUS Server For Dial-Up Or VPN Connections from the drop-down list, and then click Configure VPN Or Dial-Up.
3. In the Configure VPN Or Dial-Up Wizard, on the Select Dial-Up or Virtual Private Network Connections Type page, click Virtual Private Network (VPN) Connections, and then type the name of the new NPS network policy (or use the name supplied by the wizard). Click Next.
4. On the Specify Dial-Up Or VPN Server page, add RADIUS clients as needed that correspond to your VPN servers. Click Next.
5. On the Configure Authentication Methods page, MS-CHAP v2 is already enabled. To enable and configure an EAP authentication type, select the Extensible Authentication Protocol check box, click an EAP type in the drop-down list, and then click Configure as needed (for example, to select the specific computer certificate to use for EAP-TLS, PEAP-TLS, or PEAP-MS-CHAP v2 authentication). Click Next.

6. On the Specify User Groups page, add the groups containing the user accounts allowed to make VPN remote access connections (for example, VPNUsers), and then click Next.
7. On the Specify IP Filters page, add IPv4 and IPv6 input and output packet filters to apply to all remote access VPN connections as needed. Click Next.
8. On the Specify Encryption Settings page, enable the allowed encryption strengths, and then click Next.
9. On the Specify A Realm Name page, specify the realm name and select the Before Authentication check box as needed. For more information about realm names, see Chapter 9. Click Next.
10. On the Completing New Dial-Up Or Virtual Private Network Connections And RADIUS Clients page, click Finish.

The Configure VPN Or Dial-Up Wizard creates a connection request policy and a network policy for remote access VPN connections. The Configure VPN Or Dial-Up Wizard configures the network policy with a single EAP method. For additional EAP methods, you can configure additional methods from the Settings tab for the properties of the network policy.

After you have configured the primary NPS server with the appropriate logging, RADIUS client, and policy settings, copy the configuration to the secondary or other NPS servers. For more information, see Chapter 9.

## Deploying VPN Servers

To deploy the VPN servers for remote access VPN connections, perform the following steps:

1. Install computer certificates.
2. Configure the VPN server's connection to the intranet.
3. Install the Network Access Services role.
4. Run the Routing and Remote Access Server Setup Wizard.
5. Add native IPv6 capability (optional).

### Installing Computer Certificates

For L2TP/IPsec or SSTP-based connections, or if the VPN server is the authentication server and you are using PEAP-MS-CHAP v2, EAP-TLS, or PEAP-TLS authentication, you must install a computer certificate on the VPN server. You can install a computer certificate using the methods described in "Deploying Certificates" earlier in this chapter.

### Configuring the VPN Server's Connection to the Intranet

For IPv4, configure the VPN server's connection to the intranet with a manual TCP/IP (IPv4) configuration consisting of IPv4 address, subnet mask, intranet DNS servers, and intranet



WINS servers. For IPv6, configure the VPN server's connection to the intranet with a manual TCP/IP (IPv6) configuration consisting of IPv6 address, 64-bit prefix length, and intranet DNS servers. In both cases, to prevent default route conflicts with the default route pointing to the IPv4 or IPv6 Internet, you must not configure the default gateway on the intranet connection.

## Installing the Network Access Services Role

To install Routing and Remote Access and the Connection Manager Administration Kit, use the Server Manager tool to install the Network Access and Policy Services role and the Connection Manager Administration Kit feature.

## Running the Routing and Remote Access Server Setup Wizard

The Routing and Remote Access Server Setup Wizard automates the configuration of many elements of the VPN server. The resulting default configuration can then be customized to fit your specific deployment needs.

### To Run the Routing and Remote Access Server Setup Wizard

1. On the Start menu, point to Administrative Tools, and then click Routing and Remote Access.
2. Right-click your server name, and then click Configure And Enable Routing And Remote Access. In the Routing and Remote Access Server Setup Wizard, on the Welcome page, click Next.
3. On the Configuration page, select Remote Access (Dial-up Or VPN), and then click Next.
4. On the Remote Access page, select VPN. If you also want the VPN server to support dial-up remote access connections, click Dial-up. Click Next.
5. On the VPN Connection page, click the connection that is connected to the Internet or your perimeter network. Ensure that the Enable Security On The Selected Interface By Setting Up Static Packet Filters check box is selected, and then click Next. Figure 12-5 shows an example.
6. On the Network Selection page (displayed only if you have multiple network adapters attached to the site), select the connection from which you want Routing and Remote Access to obtain DHCP, DNS, and WINS configuration for remote access VPN clients. Click Next if this page appeared.
7. On the IP Address Assignment page, select Automatically if the VPN server should use DHCP to obtain IPv4 addresses for remote access VPN clients. Alternatively, select From A Specified Range Of Addresses to use one or more static ranges of addresses. If any of the static address ranges is an off-subnet address range, routes must be added to the routing infrastructure for the VPN clients to be reachable. When you have completed IPv4 address assignment, click Next.

**Routing and Remote Access Server Setup Wizard**

**VPN Connection**  
To enable VPN clients to connect to this server, at least one network interface must be connected to the Internet.

Select the network interface that connects this server to the Internet.

Network interfaces:

| Name     | Description             | IP Address  |
|----------|-------------------------|-------------|
| CorpNet  | Linksys LNE100TX(v5)... | 10.0.0.1    |
| Internet | Broadcom NetXtreme 5... | 131.107.0.1 |

☒ Enable security on the selected interface by setting up static packet filters.  
Static packet filters allow only VPN traffic to gain access to this server through the selected interface.

[For more information about network interfaces.](#)  
[For more information about packet filtering.](#)

< Back Next > Cancel

**Figure 12-5** VPN Connection page

8. On the Managing Multiple Remote Access Servers page, if you are using the VPN server for authentication and authorization, select No, Use Routing And Remote Access To Authenticate Connection Requests. If you are using RADIUS for authentication and authorization, select Yes, Set Up This Server To Work With A RADIUS Server. Click Next.
9. If you selected RADIUS in step 8, on the RADIUS Server Selection page, configure the primary (mandatory) and alternate (optional) RADIUS servers and the RADIUS shared secret, and then click Next. Figure 12-6 shows an example.

**Routing and Remote Access Server Setup Wizard**

**RADIUS Server Selection**  
You can specify the RADIUS servers that you want to use for authentication and accounting.

Enter the primary and alternate RADIUS servers that this server will use for remote authentication and accounting.

Primary RADIUS server: CORPRAD1

Alternate RADIUS server: CORPRAD2

Type the shared secret (password) that is used to contact these RADIUS servers.

Shared secret:

< Back Next > Cancel

**Figure 12-6** RADIUS Server Selection page

10. On the Completing The Routing and Remote Access Server Setup Wizard page, click Finish.
11. If the Routing and Remote Access Server Setup Wizard cannot automatically configure the DHCP Relay Agent component with the IPv4 addresses of DHCP servers on the intranet, you are prompted with a message. Click OK, or click Help for more information.

If your VPN server is not acting as a site-to-site VPN router, you can disable demand-dial routing to create a dedicated remote access VPN server.

### To Disable Demand-Dial Routing for Site-to-Site VPN Connections

1. From the console tree of the Routing and Remote Access snap-in, right-click the name of the server, and then click Properties.
2. On the General tab, under IPv4 Router, click Local Area Network (LAN) Routing Only, and then click OK.

## Enabling Native IPv6 Capability

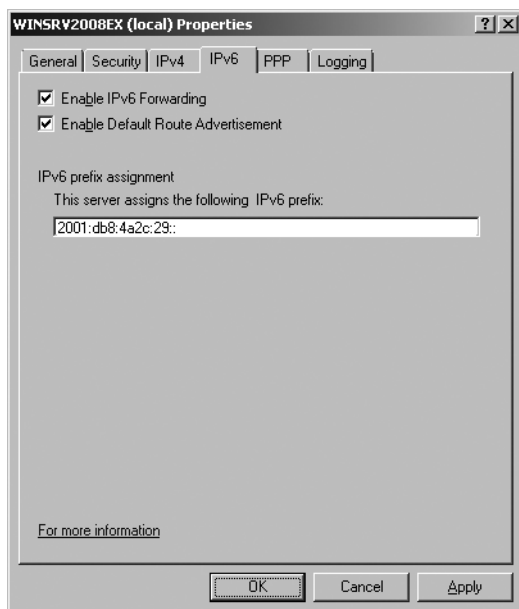
Native IPv6 capability for remote access VPN connections—IPv6 packets either inside the VPN tunnel or over a native IPv6 VPN connection—is not a current requirement for many intranets. For this reason, the Routing and Remote Access Server Setup Wizard does not automatically enable native IPv6 capability for remote access VPN connections over the IPv4 or IPv6 Internet.

To configure native IPv6 capability for remote access VPN connections in Routing and Remote Access, you need to do the following:

- Enable IPv6 routing for remote access connections.
- Configure router advertisement behavior.
- Configure the DHCPv6 Relay Agent to relay DHCPv6 messages between VPN clients and DHCPv6 servers on the intranet.

### To Configure the VPN Server to Support Native IPv6 Traffic Over VPN Connections

1. In the console tree of the Routing and Remote Access snap-in, right-click the name of the VPN server, and then click Properties.
2. On the General tab, select IPv6 Remote Access Server, and then click Apply.
3. On the IPv6 tab, ensure that the Enable IPv6 Forwarding and Enable Default Route Advertisement check boxes are selected. Type the subnet prefix that will be assigned to IPv6-based VPN clients as they connect. You do not need to specify the prefix length. For example, for the subnet prefix 2001:db8:4a2c:29::/64, type **2001:db8:4a2c:29::**. Figure 12-7 shows an example.
4. Click OK. You will be prompted to restart the Routing and Remote Access service.
5. In the console tree of the Routing and Remote Access snap-in, expand the IPv6 node.



**Figure 12-7** The IPv6 tab of the Routing and Remote Access Server Properties page

6. Right-click General, and then click New Routing Protocol.
7. In the New Routing Protocol dialog box, click OK to add the DHCPv6 Relay Agent component.
8. In the console tree, right-click DHCPv6 Relay Agent, click New Interface, select Internal, and then click OK twice.
9. Right-click DHCPv6 Relay Agent, and then click Properties.
10. On the Servers tab, type the global addresses of your DHCPv6 servers on the intranet, and then click OK.

## Configuring Intranet Network Infrastructure

To deploy the intranet network infrastructure for remote access VPN connections, perform the following steps:

1. Configure routing on the VPN server.
2. Verify name resolution and intranet reachability from the VPN server.
3. Configure routing for off-subnet address pools (if needed).
4. Configure routing for the IPv6 subnet prefix for remote access clients.

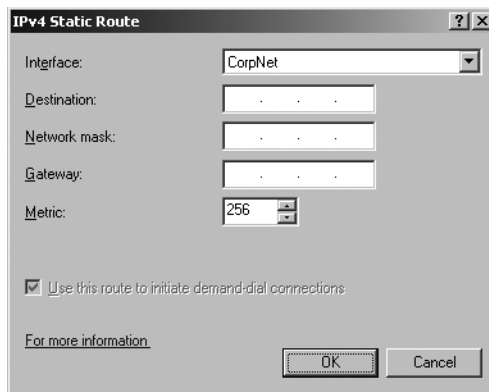
## Configuring Routing on the VPN Server

For your VPN servers to properly forward traffic to locations on the intranet, you must do one of the following:

- Add static routes that summarize the IPv4 and IPv6 address space used on the intranet.
- If you are using a RIP-capable IPv4 router on the intranet subnet connected to the VPN server, add the RIP routing protocol so that the VPN server can exchange routes with neighboring RIP routers and automatically add routes for intranet subnets to its routing table.

### To Add IPv4 Static Routes

1. In the console tree of the Routing and Remote Access snap-in, expand the IPv4 node.
2. Right-click Static Routes, and then click New Static Route.
3. In the IPv4 Static Route dialog box, shown in Figure 12-8, select the appropriate interface, and then type the destination, network mask, gateway, and metric for the static route. Click OK.



**Figure 12-8** IPv4 Static Route dialog box

4. Repeat steps 2 and 3 for additional IPv4 static routes.

### To Add IPv6 Static Routes

1. In the console tree of the Routing and Remote Access snap-in, expand the IPv6 node.
2. Right-click Static Routes, and then click New Static Route.
3. In the IPv6 Static Route dialog box, shown in Figure 12-9, select the appropriate interface and type the destination, prefix length, gateway, and metric for the static route. Click OK.
4. Repeat steps 2 and 3 for additional IPv6 static routes.

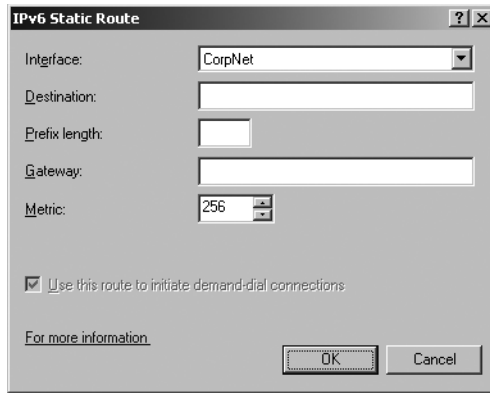


Figure 12-9 IPv6 Static Route dialog box



**Note** You must add IPv6 static routes only if you have configured your VPN server for native IPv6 capability.

### To Configure the VPN Server as a RIP Router

1. In the console tree of the Routing and Remote Access snap-in, expand the IPv4 node.
2. Right-click General, and then click New Routing Protocol.
3. In the New Routing Protocols dialog box, click RIP Version 2 For Internet Protocol, and then click OK.
4. Right-click RIP, and then click New Interface.
5. Select the intranet interface of the VPN server, and then click OK.
6. In the RIP Properties dialog box, configure the RIP routing protocol as used by the neighboring RIP router on the intranet subnet of the VPN server, and then click OK.

### Verifying Name Resolution and Reachability from the VPN Server

Verify that the VPN server can resolve names and successfully communicate with intranet resources by using the Ping command, Windows Internet Explorer, and making drive and printer connections to known intranet servers.

### Configuring Routing for Off-Subnet Address Pools

If you configured the VPN server with IPv4 address pools, and any of the pools are off-subnet, you must ensure that the routes representing the off-subnet address pools are present in your intranet IPv4 routing infrastructure. You can add static routes representing the off-subnet address pools to the neighboring routers of the VPN server and then propagate the routes to other routers by using the routing protocol of your intranet. When you add the static routes, you must specify that the gateway or next hop address is the intranet interface of the VPN server.

## Configuring Routing for the IPv6 Subnet Prefix for Remote Access Clients

To ensure that IPv6-capable remote access clients are reachable from the intranet, you must add a static route representing the subnet prefix for remote access clients to the neighboring IPv6 routers of the VPN server and then propagate the routes to other routers by using the routing protocol of your intranet. When you add the static routes, you must specify that the gateway or next hop address is the link-local address of the intranet interface of the VPN server.

## Deploying VPN Clients

To deploy VPN clients for remote access VPN connections, do the following as needed:

- Manually configure VPN clients.
- Configure and deploy CM profiles by using the CMAK.
- Configure concurrent access to the Internet and intranet.

### Manually Configuring VPN clients

If you have a small number of VPN clients, you can manually configure VPN connections for each VPN client. For Windows Server 2008 and Windows Vista VPN clients, use the Set Up A Connection Or Network Wizard. For Windows XP and Windows 2003 VPN clients, use the New Connection Wizard.

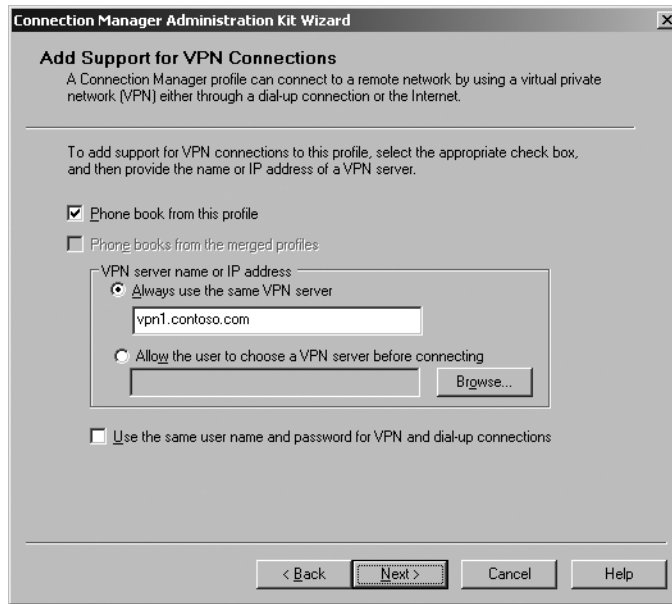
### Configuring and Deploying CM Profiles by Using the CMAK

For a large number of VPN clients running different versions of Windows, you should use the CMAK to create a CM profile for your users. After it is created, you must distribute the CM profile (a self-extracting executable file) to your users. Each user must execute the CM profile, which automatically creates a VPN connection in that user's Network Connections folder.

#### To Configure a CM Profile for a VPN Connection

1. On the Start menu, click Administrative Tools, and then click Connection Manager Administration Kit. If the CMAK is not already installed, you can install it by clicking Add Features in Server Manager and selecting it from the Features list.
2. On the Welcome page of the CMAK Wizard, click Next.
3. On the Select The Target Operating System page, select either Windows Vista or Windows Server 2003, Windows XP, Or Windows 2000, depending on which set of VPN client computers this CM profile will be distributed. Click Next.
4. On the Create Or Modify A Connection Manager Profile page, click Next to create a new profile.
5. On the Specify The Service Name And The File Name page, type the name of the profile as it will appear in the Network Connections folder and the name of the profile as it will be stored on the disk. Click Next.


6. On the Specify A Realm Name page, configure a realm name and where it will appear relative to the user name if needed. A realm name typically indicates where the user account is stored, as identified by a domain or an organization name. If you do not need to specify a realm name, click Next.
7. On the Merge Information From Other Profiles page, specify which existing profiles need to be merged into this new profile as needed, and then click Next.
8. On the Add Support For VPN Connections page, shown in Figure 12-10, select Phone Book From This Profile. In the VPN Server Name Or IP Address section, type the fully qualified domain name (FQDN), the public IPv4 address, or the global IPv6 address of the VPN server's Internet interface. Alternatively, select Allow The User To Choose A VPN Server Before Connecting, and then specify a text file containing a list of names or addresses of your VPN servers. Click Next.



**Figure 12-10** Add Support For VPN Connections page

9. On the Create Or Modify A VPN Entry page, click Edit to modify the settings of the default VPN entry. In the Edit VPN Entry dialog box, specify appropriate settings on the General, IPv4, IPv6, Security (authentication protocols and encryption requirements), and Advanced tabs. Figure 12-11 shows the default settings on the Security tab for a new entry. Click OK, and then click Next.
10. On the Add a Custom Phone Book page, clear the Automatically Download Phone Book Updates check box. (The VPN connection does not need to automatically create a dial-up connection.) Click Next.
11. On the Configure Dial-up Networking Entries page, click Next.





**Figure 12-11** New VPN Entry dialog box

12. On the Specify Routing Table Updates page, if you are using the CM profile to add routes to the VPN clients for concurrent Internet and intranet access, select Define A Routing Table Update, and then specify the file containing the routes or a URL that contains the routes. Click Next.
13. On the Configure Proxy Settings For Internet Explorer page, if you want to configure the VPN clients with a proxy server on the intranet, select either Automatically Copy The Internet Explorer Proxy Settings For The Current Use To The Tunnel Interface or Automatically Configure Proxy Settings, and then specify the file containing the proxy settings. Click Next.
14. On the Add Custom Actions page, configure custom actions as needed. Click Next.
15. On the Display A Custom Logon Bitmap page, if you want to use a custom bitmap for the user logon dialog box, click Custom Graphic, and then specify the location of the bitmap file that is 330 by 140 pixels. Click Next.
16. On the Display A Custom Phone Book Bitmap page, if you want to use a custom bitmap for the phone book dialog box, click Custom Graphic, and then specify the location of the bitmap file that is 114 by 309 pixels. Click Next.
17. On the Display Custom Icons page, if you want to use custom bitmaps in the Network and Sharing Center or Network Connections folder, click Custom Icons, and then specify the location of the bitmap files that are 32 by 32 and 16 by 16 pixels. Click Next.
18. On the Include A Custom Help File page, if you want to include a custom help file with the profile, click Custom Help File, and then specify the location of the CHM file. Click Next.

19. On the Display Custom Support Information page, if you want to include standard support text that appears in the logon dialog box, in the Support Information text box, type the desired text. Click Next.
20. On the Display A Custom License Agreement page, if you want to display a custom license agreement during the installation of the CM profile, specify the file containing the license agreement text. Click Next.
21. On the Install Additional Files With The Connection Manager Profile page, if you want to include with the CM profile additional files that are installed on the user's computer with the profile, specify their locations. Click Next.
22. On the Build The Connection Manager Profile And Its Installation Program page, click Next.
23. On the Your Connection Manager Profile Is Complete And Ready To Distribute page, click Finish.

### **Direct from the Source: Enhancements to CM Profiles**

Remote access connections using CM profiles made with Windows Server 2003 do not support DNS dynamic updates by the remote access clients. As a workaround, it is necessary to specify a post-connect action script within the CM profile to register with the intranet DNS server after the VPN connection has been established. The new version of Connection Manager included with Windows Server 2008 adds client DNS dynamic update registration functionality. You can configure DNS dynamic update from the Advanced tab for a dial-up or VPN entry on the Dial-Up Or VPN Entry Page within the CMAK Wizard. CM profiles for Windows Vista also includes support for profile authoring with IPv6 configuration options.

*Tim Quinn, Support Escalation Engineer*

*Enterprise Platform Support*

## **Distributing Your CM Profiles**

There are several ways to distribute your CM profile. Choose one of the following methods, or provide more than one method to give your users a choice.

**Distributing CM Profiles on CD or Disk** You can distribute CDs or disks containing your self-installing CM profile. A disk can include a floppy disk or, more commonly with new computers that do not include floppy disk drives, a Universal Serial Bus (USB) flash drive (UFD).

The benefit of distributing this way is that you can physically give a copy to all users or send them easily through the mail. However, this solution might be costly and has little inherent security.

**Distributing CM Profiles by E-Mail** You can send a CM profile through e-mail to your users. If you choose to send the CM profile through e-mail, ensure that users are able to receive .exe files, because not all e-mail systems allow executable files as attachments. A workaround is to compress the CM profile in a Zip format before sending.

**Distributing CM Profiles by Download** You can set up a Web site from which users can download the CM profile. Desktop users and portable-computer users can download directly to their computers from a Web site on your intranet.

It is also possible to make the CM profile available by download from a Web site over the Internet. However, identify any security risks to your organization before posting your CM profile on an Internet site.

**Pre-Installing CM Profiles** You can pre-install the CM profile on each client computer individually. The benefit of this method is that users are not required to install anything themselves, which can reduce user frustration and calls to your help desk. However, this method requires administrator or help desk resources during the initial installation, which might be a large resource hit during the rollout phase of your deployment. This method is useful when there are a small number of client computers or when all the client computers and devices are controlled by your organization.

**Combining Distribution Methods** You can also use a combination of distribution methods. For example, a company could distribute the CM profiles on CD to users who work from their own computers from remote locations, provide downloads for local employees who have portable computers, and pre-install the CM profile on any new portable computers before distribution.

## Configuring Concurrent Access to the Internet and Intranet

To configure concurrent access to the IPv4 Internet and your intranet, you can use the following:

- Classless Static Routes DHCP option
- Connection Manager Administration Kit

**Using the Classless Static Routes DHCP Option** VPN clients running Windows Server 2008, Windows Vista, Windows XP, or Windows Server 2003 send a DHCPInform message to the VPN server after the PPP negotiation is complete, requesting a set of DHCP options. This is done so that the VPN client can obtain an updated list of DNS and WINS servers and a DNS domain name that is assigned to the VPN connection. The DHCPInform message is forwarded to a DHCP server on the intranet by the VPN server, and the response is sent back to the VPN client.

The DHCPInform message includes a request for the Classless Static Routes DHCP option. For concurrent access, the Classless Static Routes DHCP option contains a set of routes that represent the address space of your intranet and that are automatically added to the routing

table of the requesting VPN client and automatically removed when the VPN connection is terminated. The Classless Static Routes DHCP option (option number 121) must be manually configured on a DHCP server running Windows Server 2008 or Windows Server 2003.

To use the Classless Static Routes option for concurrent access, configure this option for the scope that corresponds to the intranet subnet to which the VPN server is connected, and add the set of routes that correspond to the summarized IPv4 address space of your organization intranet. For example, if you use the private IPv4 address space for your organization intranet, the Classless Static Routes option would have the following three routes:

- 10.0.0.0 with the subnet mask of 255.0.0.0
- 172.16.0.0 with the subnet mask of 255.240.0.0
- 192.168.0.0 with the subnet mask of 255.255.0.0

The Router IP address for each route added to the Classless Static Routes option should be set to the IPv4 address of a router interface on the intranet subnet to which the VPN server is connected. For example, if the VPN server is connected to the intranet subnet 10.89.192.0/20, and the IPv4 address of the intranet router on this subnet is 10.89.192.1, set the Router IP address for each route to 10.89.192.1.

**Using the Connection Manager Administration Kit** You can use the Connection Manager Administration Kit (CMAK) for Windows Server 2008 to configure specific routes as part of the CM profile that is distributed to VPN clients. For more information about the CMAK and CM profiles, see “VPN Clients” earlier in this chapter.



**More Info** For more information about configuring concurrent access with the CMAK, see “Split Tunneling for Concurrent Access to the Internet and an Intranet” at <http://technet.microsoft.com/en-us/library/bb878117.aspx>.

## Ongoing Maintenance

The areas of maintenance for a remote access VPN solution are as follows:

- Management of user accounts
- Management of VPN servers
- Updating of CM profiles

## Managing User Accounts

When a new user account is created in Active Directory and that user is allowed to create remote access VPN connections, add the new user account to the appropriate group for VPN access. For example, add the account to the Wcoast\_VPNUsers security group, which is a

member of the VPNUsers universal group. The network policy for VPN connections is configured to use membership in the VPNUsers group as a condition for granting access.

When user accounts are deleted in Active Directory, no additional action is necessary to prevent remote access VPN connections.

As needed, you can create additional universal groups and network policies to define remote access for different sets of users. For example, you can create a global Contractors group and a network policy that allows remote access VPN connections to members of the Contractors group only during normal business hours or for access to specific intranet resources.

## Managing VPN Servers

You might need to manage VPN servers when adding or removing a VPN server from your remote access VPN solution. Once deployed, VPN servers do not need a lot of ongoing maintenance. Most of the ongoing changes to VPN server configuration are because of capacity and changes in network infrastructure.

### Adding a VPN Server

1. Follow the design points and deployment steps in this chapter to create a new VPN server on the Internet.
2. Update or add the FQDN in the Internet DNS for the IPv4 or IPv6 address of the new VPN server.
3. Update your RADIUS server configuration to add the VPN server as a RADIUS client.

### Removing a VPN Server

To remove a VPN server:

1. Update or remove the FQDN in the Internet DNS for the IPv4 or IPv6 address of the VPN server.
2. Update your RADIUS server configuration to remove the VPN server as a RADIUS client.
3. Shut down and remove the VPN server.

### Adding Possible Connections

By default, the Routing and Remote Access Server Setup Wizard configures Routing and Remote Access with up to the following ports (each port can support a single VPN connection):

- 128 PPTP ports
- 128 L2TP ports
- 128 SSTP ports

To increase the maximum number of connections for a VPN protocol, do the following:

1. In the console tree of the Routing and Remote Access snap-in, right-click Ports, and then click Properties.
2. In the Ports Properties dialog box, double-click the WAN Miniport device corresponding to the VPN protocol.
3. In the Configure Device dialog box, in the Maximum Ports spin-box, type the maximum number of ports, and then click OK twice.

## Configuration for Changes in Infrastructure Servers

Infrastructure servers include DHCP, DNS, WINS, and RADIUS (NPS) servers. If the changes to these types of infrastructure servers affect the configuration of the VPN server, you will need to change the configuration of the VPN server for the new infrastructure.

**DHCP** The Routing and Remote Access service on the VPN server uses the DHCP Relay Agent and DHCPv6 Relay Agent routing protocol components to forward DHCP and DHCPv6 messages between VPN clients and DHCP or DHCPv6 servers on the intranet. If the IPv4 or IPv6 addresses of the configured DHCP or DHCPv6 servers change (for example, because of additions or removals of DHCP or DHCPv6 servers on the intranet), you must change the list of DHCP and DHCPv6 addresses for the DHCP Relay Agent and DHCPv6 Relay Agent routing protocol components on the VPN server.

**DNS** The VPN server sends the IPv4 addresses of its configured DNS servers to VPN clients during the PPP negotiation. Additional IPv4 addresses of DNS servers might be configured on the VPN client from the response to the DHCPInform message. If the IPv4 addresses of the configured DNS servers change (for example, because of additions or removals of DNS servers on the intranet), you must change the DNS server configuration on the VPN server and the DNS server option on the DHCP server to prevent VPN clients from configuring incorrect DNS server IPv4 addresses.

For native IPv6-based VPN connections, VPN clients obtain the IPv6 addresses from the response to the DHCPv6 Information-Request message. If the IPv6 addresses of the configured DNS servers change (for example, because of additions or removals of DNS servers on the intranet), you must change the IPv6 DNS server option on the DHCPv6 server to prevent it from configuring VPN clients with incorrect DNS server IPv6 addresses.

**WINS** The VPN server sends the IPv4 addresses of its configured WINS servers to VPN clients during the PPP negotiation. Additional IPv4 addresses of WINS servers might be configured on the VPN client based on the response to the DHCPInform message. If the IPv4 addresses of the configured WINS servers change (for example, because of additions or removals of WINS servers on the intranet), you must change the WINS server configuration on the VPN server and the NetBIOS name server option on the DHCP server to prevent VPN clients from configuring an incorrect WINS server IPv4 address.

**RADIUS** If the VPN server is configured to use RADIUS authentication, and the IPv4 addresses of the RADIUS servers change (for example, because of additions or removals of RADIUS servers on the intranet), you must do the following:

1. Ensure that the new RADIUS servers are configured with a RADIUS client corresponding to the VPN servers.
2. Update the configuration of the VPN servers to include the IPv4 addresses of the new RADIUS servers.

## Updating CM Profiles

To update a CM profile, do the following:

1. Create an updated CM profile by using the CMAK.
2. Distribute the updated CM profile to your VPN client users through e-mail, a file share, or other means with the instructions or automated process to execute the profile and update their VPN connection settings.

## Troubleshooting

Because of the different components and processes involved, troubleshooting remote access VPN connections can be a difficult task. This section describes the many tools that are provided with Windows Server 2008 and Windows Vista to troubleshoot remote access VPN connections and the most common problems with remote access VPN connections.

## Troubleshooting Tools

Microsoft provides the following tools to troubleshoot VPN connections from the VPN server:

- TCP/IP troubleshooting tools
- Authentication and accounting logging
- Event logging
- NPS event logging
- PPP logging
- Tracing
- Network Monitor 3.1

Additionally, Windows Server 2008 and Windows Vista provide the following tools to troubleshoot VPN connections from the VPN client:

- TCP/IP troubleshooting tools
- Network Diagnostics Framework support for remote access connections

## TCP/IP Troubleshooting Tools

The Ping, Tracert, and Pathping tools use ICMP Echo and Echo Reply and ICMPv6 Echo Request and Echo Reply messages to verify connectivity, display the path to a destination, and test path integrity. The **route print** command can be used to display the IPv4 and IPv6 routing tables. Alternatively, on the VPN server, you can use the **netsh routing ip show rtmroutes** command or the Routing and Remote Access snap-in to display routes. The Nslookup tool can be used to troubleshoot DNS and name resolution issues.

## Authentication and Accounting Logging

A VPN server running Windows Server 2008 supports the logging of authentication and accounting information for remote access VPN connections in local logging files when Routing and Remote Access is configured to perform authentication and accounting locally. This logging is separate from the events recorded in the Windows Logs\Security event log. You can use the information that is logged to track remote access usage and authentication attempts. Authentication and accounting logging is especially useful for troubleshooting network policy issues. For each authentication attempt, the name of the network policy that either accepted or rejected the connection attempt is recorded.

To enable authentication and accounting logging, open the Network Policy Server snap-in, click Accounting, and then click Configure Local File Logging. On the Settings tab, configure the appropriate settings.

The authentication and accounting information is stored in a configurable log file or files stored in the %SystemRoot%\System32\LogFiles folder. The log files are saved in Internet Authentication Service (IAS) or database-compatible format, meaning that any database program can read the log file directly for analysis. Routing and Remote Access can also send authentication and accounting information to a Structured Query Language (SQL) database.

If the VPN server is configured for RADIUS authentication and accounting, and the RADIUS server is a computer running Windows Server 2008 and NPS, the authentication and accounting logs are stored in the %SystemRoot%\System32\LogFiles folder on the NPS server computer. NPS for Windows Server 2008 can also send authentication and accounting information to a Microsoft SQL Server database.

## Event Logging

In the Routing and Remote Access snap-in, in the properties dialog box of a VPN server, on the Logging tab, there are four levels of logging for creating entries in the Windows Logs\System event log. To obtain the maximum amount of information, select Log All Events, and then try to complete the connection again. When the connection fails, check the Windows Logs\System event log for events with the event sources of RasServer, RemoteAccess, or RasSSTP that were logged during the connection process. After you are finished viewing the events, on the Logging tab, select Log Errors And Warnings to conserve system resources.



## NPS Event Logging

If your VPN servers are configured for RADIUS authentication, and your RADIUS servers are computers running Windows Server 2008 and NPS, check the Windows Logs\Security event log for NPS events corresponding to rejected (event ID 6273) or accepted (event ID 6272) connection attempts. NPS event log entries contain a lot of information on the connection attempt, including the name of the connection request policy that matched the connection attempt (the Proxy Policy Name field in the description of the event) and the network policy that accepted or rejected the connection attempt (the Network Policy Name field in the description of the event). NPS event logging for rejected or accepted connection attempts is enabled by default and configured in the Network Policy Server snap-in, in the properties dialog box of an NPS server, on the Service tab.

## PPP Logging

PPP logging records the series of programming functions and PPP control messages during a PPP connection and is a valuable source of information when you are troubleshooting the failure of a PPP connection. To enable PPP logging, in the Routing and Remote Access snap-in, in the properties dialog box of a VPN server, on the Logging tab, select the Log Additional Routing And Remote Access Information check box.

By default, the PPP log is stored as the Ppp.log file in the %SystemRoot%\Tracing folder.

## Tracing

The Routing and Remote Access service has an extensive tracing capability that you can use to troubleshoot complex network problems. You can enable components of Windows Server 2008 to log tracing information to files by using the Netsh tool or by setting registry values.

**Enabling Tracing with Netsh** You can use the Netsh tool to enable and disable tracing for specific components or for all components. To enable and disable tracing for a specific component, use the following syntax:

```
netsh ras diagnostics set rastracing Component enabled|disabled
```

where ***Component*** is a component in the list of Routing and Remote Access service components found in the Windows Server 2008 registry under HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Tracing. For example, to enable tracing for the RASAUTH component, the command is:

```
netsh ras diagnostics set rastracing rasauth enabled
```

To enable tracing for all components, run the following command:

```
netsh ras diagnostics set rastracing * enabled
```

**Enabling Tracing Through the Registry** You can configure the tracing function by changing settings in the Windows registry under HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Tracing.

You can enable tracing for each Routing and Remote Access service component by setting the registry values described later. You can enable and disable tracing for components while the Routing and Remote Access service is running. Each component is capable of tracing and appears as a subkey under the Tracing registry key.

To enable tracing for each component, you can configure the following registry value entries for each protocol key:

- **EnableFileTracing (REG\_DWORD) Flag** You can enable logging tracing information to a file by setting EnableFileTracing to **1**. The default value is 0.
- **FileDirectory (REG\_EXPAND\_SZ) Path** You can change the default location of the tracing files by setting FileDirectory to the path you want. The file name for the log file is the name of the component for which tracing is enabled. By default, log files are placed in the %SystemRoot%\Tracing folder.
- **FileTracingMask (REG\_DWORD) LevelOfTracingInformationLogged** FileTracingMask determines how much tracing information is logged to the file. The default value is 0xFFFF0000.
- **MaxFileSize (REG\_DWORD) SizeOfLogFile** You can change the size of the log file by setting different values for MaxFileSize. The default value is 0x10000 (64K).



**Note** Tracing consumes system resources and should be used sparingly to help identify network problems. After the trace is captured or the problem is identified, you should immediately disable tracing. Do not leave tracing enabled on multiprocessor computers.

Tracing information can be complex and detailed. Most of the time, this information is useful only to Microsoft support professionals or to network administrators who are experienced with Routing and Remote Access. The tracing log files can be sent to Microsoft support for analysis if necessary.

## Network Monitor 3.1

You can use Microsoft Network Monitor 3.1 or a commercial packet analyzer (also known as a *network sniffer*) to capture and view the traffic sent between a VPN server and VPN client during the VPN connection process and during data transfer or the RADIUS traffic sent between a VPN server and a RADIUS server. Network Monitor 3.1 includes RADIUS, PPTP, PPP, L2TP, IPsec, HTTP, SSL, and EAP parsers. A *parser* is a component included with Network Monitor 3.1 that can separate the fields of a protocol header and display their structure and values. Without a parser, Network Monitor 3.1 displays the hexadecimal bytes of a header, which you must parse manually.



**On the Disc** You can link to the download site for Network Monitor from the companion CD-ROM.

The proper interpretation of the remote access and VPN traffic with Network Monitor 3.1 requires an in-depth understanding of PPP, PPTP, IPsec, SSL, RADIUS, and other protocols. You can save Network Monitor 3.1 captures as files and send them to Microsoft support for analysis.

## Network Diagnostics Framework Support for Remote Access Connections

To provide a better user experience when encountering network connectivity issues, Windows Vista includes the Network Diagnostics Framework (NDF), a set of technologies and guidelines that enable a set of troubleshooters (also known as a *helper classes*) to assist in the diagnosis and possible automatic correction of networking problems. When a user experiences a networking problem in Windows Vista, NDF will provide the user the ability to diagnose and repair the problem within the context of that problem. This means that the diagnostics assessment and resolution steps are presented to the users within the application or dialog box that they were using when the problem occurred or based on the failed network operation.

With NDF, when the user tries to complete a task that depends on network connectivity, such as browsing to a Web site or sending an e-mail message, an error message might appear indicating failure to complete the task (such as “Page cannot be displayed” or “Server is not available”). With NDF, the error message might include an option to diagnose the problem. During the diagnosis, NDF will analyze why the user’s task has failed and present a solution to the problem or possible list of causes and corrections in clear language, allowing the user to take action to fix the problem.

Windows Vista includes a troubleshooter to diagnose failed remote access connections. If a remote access connection fails, Windows displays a dialog box with information about the error. The dialog box includes a Diagnose button that launches the remote access NDF troubleshooter. From the diagnosis session, users can repair their remote access connection problem without needing to involve IT support staff.

## Troubleshooting Remote Access VPNs

Remote access VPN problems typically fall into the following categories:

- Connection attempt is rejected when it should be accepted.
- L2TP/IPsec authentication issues.
- SSTP authentication issues.
- Connection attempt is accepted when it should be rejected.

- Unable to reach locations beyond the VPN server.
- Unable to establish a tunnel.

Use the following troubleshooting tips to isolate the configuration or infrastructure issue that is causing the problem.

## Connection Attempt Is Rejected When It Should Be Accepted

If a connection attempt is being rejected when it should be accepted, check the following:

- Using the Ping command, verify that the FQDN of the VPN server is being resolved to its correct IPv4 address. The ping itself might not be successful because of packet filtering that is preventing the delivery of ICMP messages to and from the VPN server.
- For password-based authentication, verify that the VPN client's user credentials—consisting of user name, password, and domain name—are correct and can be validated by the authentication server (the VPN server or the RADIUS server).
- Verify that the user account of the VPN client is not locked out, expired, or disabled, and that the time the connection is being made corresponds to the configured logon hours. If the password on the account has expired, verify that the remote access VPN client is using PEAP-MS-CHAP v2 or MS-CHAP v2. PEAP-MS-CHAP v2 and MS-CHAP v2 are the only authentication protocols provided with Windows Server 2008 that allow you to change an expired password during the connection process.

For an administrator-level account whose password has expired, reset the password using another administrator-level account.

- Verify that the user account has not been locked out because of remote access account lockout.
- Verify that the Routing and Remote Access service is running on the VPN server.
- For SSTP based VPN connections, verify that the Secure Socket Tunneling Protocol Service is running on the VPN server.
- In the Routing and Remote Access snap-in, in the properties dialog box of a VPN server, on the General tab, verify that the VPN server is enabled as an IPv4 or IPv6 remote access server.
- In the Routing and Remote Access snap-in, in the properties dialog box of the Ports node, verify that the WAN Miniport (PPTP), WAN Miniport (L2TP), and WAN Miniport (SSTP) devices are enabled for inbound remote access.
- Verify that the VPN client, the VPN server, and the network policy for VPN connections are configured to use at least one common authentication method.
- Verify that the VPN client and the network policy for VPN connections are configured to use at least one common encryption strength.

- Verify that the parameters of the connection have permission through network policies.

For the connection to be accepted, the parameters of the connection attempt must:

- ❑ Match all the conditions of at least one network policy.
- ❑ Be granted remote access permission through the user account (set to Allow Access), or if the user account has the Control Access Through NPS Network Policy option selected, the matching network policy must have the Grant Access policy type selected.
- ❑ Match all the settings of the network policy.
- ❑ Match all the settings of the dial-in properties of the user account.

To obtain the name of the network policy that rejected the connection attempt, scan the Windows Logs\Security event log for events corresponding to rejected (event ID 6273) or accepted (event ID 6272) connection attempts. The network policy that accepted or rejected the connection attempt is the Network Policy Name field in the description of the event.

- If you are logged on using an account with domain administrator permissions when you run the Routing and Remote Access Server Setup Wizard and configure Routing and Remote Access to perform authentication locally, the wizard automatically adds the computer account of the VPN server to the RAS and IAS Servers domain-local security group. This group membership allows the VPN server computer to access user account information. If the VPN server is unable to access user account information, verify that:
- The computer account of the VPN server computer is a member of the RAS and IAS Servers security group for all the domains that contain user accounts for which the VPN server is authenticating remote access. You can run the **netsh nps show registered-server** command at a command prompt to view the current registration. You can run the **netsh nps add registeredserver** command to register the server in a domain in which the VPN server is a member or other domains. Alternatively, you or your domain administrator can add the computer account of the VPN server computer to the RAS and IAS Servers security group of all the domains that contain user accounts for which the VPN server is authenticating remote access.
- If you add or remove the VPN server computer to the RAS and IAS Servers security group, the change does not take effect immediately (because of the way that Windows Server 2008 caches Active Directory information). For the change to take effect immediately, you must restart the VPN server computer.
- Verify that all the PPTP, L2TP, or SSTP ports on the VPN server are not already being used. If necessary to allow more connections, in the Routing and Remote Access snap-in, in the properties dialog box of the Ports object, increase the number of PPTP, L2TP, or SSTP ports.
- Verify that the VPN server supports the VPN protocol of the VPN client.

By default, a VPN client running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP has the Automatic VPN Type option selected. If the PPTP, L2TP IPsec, or SSTP VPN type is selected, verify that the VPN server supports the selected tunneling protocol.

When you run the Routing and Remote Access Server Setup Wizard and configure a VPN server, a Windows Server 2008-based computer running the Routing and Remote Access service is a PPTP, L2TP, and SSTP server with 128 PPTP ports, 128 L2TP ports, and 128 SSTP ports. To create a PPTP-only server, set the number of L2TP and SSTP ports to zero. To create an L2TP-only server, set the number of SSTP ports to 0 and the PPTP ports to 1, and disable remote access inbound connections and demand-dial connections for the WAN Miniport (PPTP) device. Do this in the Routing and Remote Access snap-in, in the Ports object dialog box. To create an SSTP-only server, set the number of L2TP ports to 0 and the PPTP ports to 1 and disable remote access inbound connections and demand-dial connections for the WAN Miniport (PPTP) device.

- If the VPN server is configured with static IPv4 address pools, verify that there are enough addresses for all the possible connections. If all the addresses in the static pools have been allocated to connected VPN clients, the VPN server will be unable to assign an IPv4 address for TCP/IP-based connections, and the connection attempt will be rejected.
- Verify how the VPN server is performing authentication. The VPN server can be configured to authenticate the credentials of the VPN client either locally or use RADIUS.
  - ❑ For RADIUS-based authentication, verify that the VPN server computer can communicate with the RADIUS server.
  - ❑ For local authentication, verify that the VPN server has joined the Active Directory domain and that the computer account of the VPN server computer has been added to the RAS and IAS Servers security group.

## L2TP/IPsec Authentication Issues

The following are the most common problems that cause L2TP/IPsec connections to fail:

- **No certificate** By default, L2TP/IPsec connections require that the VPN server and VPN client exchange computer certificates for IPsec peer authentication. Use the Certificates snap-in to check the local computer certificate stores of both the VPN client and VPN server to ensure that a suitable certificate exists.
- **Incorrect certificate** If certificates exist, they must be verifiable. Unlike manually configuring IPsec rules, the list of certification authorities (CAs) for L2TP/IPsec connections is not configurable. Instead, each computer in the L2TP connection sends a list of root CAs to its IPsec peer from which it accepts a certificate for authentication. The root CAs in this list correspond to the root CAs that issued computer certificates to the computer. For example, if Computer A is issued computer certificates by root CAs

CertAuth1 and CertAuth2, it notifies its IPsec peer during main mode negotiation that it will accept certificates for authentication from only CertAuth1 and CertAuth2. If the IPsec peer, Computer B, does not have a valid computer certificate issued from either CertAuth1 or CertAuth2, IPsec security negotiation fails.

The VPN client must have a valid computer certificate for IPsec authenticate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN server trusts. Additionally, the VPN server must have a valid computer certificate installed that was issued by a CA that follows a valid certificate chain from the issuing CA up to a root CA that the VPN client trusts.

- **A NAT is between the remote access client and remote access server** If there is a NAT between the VPN client and the VPN server, both computers must support IPsec NAT-T. VPN clients running Windows Server 2008, Windows Vista, Windows Server 2003, or Windows XP SP2 support IPsec NAT-T. VPN servers running Windows Server 2008 or Windows Server 2003 support IPsec NAT-T.
- **A firewall is between the remote access client and remote access server** If there is a firewall between a Windows VPN client and a Windows Server 2008 VPN server and you cannot establish an L2TP/IPsec connection, verify that the firewall allows L2TP/IPsec traffic to be forwarded. For more information, see “Firewall Packet Filtering for VPN Traffic” earlier in this chapter.

## SSTP Authentication Issues

The following are the most common problems that cause SSTP connections to fail:

- **No certificate** SSTP connections require that the VPN server send a computer certificate to the VPN client during the SSL authentication. Using the Certificates snap-in, verify that the VPN server has a suitable computer certificate installed.
- **Certificate validation fails** The VPN client must have the root CA certificate for the issuing CA of the VPN server’s computer certificate installed. Obtain the name of the root CA certificate of the VPN server’s computer certificate, and then verify that it is installed on your VPN clients. Also, do the following:
  - Verify that the computer certificate of the VPN server has not expired or been revoked.
  - Verify that the CRL distribution points listed in the CRL Distribution Points property of the VPN server’s computer certificate are reachable on the Internet.
  - Verify that the name of the VPN server, on the General tab in the properties dialog box of the VPN connection in the Network Connections folder, on the VPN client matches the Subject property of the VPN server’s computer certificate. This name must match whether you are using DNS host names, IPv4 addresses, or IPv6 addresses for the VPN server.

## Connection Attempt Is Accepted When It Should Be Rejected

If a connection attempt is being accepted when it should be rejected, check the following:

- Verify that the remote access permission on the user account is set to either Deny access or Control Access Through NPS Network Policy. If set to the latter, verify that the first matching network policy's type is set to Deny Access. To obtain the name of the network policy that accepted the connection attempt, scan the Windows Logs\Security event log for an event that corresponds to the connection attempt. The text of the event contains the policy name. The network policy that accepted or rejected the connection attempt is the Network Policy Name field in the description of the event
- If you have created a network policy to explicitly reject all connections, verify the policy conditions, type, and settings, and its location in the list of network policies.

## Unable to Reach Locations Beyond the VPN Server

If a VPN client cannot reach locations on the intranet beyond the VPN server, check the following:

- In the Routing and Remote Access snap-in, in the properties dialog box of a VPN server, on the General tab, verify that the IPv4 Remote Access Server and IPv6 Remote Access Server check boxes are selected.
- Verify that the IPv4 or IPv6 protocol is enabled for forwarding on the IPv4 and IPv6 tabs for the properties of a VPN server in the Routing and Remote Access snap-in.
- Verify the IPv4 address pools of the VPN server.

If the VPN server is configured to use an off-subnet IPv4 address pool, verify that the range of addresses set by the IPv4 address pool are reachable by the hosts and routers of the intranet. If not, you must either add the IPv4 routes for the VPN server's IPv4 address pools to the routers of the intranet or, if you are using the RIP routing protocol, enable RIP on the VPN server. If the routes for the off-subnet address pools are not present, remote access VPN clients cannot receive traffic from locations on the intranet.

If the VPN server is configured to use DHCP to obtain IPv4 addresses for remote access clients, and no DHCP server is available, the VPN server assigns addresses from the Automatic Private IP Addressing (APIPA) address range from 169.254.0.1 through 169.254.255.254. Allocating APIPA addresses for remote access clients works only if the network to which the VPN server is attached is also using APIPA addresses.

If the VPN server is using APIPA addresses when a DHCP server is available, verify that the proper adapter is selected from which to obtain DHCP-allocated IPv4 addresses. This selection is done through the Routing and Remote Access Server Setup Wizard. In the Routing and Remote Access snap-in, in the properties dialog box of a VPN server, on the IPv4 tab, you can manually choose a LAN adapter from the Adapter list.



If the IPv4 address pools are on-subnet—a range of IPv4 addresses that are a subset of the range of IP addresses for the network to which the VPN server is attached—verify that the range of IPv4 addresses in the IPv4 address pools are not assigned to other TCP/IP nodes either through manual configuration or through DHCP.

- Verify that the IPv6 subnet prefix that is being assigned to IPv6-capable VPN clients is a route in your IPv6 routing infrastructure that points back to the intranet interface of the VPN server.
- Verify that there are no IPv4 or IPv6 input or output packet filters in the settings of the network policy for VPN connections that are preventing the sending or receiving of traffic.

## Unable to Establish Tunnel

If a VPN client cannot create a tunnel to the VPN server, check the following:

- Verify that packet filtering on a router interface between the VPN client and the VPN server is not preventing the forwarding of VPN traffic. See “Firewall Packet Filtering for VPN Traffic” earlier in this chapter for information about the types of traffic that must be allowed for VPN connections.

On a Windows Server 2008–based VPN server, IPv4 packet filtering can be separately configured on the Windows Firewall with Advanced Security and the Routing and Remote Access snap-in. Check both places for filters that might be excluding VPN connection traffic.

- Verify that the Winsock Proxy client is not currently running on the VPN client.

When the Winsock Proxy client is active, Windows Sockets (Winsock) API calls such as those used to create tunnels and send tunneled data are intercepted and forwarded to a configured proxy server.

A proxy server–based computer allows an organization to access specific types of Internet resources (typically Web and FTP) without directly connecting that organization to the Internet. The organization can instead use private IP address prefixes such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16.

Proxy servers are typically used so that private users in an organization can have access to public Internet resources as if they were directly attached to the Internet. VPN connections are typically used so that authorized public Internet users can gain access to private organization resources as if they were directly attached to the private network. A single computer can act as a proxy server (for private users) and a VPN server (for authorized Internet users) to facilitate both exchanges of information.

## Chapter Summary

Deploying a remote access VPN solution involves configuration of Active Directory, PKI, Group Policy, and RADIUS elements of a Windows-based authentication infrastructure and planning and deployment of VPN servers on the Internet. Once deployed, ongoing maintenance of a remote access VPN solution consists of managing VPN servers and their configuration for changes in infrastructure servers and updating and deploying CM profiles. Common problems with VPN connections include the inability to connect because of an authentication or authorization failure and the inability to reach intranet resources from the VPN client.

## Additional Information

For additional information about VPN support in Windows, see the following:

- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Virtual Private Networks” (<http://www.microsoft.com/vpn>)

For additional information about VPN Internet standards, see the following:

- RFC 2637, “Point-to-Point Tunneling Protocol (PPTP)”
- RFC 2661, “Layer Two Tunneling Protocol (L2TP)”
- RFC 3193, “Securing L2TP using IPsec”

For additional information about Active Directory, see the following:

- Chapter 9, “Authentication Infrastructure”
- *Windows Server 2008 Active Directory Resource Kit* by Stan Reimer, Mike Mulcare, Conan Kezema, and Byron Wright, with the Microsoft Active Directory Team (Microsoft Press, 2008)
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support

For additional information about PKI, see the following:

- Chapter 9, “Authentication Infrastructure”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>

- Windows Server 2008 Help and Support
- “Public Key Infrastructure for Windows Server” (<http://www.microsoft.com/pki>)
- *Windows Server 2008 PKI and Certificate Security* by Brian Komar (Microsoft Press, 2008)

For additional information about Group Policy, see the following:

- Chapter 9, “Authentication Infrastructure”
- *Windows Group Policy Resource Kit: Windows Server 2008 and Windows Vista* by Derek Melber, Group Policy MVP, with the Windows Group Policy Team (Microsoft Press, 2008)
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Microsoft Windows Server Group Policy” (<http://www.microsoft.com/gp>)

For additional information about RADIUS and NPS, see the following:

- Chapter 9, “Authentication Infrastructure”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Network Policy Server” (<http://www.microsoft.com/nps>)

For additional information about NAP and VPN Enforcement, see the following:

- Chapter 14, “Network Access Protection Overview”
- Chapter 15, “Preparing for Network Access Protection”
- Chapter 18, “VPN Enforcement”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Network Access Protection” (<http://www.microsoft.com/nap>)

## Chapter 14

# Network Access Protection Overview

This chapter describes the need for the new Network Access Protection (NAP) platform in the Windows Server 2008, Windows Vista, and Windows XP SP3 operating systems, the components of NAP on an example intranet, and how NAP works for different types of NAP enforcement methods.

This chapter assumes that you understand the role of Active Directory, public key infrastructure (PKI), Group Policy, and Remote Authentication Dial-In User Service (RADIUS) elements of a Microsoft Windows-based authentication infrastructure for network access. For more information, see Chapter 9, “Authentication Infrastructure.”

## The Need for Network Access Protection

To understand the need for NAP, it is important to review the measures that must be taken to prevent the spread of malicious software (malware). This section provides an overview of malware threats and methods, malware prevention technologies, and how NAP provides centralized definition, integration, and enforcement of system health requirements to help prevent the exposure to malware on a private network.

## Malware and Its Impact on Enterprise Computing

It is an unfortunate fact of life that modern computer networks are hostile environments. The same computer networking technologies that allow seamless communication between computers for e-mail, file transfers, Web access, and real-time collaboration are also used by malware to access and infect vulnerable computers. Malware is designed to install on a computer without the knowledge or consent of the computer user for the purposes of damage, data access, to report on the activities of the computer, or to allow the computer to be controlled by other computers. Malware can take the form of computer viruses (programs that propagate from one computer to another through media exchange or automatically over a network), Trojan horses (malware concealed inside programs that have another primary purpose), spyware (malware that records and reports on how the computer is being used), or adware (malware that displays advertising material to the user).

The Internet is an especially hostile environment, where a vulnerable computer can be attacked and infected in minutes by address and port scanning malware. Home networks also can be hostile environments because home computers are more likely to be vulnerable not only to address-scanning and port-scanning malware but also to malware that is installed on

home computers through Trojan horse techniques such as e-mail attachments, Web controls, and free software exchanged through the computer enthusiast community.

Private organization networks, also known as intranets, are less hostile because they are typically not directly connected to the Internet. Additionally, at least for enterprise networks, an information technology (IT) staff has typically deployed malware prevention software. However, enterprise networks are still vulnerable to infection by Trojan horse–based malware that is downloaded and installed by users from the Internet.

## How Malware Enters the Enterprise Network

Typical enterprise networking environments are not directly connected to the Internet. There is a small set of computers that are directly connected to the Internet to provide Internet services to customers or business partners. Most intranet computers are separated from the Internet by perimeter systems such as firewalls and proxy servers. Therefore, the computers of the enterprise network are typically protected from scanning attacks by network-level viruses emanating from the Internet.

However, the following can circumvent the perimeter security provided by firewalls or proxy servers:

- **Trojan horse–based viruses that are installed through code that is executed on a computer** Users on the enterprise network can inadvertently obtain viruses from e-mail, Web pages, and other types of files that are downloaded from the Internet. E-mail attachments are a common method of delivering Trojan horse–based viruses. Web pages are another common method because the proxy server for Internet Web access is designed to transfer the files that comprise a Web page. Enterprise network users can obtain viruses from Web pages and their associated files.
- **Mobile computers that can be moved and connected to other networks** The obvious example of a mobile computer is a laptop computer. A user takes a laptop home, on business trips, and to other public network locations such as wireless hot spots. Each time the user connects the laptop computer to a network that is not the enterprise network, the laptop runs the risk of being exposed to network-level viruses.
- **Employee remote access** When employees use remote access connections to connect to an enterprise network, they are logically connected to the enterprise network as if there were an Ethernet cable from the employee’s location to a switch port on the enterprise network. Through this logical connection, the organization network can be exposed to network-level viruses.
- **Guest computers** When guests of the organization—such as consultants, vendors, or business partners—connect their computers to the organization network, they can expose it to network-level viruses.

## Malware Impact

Malware can have a direct financial impact on networking operations for both the Internet and private networks because of exposure of confidential information, loss of intellectual property, bandwidth consumed, lost productivity to computers that have become unusable because of the malware, and the time required to remove the malware from all the infected computers. Malware has disrupted networking communications in the past and has the potential of doing so in the future.

## Preventing Malware on Enterprise Networks

Based on previous malware infections (such as Love Bug in 2000 and Code Red in 2001), the IT industry began to work to prevent future infections. The result is a set of malware prevention technologies and techniques that many organization networks and end users employ today.

### Malware Prevention Technologies

Because malware is inherently software, malware prevention software has evolved to prevent its installation and spread. Malware prevention software has the following forms:

- **Antivirus** Software that monitors for known malware in files copied or downloaded to a computer. Antivirus software typically uses a local database of known signatures that identify malware stored in files and e-mail. If malware is detected, the antivirus software can remove the malware or prevent the file from being stored or executed. Because new viruses are created and distributed, the database of known antivirus signatures must be periodically updated.
- **Antispam** Software that prevents unwanted e-mail messages from being stored in your e-mail inbox. Spam is a very common way to spread viruses or spyware.
- **Antispyware** Software that detects and removes known spyware and adware from your computer. Just like antivirus software, antispyware software must be periodically updated to prevent new spyware from being installed. An example of antispyware software is Windows Defender from Microsoft, included with Windows Vista.

In addition to malware prevention software, the following technologies also help prevent malware:

- **Automatic updates for Windows-based computers** For computers running a version of Windows, some types of viruses are designed to exploit a known security issue that has been identified by Microsoft and for which a security update is available. The virus attempts to infect those computers that have not yet been updated. To automate the installation of security updates from Microsoft before virus writers have a chance to write malware and spread it across the Internet, current versions of Windows support automatic updates. Based on a user-specified schedule, a computer running the

Windows Vista, Windows Server 2008, Windows XP, or Windows Server 2003 operating systems can poll the Windows Update Web site and download the latest security updates and automatically install them. Windows Update reduces the administrative burden on IT administrators to keep their computers current with the latest operating system updates.

- **Host-based stateful firewalls** A host-based stateful firewall runs on a computer and monitors network traffic at the packet level to help prevent malicious traffic from being either received or sent by the computer. Some viruses attempt to automatically propagate themselves by scanning the local subnet for available computers and then attacking the computers that are found. If successful, the virus automatically propagates from one computer to another. If an infected computer is moved, the virus begins attacking the computers on the newly attached subnet. An example is when a laptop computer that was infected on a home network is plugged into an organization's private network.

A stateful host-based firewall, such as Windows Firewall included with Windows Vista, Windows Server 2008, Windows XP SP2, and Windows Server 2003 SP1 or SP2, discards all unsolicited incoming traffic that does not correspond to either traffic sent in response to a request of the computer (solicited traffic) or unsolicited traffic that has been specified as allowed (excepted traffic). An example of solicited incoming traffic is the traffic corresponding to a Web page requested by a user of the computer. An example of excepted traffic is traffic that is allowed because the computer is running a server service, such as a Web server, and must receive unsolicited requests.

Because typical network-based viruses rely on unsolicited incoming traffic to scan and attack computers, enabling a host-based stateful firewall on all computers connected to the Internet and an intranet can help prevent the spread of these types of viruses.

To prevent malware from entering and spreading on an enterprise network, IT administrators should do the following:

- Ensure that your host computers are using the correct privilege levels for network services and user accounts. By minimizing the privilege level, you can help prevent malware from installing itself on and exploiting a host computer. For example, computers running Windows Vista use User Account Control (UAC) to reduce the risk of exposure by limiting administrator-level access to processes requiring authorization.
- Use malware prevention software and keep it updated.
- Enable automatic update to install Windows updates as they become available. An organization network can also deploy approved updates through a central server, such as through Windows Server Updates Services (WSUS).
- Use a host-based stateful firewall, such as Windows Firewall, to help prevent infection by network-level viruses that depend on unsolicited incoming traffic.

## Computer System Health and Monitoring

The use of malware prevention technologies brings to light a new issue for IT administrators to determine and monitor: the system health of computers on the intranet. The system health is defined by a computer's current configuration state, which includes the set of installed malware prevention technologies, their current state (such as *enabled* or *disabled* and *current* or *delinquent* with the latest updates), and other configuration settings.

**Determining System Health Requirements** The definition of system health will vary based on an organization's installed malware prevention technologies, computer configuration settings, and other security requirements. To help set the parameters of required system health, an IT administrator should consider the following:

- Antivirus software
  - ❑ Is an antivirus program deployed throughout the organization network?
  - ❑ If so, how current must the antivirus signature file or other updates be for a computer to be considered healthy?
- Antispam software
  - ❑ Is an antispam program deployed throughout the organization network?
  - ❑ If so, how current should the antispam updates be for a computer to be considered healthy?
- Antispyware software
  - ❑ Is an antispyware program deployed throughout the organization network?
  - ❑ If so, how current should the antispyware updates be for a computer to be considered healthy?
- Automatic operating system updates
  - ❑ Is Windows Automatic Update used throughout the organization network?
  - ❑ If so, must automatic updates be enabled for a computer to be considered healthy?
  - ❑ How current do the installed updates have to be for a computer to be considered healthy?
- Host-based stateful firewall
  - ❑ Is a host-based stateful firewall deployed throughout the organization network?
  - ❑ If so, must the firewall be enabled for a computer to be considered healthy? Which exceptions can be configured for a computer to be considered healthy?
- Other configuration settings
  - ❑ Are there other configuration settings required for adherence to the organization's security policies?
  - ❑ If so, which settings are required for a computer to be considered healthy?



For example, an IT administrator can create a system health policy that requires that all computers meet all the following requirements:

- All critical operating system updates must have been installed as of a specific date.
- The antivirus software must have been installed and be running to monitor incoming and outgoing files.
- The most recent signature for the antivirus software must have been installed.
- The antispyware software must have been installed and be running to monitor running services and incoming files.
- The most recent updates to the antispyware software must have been installed.
- The antispyware software must have been installed and be running to monitor incoming e-mail messages.
- The most recent updates to the antispyware software must have been installed.
- The host-based stateful firewall has been installed and is enabled.
- The host-based firewall must have an approved list of exceptions.
- The Transmission Control Protocol/Internet Protocol (TCP/IP) protocol stack on the computer must have IP routing disabled.
- The TCP/IP protocol stack on the computer must have automatic configuration enabled.

However, the biggest problem facing IT administrators is not in setting the requirements for system health but ensuring that all the computers on the organization network meet those requirements and implementing an enforcement mechanism for those computers that do not meet the requirements.

**Enforcing System Health Requirements** Coupled with the problem of determining whether the requirements for system health are being met is enforcing system health requirements for the computers on an organization network. In other words, if a computer on the organization network does not meet the requirements for system health, there should be consequences. For example, a computer that is not compliant with system health requirements should not be allowed to communicate with other computers on the network.

Although most malware prevention software has its own mechanisms for keeping current, there is no enforcement of system health requirements. For example, if an antivirus program does not have the latest updates, there are no consequences for the computer and the user of the computer.

To make system health enforceable, there must be a central computer on the intranet that evaluates system health and is configured with the organization's system health requirements. Client computers that attempt to connect to communicate on the network must have their system health evaluated so that noncompliant computers can be detected. The central

system health evaluation computer must impose a consequence on noncompliant computers. An obvious consequence for a noncompliant computer is that it is refused a connection to the network. However, this dire consequence does not allow the noncompliant computer an opportunity to correct its configuration state.

Rather than preventing all access to the intranet, a solution that allows the noncompliant computer to correct its state, an action known as *remediation*, is to allow limited access to a subset of intranet servers that contain the needed updates, software, scripts, or other resources. Examples of servers on this limited access logical network can include antivirus or software update servers. By using these resources and instructions from the central computer that is evaluating system health, a noncompliant computer can automatically correct its configuration.

## The Role of NAP

NAP for Windows Server 2008, Windows Vista, and Windows XP SP3 provides components and an application programming interface (API) set that can help IT administrators enforce compliance with health requirement policies for network access or communication. With NAP, developers and administrators can create solutions for validating computers that connect to their networks, provide needed updates or access to required health update resources, and limit the access or communication of noncompliant computers. Third-party vendors can leverage the powerful capabilities of NAP to create custom solutions for enforcing system health requirements. Administrators can customize the health maintenance solution they develop and deploy, whether for monitoring the computers accessing the network for health policy compliance, automatically updating computers with software updates to meet health policy requirements, or limiting the access of computers that do not meet health policy requirements.

With NAP, Windows-based networks now have an infrastructure that allows the following:

- IT administrators can configure system health requirements for NAP-capable computers.
- IT administrators can specify access enforcement behaviors for NAP-capable and non-NAP-capable computers, which include the following:
  - ❑ Monitoring of the access and communication attempts of computers and recording the access attempts in server event logs for ongoing or forensic analysis
  - ❑ Enforcement of network access restrictions for noncompliant or non-NAP-capable computers
- NAP-capable computers can automatically update themselves to become compliant (upon initial network access or communication) and remain compliant (automatically download updates or change settings on an ongoing basis).

## Aspects of NAP

NAP has three important and distinct aspects:

- **Health state validation** When a computer attempts to connect to the network, the computer's health state is validated against the health requirement policies as specified by the administrator. Administrators can also specify what to do if a computer is not compliant. In a monitoring-only environment, all computers have their health state evaluated, and the compliance state of each computer is logged for analysis. In a limited access environment, computers that comply with the health requirement policies are allowed unlimited access to the network. Computers that do not comply with health requirement policies can have their access limited.
- **Health policy compliance** Administrators can help ensure compliance with health requirement policies by configuring settings to automatically update noncompliant computers with missing software updates or configuration changes through separate management software products, such as Microsoft Systems Management Server or Microsoft System Center Configuration Manager 2007. In a monitoring-only environment, computers will have access to the network before they are updated with required updates or configuration changes. In a limited access environment, noncompliant computers have limited access until the updates and configuration changes are completed. In both environments, computers that are compatible with NAP can automatically become compliant, and administrators can specify exceptions for computers that are not compatible with NAP.
- **Limited access** Administrators can protect their networks by limiting the access of noncompliant computers, as specified by the administrator. Administrators can create a restricted network containing health update resources and other servers, and noncompliant computers can only access the restricted network. Administrators can also configure exceptions so that computers that are not compatible with NAP do not have their network access limited.

## Typical NAP Scenarios

NAP helps provide a solution for the following common needs:

- **Verification of the health state of roaming laptops** Portability and flexibility are two primary advantages of laptops, but these features also present a health threat. Company laptops frequently leave and return to the company network. While laptops are away from the company, they might not receive the most recent software updates or configuration changes. Laptops might also become infected while they are exposed to unprotected networks such as the Internet. By using NAP, network administrators can check the health state of any laptop when it reconnects to the company network, whether by creating a virtual private network (VPN) connection to the company network or by physically returning to the office.

- **Verification of the health state of desktop computers** Although desktop computers do not usually leave the premises, they still can present a threat to a network. To minimize this threat, administrators must maintain these computers with the most recent updates and required software. Otherwise, these computers are at higher risk of infection from Web sites, e-mail, files from shared folders, and other publicly accessible resources. By using NAP, network administrators can automate health state checks to verify each desktop computer's compliance with health requirement policies. Administrators can check log files to determine which computers do not comply. With the addition of management software, administrators can generate automatic reports and automatically update noncompliant computers. When administrators change health requirement policies, computers can be automatically provided with the most recent updates.
- **Verification of the health state of visiting laptops** Organizations sometimes must allow consultants, business partners, and guests to connect to their private networks. The laptops that these visitors bring might not meet system health requirements and can present health risks. By using NAP, administrators can determine that the visiting laptops are not compliant and allow only access to the Internet. Administrators would not typically require or provide any updates or configuration changes to the visiting laptops.
- **Verification of the health state of unmanaged home computers** Unmanaged home computers that are not a member of the company's Active Directory domain can connect to a managed company network through a VPN connection. Unmanaged home computers provide an additional challenge to administrators because they do not have physical access to these computers. Lack of physical access makes enforcing compliance with health requirements, such as the use of antivirus software, even more difficult. However, with NAP, network administrators can verify the health state of a home computer every time it makes a VPN connection to the company network and limit the access to a restricted network until system health requirements are met.

## Extensibility of NAP

NAP is an extensible platform that provides an infrastructure and an API set for adding components that verify and amend a computer's health state and that enforce access restrictions. For a more detailed explanation of NAP architecture and its extensibility, see "Network Access Protection Platform Architecture" at <http://go.microsoft.com/fwlink/?LinkID=90197>.

## Limitations of NAP

NAP is not designed to protect a network from malicious users. It is designed to help administrators automatically maintain the health of the computers on the network, which in turn helps maintain the network's overall integrity. For example, if a computer has all the software and configuration settings that the health policies require, the computer is compliant and will be granted the appropriate access to the network. NAP does not prevent an authorized

user with a compliant computer from uploading a malicious program to the network or engaging in other inappropriate behavior.

## Business Benefits of NAP

The following are the business benefits of NAP:

- **Lower total cost of ownership through centralized configuration and management of system requirements for connection or communication** NAP provides a central point of configuration to specify the following:

- The system health requirements for computers that are connecting to or communicating on your network, which can include malware prevention, software settings, or system configuration settings.
- The enforcement behavior for computers that do not meet the requirements. Enforcement behavior can be passive, allowing unlimited access but recording each connection or communication attempt; or active, limiting the access of the noncompliant computer.

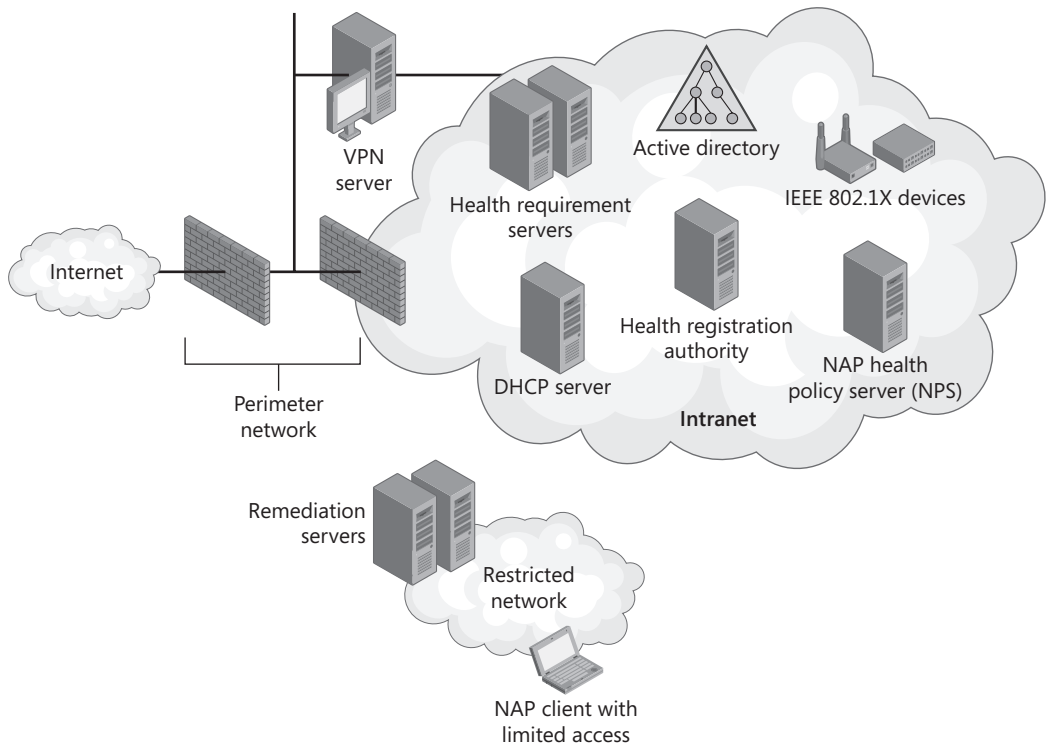
The system requirements and enforcement behavior are centrally configured in the form of health requirement policies on the server that evaluates the client's system settings.

- **Lower total cost of ownership through automated system health or configuration remediation** NAP-capable computers will automatically install updates for their malware prevention software and make required configuration settings prior to being granted unlimited access to the network. Although most malware prevention software periodically checks for updates to install, NAP requires the updates for network connectivity. Once a NAP-capable computer is compliant, NAP components will automatically perform updates to ensure ongoing compliance.
- **Reduced chance of infection by malware** Because the NAP platform can enforce system health requirements, NAP-capable computers can be updated and protected against known malware attacks through operating system and antivirus updates on computers prior to allowing them unlimited access. Appropriately configured NAP-enabled networks will have a reduced exposure to malware.
- **Utilization of existing system health and configuration requirements infrastructure** NAP does not replace your existing system health and configuration infrastructure. Rather, it adds value to the existing components of system health and configuration and extends their role by tying them all together with the common goal of setting and enforcing system health requirements on connecting or communicating computers. Many system configuration, malware prevention, and network security infrastructure vendors support NAP. For a complete list, see Network Access Protection Partners at <http://www.microsoft.com/windowsserver2003/partners/nappartners.mspx>.

## Components of NAP

The following sections describe some of the components of the NAP infrastructure to provide a basic understanding of NAP processes. For a more detailed explanation of NAP components and architecture, see the “Network Access Protection Platform Architecture” white paper at <http://go.microsoft.com/fwlink/?LinkID=90197>.

Figure 14-1 shows the components of a NAP-enabled network infrastructure.



**Figure 14-1** Components of a NAP-enabled network infrastructure

The components of a NAP-enabled network infrastructure consist of the following:

- **NAP clients** Computers that support the NAP platform and include computers running Windows Server 2008, Windows Vista, or Windows XP SP3.
- **NAP enforcement points** Computers or network access devices that use NAP or can be used with NAP to require the evaluation of a NAP client's health state and provide restricted network access or communication. NAP enforcement points use a Network Policy Server (NPS) that is acting as a NAP health policy server to evaluate the health state of NAP clients, whether network access or communication is allowed, and the set

of remediation actions that a noncompliant NAP client must perform. Examples of NAP enforcement points are the following:

- ❑ **Health Registration Authority (HRA)** A computer running Windows Server 2008 and Internet Information Services (IIS) that obtains health certificates from a certification authority (CA) for compliant NAP clients.
  - ❑ **Network access devices** Ethernet switches or wireless access points (APs) that support IEEE 802.1X authentication
  - ❑ **VPN server** A computer running Windows Server 2008 and Routing and Remote Access that allows remote access VPN connections to an intranet
  - ❑ **DHCP server** A computer running Windows Server 2008 and the Dynamic Host Configuration Protocol (DHCP) Server service that provides automatic Internet Protocol version 4 (IPv4) address configuration to intranet clients
- **NAP health policy servers** Computers running Windows Server 2008 and the NPS service that store health requirement policies and provide health state validation for NAP. NPS is the replacement for the Internet Authentication Service (IAS), the Remote Authentication Dial-In User Service (RADIUS) server and proxy provided with Windows Server 2003. NPS can also act as an authentication, authorization, and accounting (AAA) server for network access. When acting as a AAA server or NAP health policy server, NPS is typically run on a separate server for centralized configuration of network access and health requirement policies, as Figure 14-1 shows. The NPS service is also run on Windows Server 2008–based NAP enforcement points, such as an HRA or DHCP server. However, in these configurations, the NPS service is acting as a RADIUS proxy to exchange RADIUS messages with a NAP health policy server.
  - **Health requirement servers** Computers that provide current system health state for NAP health policy servers. For example, a health requirement server for an antivirus program tracks the latest version of the antivirus signature file.
  - **Active Directory Domain Services** The Windows directory service that stores account credentials and properties and Group Policy settings. Although not required for health state validation, Active Directory is required for Internet Protocol Security (IPsec)–protected communications, 802.1X-authenticated connections, and remote access VPN connections.
  - **Restricted network** A separate logical or physical network that contains:
    - ❑ **Remediation servers** Network infrastructure servers and health update servers that NAP clients can access to remediate their noncompliant state. Examples of network infrastructure servers include Domain Name System (DNS) servers and Active Directory domain controllers. Examples of health update servers include antivirus signature distribution servers and software update servers.
    - ❑ **NAP clients with limited access** Computers that are placed on the restricted network when they do not comply with health requirement policies.

- ❑ **Non-NAP-capable computers** Optionally, computers that do not support NAP can be placed on the restricted network (not shown in Figure 14-1).

## System Health Agents and System Health Validators

Components of the NAP infrastructure known as system health agents (SHAs) on NAP clients and system health validators (SHVs) on NAP health policy servers provide health state tracking and validation for attributes of system health. Windows Vista and Windows XP SP3 include a Windows Security Health Validator SHV that monitors the settings of the Windows Security Center. Windows Server 2008 includes the corresponding Windows Security Health Validator SHV. NAP is designed to be flexible and extensible. It can interoperate with any vendor who provides SHAs and SHVs that use the NAP API.

An SHA creates a statement of health (SoH) that contains the current status information about the attribute of health being monitored by the SHA. For example, an SHA for an antivirus program might contain the state of the program (installed and running) and the version of the current antivirus signature file. Whenever an SHA updates its status, it creates a new SoH. To indicate its overall health state, a NAP client uses a System Statement of Health (SSoH), which includes version information for the NAP client and the set of SoHs for the installed SHAs.

When the NAP client validates its system health, it passes its SSoH to the NAP health policy server for evaluation through a NAP enforcement point. The NAP health policy server uses the SSoH, its installed SHVs, and its health requirement policies to determine whether the NAP client is compliant with system health requirements, and if it is not, the remediation actions that must be taken to achieve compliance. Each SHV produces a statement of health response (SoHR), which can contain remediation instructions. For example, the SoHR for an antivirus program might contain the current version number of the antivirus signature file and the name or IP address of the antivirus signature file server on the intranet.

Based on the SoHRs from the SHVs and the configured health requirement policies, the NAP health policy server creates a System Statement of Health Response (SSoHR), which indicates whether the NAP client is compliant or noncompliant and includes the set of SoHRs from the SHVs. The NAP health policy server passes the SSoHR back to the NAP client through a NAP enforcement point. The NAP client passes the SoHRs to its SHAs. The noncompliant SHAs automatically remediate their health state and create updated SoHs, and the health validation process begins again.

## Enforcement Clients and Servers

A NAP Enforcement Client (EC) is a component on a NAP client that requests some level of access to a network, passes the computer's health status to a NAP enforcement point that is providing the network access, and indicates health evaluation information to other components



of the NAP client architecture. The NAP ECs for the NAP platform supplied in Windows Vista, Windows XP SP3, and Windows Server 2008 are the following:

- An IPsec EC for IPsec-protected communications
- An EAPHost EC for 802.1X-authenticated connections
- A VPN EC for remote access VPN connections
- A DHCP EC for DHCP-based IPv4 address configuration
- A TS Gateway EC for connections to a TS Gateway server

A NAP Enforcement Server (ES) is a component on a NAP enforcement point running Windows Server 2008 that allows some level of network access or communication, can pass a NAP client's health status to NPS for evaluation, and, based on the response from NPS, can provide the enforcement of limited network access. The NAP ESs included with Windows Server 2008 are the following:

- An IPsec ES for IPsec-protected communications
- A DHCP ES for DHCP-based IPv4 address configuration
- A TS Gateway ES for TS Gateway server connections

For 802.1X-authenticated and remote access VPN connections, there is no separate ES component running on the 802.1X switch or wireless AP or VPN server.

Together, ECs and ESs require health state validation and enforce limited network access for noncompliant computers for specific types of network access or communication.

## NPS

NPS is a RADIUS server and proxy in Windows Server 2008. As a RADIUS server, NPS provides AAA services for various types of network access. For authentication and authorization, NPS uses Active Directory to verify user or computer credentials and obtain user or computer account properties when a computer attempts an 802.1X-authenticated connection or a VPN connection.

NPS also acts as a NAP health policy server. Administrators set system health requirements in the form of health requirement policies on the NAP health policy server. NAP health policy servers evaluate health state information provided by NAP clients to determine health compliance, and for noncompliance, the set of remediation actions that must be taken by the NAP client to become compliant.

The role of NPS as an AAA server is independent from its role as a NAP health policy server. These roles can be used separately or combined as needed. For example:

- NPS can be an AAA server on an intranet that has not yet deployed NAP.

- NPS can be a combination of AAA server and health policy server for 802.1X-authenticated connections on an intranet that has deployed NAP for 802.1X-authenticated connections.
- NPS can be a health policy server for DHCP configuration on an intranet that has deployed NAP for DHCP configuration.

For more information about NPS and RADIUS, see Chapter 9.

## Enforcement Methods

Windows Vista, Windows XP SP3, and Windows Server 2008 include NAP support for the following types of network access or communication:

- IPsec-protected traffic
- IEEE 802.1X-authenticated network connections
- Remote access VPN connections
- DHCP address configurations

Windows Server 2008 and Windows Vista also include NAP support for connections to a TS Gateway server.

Administrators can use these types of network access or communication, known as *NAP enforcement methods*, separately or together to limit the access or communication of noncompliant computers. NPS acts as a health policy server for all these NAP enforcement methods.

The following sections describe the IPsec, 802.1X, VPN, and DHCP enforcement methods.

## IPsec Enforcement

With IPsec enforcement, a computer must be compliant to initiate communications with other compliant computers on an intranet in a server isolation or domain isolation IPsec deployment, which require that incoming communications be protected with IPsec. Because IPsec enforcement utilizes IPsec, you can specify requirements for protected communications with compliant computers on a per-IP address or per-TCP/UDP port number basis. IPsec enforcement confines communication to compliant computers after they have successfully connected and obtained a valid IP address configuration. IPsec enforcement one of the strongest forms of limited network access or communication in NAP.

The components of IPsec enforcement consist of an IPsec ES on an HRA running Windows Server 2008 and an IPsec EC in Windows Vista, Windows XP SP3, or Windows Server 2008. The HRA obtains X.509-based health certificates for NAP clients when they prove that they are compliant. These health certificates are then used in conjunction with IPsec policy settings to authenticate NAP clients when they initiate IPsec-protected communications with other compliant NAP clients on an intranet.

For more information about server isolation and domain isolation with IPsec, see Chapter 4, “Windows Firewall with Advanced Security.”

## 802.1X Enforcement

With 802.1X enforcement, a computer must be compliant to obtain unlimited network access through an 802.1X-authenticated network connection, such as to an authenticating Ethernet switch or an IEEE 802.11 wireless AP. For noncompliant computers, network access is limited through a restricted access profile placed on the connection by the Ethernet switch or wireless AP. The restricted access profile can specify an access control list (ACL), which corresponds to a set of IP packet filters configured on the Ethernet switch or wireless AP, or a virtual LAN (VLAN) identifier (ID) that corresponds to the restricted network VLAN. With 802.1X enforcement, health policy requirements are enforced every time a computer attempts an 802.1X-authenticated network connection. 802.1X enforcement also actively monitors the health status of the connected NAP client and applies the restricted access profile to the connection if the client becomes noncompliant.

The components of 802.1X enforcement consist of NPS in Windows Server 2008 and an EAPHost EC in Windows Vista, Windows XP SP3, and Windows Server 2008. 802.1X enforcement provides strong limited network access for all computers accessing the network through an 802.1X-authenticated connection.

## VPN Enforcement

With VPN enforcement, a computer must be compliant to obtain unlimited network access through a remote access VPN connection. For noncompliant computers, network access is limited through a set of IP packet filters that are applied to the VPN connection by the VPN server. With VPN enforcement, health policy requirements are enforced every time a computer attempts to obtain a remote access VPN connection to the network. VPN enforcement also actively monitors the health status of the NAP client and applies the IP packet filters for the restricted network to the VPN connection if the client becomes noncompliant.

The components of VPN enforcement consist of NPS in Windows Server 2008 and a VPN EC that is part of the remote access client in Windows Vista, Windows XP SP3, and Windows Server 2008. VPN enforcement provides strong limited network access for all computers accessing the network through a remote access VPN connection.



**Note** VPN enforcement with NAP is different than Network Access Quarantine Control, a feature in Windows Server 2003.

## DHCP Enforcement

With DHCP enforcement, a computer must be compliant to obtain an IPv4 address configuration that has unlimited network access from a DHCP server. For noncompliant computers,

network access is limited by an IPv4 address configuration that allows limited access only to the restricted network. With DHCP enforcement, health policy requirements are enforced every time a DHCP client attempts to lease or renew an IPv4 address configuration. DHCP enforcement also actively monitors the health status of the NAP client and renews the IPv4 address configuration for access only to the restricted network if the client becomes non-compliant.

The components of DHCP enforcement consist of a DHCP ES that is part of the DHCP Server service in Windows Server 2008 and a DHCP EC that is part of the DHCP Client service in Windows Vista, Windows XP SP3, and Windows Server 2008. Because DHCP enforcement relies on a limited IPv4 address configuration that can be overridden by a user with administrator-level access, it is a weak form of limited network access in NAP.

## How NAP Works

NAP is designed so that administrators can configure it to meet the individual needs of their networks. Therefore, the actual configuration of NAP will vary according to the administrator's preferences and requirements. However, the underlying operation of NAP remains the same. This section describes how NAP works on the example intranet shown in Figure 14-1. This example intranet is configured for the following:

- Health state validation, health policy compliance, and limited network access for non-compliant NAP clients
- IPsec enforcement, 802.1X enforcement, VPN enforcement, and DHCP enforcement

When obtaining a health certificate, making an 802.1X-authenticated or VPN connection to the intranet, or leasing or renewing an IPv4 address configuration from the DHCP server, each NAP client is classified in one of the following ways:

- NAP clients that meet the health policy requirements are classified as compliant and are allowed unlimited access to the intranet.
- NAP clients that do not meet the health policy requirements are classified as noncompliant and have their access limited to the restricted network until they meet the requirements. A noncompliant NAP client does not necessarily have a virus or some other active threat to the intranet, but it does not have the software updates or configuration settings as required by health requirement policies. A noncompliant NAP client is at higher risk of being compromised and passing on that risk to the intranet. The SHAs on NAP clients can automatically update computers with limited access with the software or configuration settings required for unlimited access. Automatic remediation ensures that noncompliant NAP clients obtain the necessary updates and are granted unlimited access as quickly as possible.

The example intranet in Figure 14-1 contains a restricted network. A restricted network can be created logically or physically. For example, IP filters, static routes, an ACL, or a VLAN

identifier can be placed on a NAP client's connection to specify the remediation servers with which they can communicate.

Because most intranets contain a heterogeneous mixture of computers and devices, an administrator might choose to exempt some computers or devices from health policy requirements, for example, computers that require unlimited intranet access and are running Windows Server 2003, Windows 2000 or older versions of Windows, and operating systems other than Windows that do not support NAP. To prevent limited access for these computers, an administrator can optionally configure health requirement policies to grant unlimited access to the intranet for specific non-NAP-capable computers. Ideally, you should update or upgrade your non-NAP-capable computers to support NAP so that all of your computers can have their system health evaluated.

An administrator can also configure an exception policy on the NAP health policy server; exempted computers are not checked for compliance and have unlimited access to the intranet.

The following sections describe the basic processes for IPsec enforcement, 802.1X enforcement, VPN enforcement, and DHCP enforcement for a NAP client.

## How IPsec Enforcement Works

The following process describes how IPsec enforcement works for a NAP client that is starting on the example intranet shown in Figure 14-1:

1. The IPsec EC component sends its SSoH indicating its current health state to the HRA.
2. The HRA sends the NAP client's SSoH to the NAP health policy server.
3. The NAP health policy server evaluates the SSoH of the NAP client, determines whether the NAP client is compliant, and sends the resulting SSoHR to the HRA. If the NAP client is not compliant, the SSoHR includes health remediation instructions.
4. If the health state is compliant, the HRA obtains a health certificate for the NAP client. Based on its IPsec policy settings as configured by the administrator, the NAP client can now initiate IPsec-protected communication with other compliant computers using its health certificate for IPsec authentication, and it can respond to communications initiated from other compliant computers that authenticate using their own health certificate.
5. If the health state is not compliant, the HRA sends the SSoHR to the NAP client and does not issue a health certificate. The NAP client cannot initiate communication with other computers that require a health certificate for IPsec authentication. However, the NAP client can initiate communications with remediation servers to correct its health state.
6. The NAP client sends update requests to the appropriate remediation servers.
7. The remediation servers provide the NAP client with the required updates for compliance with health requirements. The NAP client updates its SSoH.

8. The NAP client sends its updated SSoH to the HRA.
9. Assuming that all the required updates were made, the NAP health policy server determines that the NAP client is compliant and sends the SSoHR indicating health compliance to the HRA.
10. The HRA obtains a health certificate for the NAP client. The NAP client can now initiate IPsec-protected communication with other compliant computers.

For information about deploying IPsec enforcement, see Chapter 15, “Preparing for Network Access Protection,” and Chapter 16, “IPsec Enforcement.”

## How 802.1X Enforcement Works

The following process describes how 802.1X enforcement works for a NAP client that is initiating an 802.1X-authenticated connection on the example intranet shown in Figure 14-1:

1. The NAP client and the Ethernet switch or wireless AP begin 802.1X authentication.
2. The NAP client sends its user or computer authentication credentials to the NAP health policy server.
3. If the authentication credentials are valid, the NAP health policy server requests the health state from the NAP client. If the authentication credentials are not valid, the connection attempt is terminated.
4. The NAP client sends its SSoH to the NAP health policy server.
5. The NAP health policy server evaluates the SSoH of the NAP client, determines whether the NAP client is compliant, and sends the results to the NAP client and the Ethernet switch or wireless AP. If the NAP client is not compliant, the results include a limited access profile for the Ethernet switch or wireless AP and the SSoHR containing health remediation instructions for the NAP client.
6. If the health state is compliant, the Ethernet switch or wireless AP completes the 802.1X authentication, and the NAP client has unlimited access to the intranet.
7. If the health state is not compliant, the Ethernet switch or wireless AP completes the 802.1X authentication but limits the access of the NAP client to the restricted network through an ACL or a VLAN ID. The NAP client can send traffic only to the remediation servers on the restricted network.
8. The NAP client sends update requests to the remediation servers.
9. The remediation servers provide the NAP client with the required updates for compliance with health requirement policies. The NAP client updates its SSoH.
10. The NAP client restarts 802.1X authentication and sends its updated SSoH to the NAP health policy server.

11. Assuming that all the required updates were made, the NAP health policy server determines that the NAP client is compliant and instructs the Ethernet switch or wireless AP to allow unlimited access.
12. The Ethernet switch or wireless AP completes the 802.1X authentication, and the NAP client has unlimited access to the intranet.

For information about deploying 802.1X enforcement, see Chapter 15 and Chapter 17, “802.1X Enforcement.”

## How VPN Enforcement Works

The following process describes how VPN enforcement works for a NAP client that is initiating a VPN connection on the example intranet shown in Figure 14-1:

1. The NAP client initiates a connection to the VPN server.
2. The NAP client sends its user authentication credentials to the VPN server.
3. If the authentication credentials are valid, the NAP health policy server requests the health state from the NAP client. If the authentication credentials are not valid, the VPN connection attempt is terminated.
4. The NAP client sends its SSoH to the NAP health policy server.
5. The NAP health policy server evaluates the SSoH of the NAP client, determines whether the NAP client is compliant, and sends the results to the NAP client and the VPN server. If the NAP client is not compliant, the results include a set of packet filters for the VPN server and the SSoHR containing health remediation instructions for the NAP client.
6. If the health state is compliant, the VPN server completes the VPN connection, and the NAP client has unlimited access to the intranet.
7. If the health state is not compliant, the VPN server completes the VPN connection but, based on the packet filters, limits the access of the NAP client to the restricted network. The NAP client can send traffic only to the remediation servers on the restricted network.
8. The NAP client sends update requests to the remediation servers.
9. The remediation servers provide the NAP client with the required updates for compliance with health requirement policies. The NAP client updates its SSoH.
10. The NAP client restarts authentication with the VPN server and sends its updated SSoH to the NAP health policy server.
11. Assuming that all the required updates were made, the NAP health policy server determines that the NAP client is compliant and instructs the VPN server to allow unlimited access.

12. The VPN server completes the VPN connection, and the NAP client has unlimited access to the intranet.

For information about deploying VPN enforcement, see Chapter 15 and Chapter 18, “VPN Enforcement.”

## How DHCP Enforcement Works

The following process describes how DHCP enforcement works for a NAP client that is attempting an initial DHCP configuration on the example intranet shown in Figure 14-1:

1. The NAP client sends a DHCP request message containing its SSoH to the DHCP server.
2. The DHCP server sends the SSoH of the NAP client to the NAP health policy server.
3. The NAP health policy server evaluates the SSoH of the NAP client, determines whether the NAP client is compliant, and sends the results to the DHCP server. If the NAP client is not compliant, the results include a limited access configuration for the DHCP server and an SSoHR containing health remediation instructions for the NAP client.
4. If the health state is compliant, the DHCP server assigns an IPv4 address configuration for unlimited access to the NAP client and completes the DHCP message exchange.
5. If the health state is not compliant, the DHCP server assigns an IPv4 address configuration for limited access to the restricted network to the NAP client and completes the DHCP message exchange, sending the SSoHR to the NAP client. The NAP client can send traffic only to the remediation servers on the restricted network.
6. The NAP client sends update requests to the remediation servers.
7. The remediation servers provide the NAP client with the required updates for compliance with health requirement policies. The NAP client updates its SSoH.
8. The NAP client sends a new DHCP request message containing its updated SSoH to the DHCP server.
9. The DHCP server sends the updated SSoH of the NAP client to the NAP health policy server.
10. Assuming that all the required updates were made, the NAP health policy server determines that the NAP client is compliant and instructs the DHCP server to assign an IPv4 address configuration for unlimited access to the intranet.
11. The DHCP server assigns an address configuration for unlimited access to the NAP client and completes the DHCP message exchange.

For information about deploying DHCP enforcement, see Chapter 15 and Chapter 19, “DHCP Enforcement.”



### How It Works: NAP Component Interaction

System health information, in the form of SSoHs and SSoHRs, between a NAP health policy server and a NAP enforcement point is sent as attributes of a RADIUS message. A NAP health policy server is a RADIUS server, and NAP enforcement points are RADIUS clients.

For IPsec enforcement, system health information between a NAP client and an HRA is sent over Hypertext Transfer Protocol (HTTP) or an encrypted HTTP over Secure Sockets Layer (SSL) session. The NAP client uses HTTP or the HTTP over SSL session to indicate its current system health state and request a health certificate. The HRA uses HTTP or the HTTP over SSL session to send the SSoHR and the health certificate to the NAP client.

For 802.1X enforcement, system health information between a NAP client and a NAP health policy server is sent as Protected Extensible Authentication Protocol (PEAP)-Type-Length-Value (TLV) messages. On the link between the NAP client and the authenticating switch or wireless AP, the PEAP-TLV messages are sent over the EAP over LAN (EAPOL) protocol. Between the authenticating switch or wireless AP and the NAP health policy server, the PEAP-TLV messages are encapsulated and sent as RADIUS attributes of RADIUS messages.

For VPN enforcement, system health information between a NAP client and a NAP health policy server is also sent as PEAP-TLV messages. The PEAP-TLV messages are sent over the Point-to-Point Protocol (PPP)-based logical link between the NAP client and the VPN server created by the VPN connection. Between the VPN server and the NAP health policy server, the PEAP-TLV messages are encapsulated and sent as RADIUS attributes of RADIUS messages.

For DHCP enforcement, system health information between a NAP client and a DHCP server is sent as DHCP options in DHCP messages.

## Chapter Summary

NAP is a new platform for Windows Vista, Windows Server 2008, and Windows XP SP3 that includes client and server components to limit the network access or communication of computers until they are compliant with system health requirements. Administrators can configure IPsec enforcement, 802.1X enforcement, VPN enforcement, DHCP enforcement, or all of them, depending on their needs.

IPsec enforcement works by not issuing health certificates to noncompliant NAP clients so that they cannot initiate protected communications with compliant NAP clients. 802.1X enforcement is done by specifying an ACL or VLAN ID that is applied to the 802.1X connection

by the Ethernet switch or wireless AP to limit the access to the restricted network. VPN enforcement is done through IP packet filters that are applied to the VPN connection by the VPN server to limit the access to the restricted network. DHCP enforcement is done through an IPv4 address configuration that limits access to the restricted network.

## Additional Information

For additional information about NAP, see the following:

- Chapter 15, “Preparing for Network Access Protection”
- Chapter 16, “IPsec Enforcement”
- Chapter 17, “802.1X Enforcement”
- Chapter 18, “VPN Enforcement”
- Chapter 19, “DHCP Enforcement”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Network Access Protection” (<http://www.microsoft.com/nap>)

For additional information about RADIUS and NPS, see the following:

- Chapter 9, “Authentication Infrastructure”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Microsoft Network Policy Server” (<http://www.microsoft.com/nps>)

For additional information about IPsec, see the following:

- Chapter 4, “Windows Firewall with Advanced Security”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- IPsec (<http://www.microsoft.com/ipsec>)

For additional information about IEEE 802.1X for wireless and wired networks, see the following:

- Chapter 10, “IEEE 802.11 Wireless Networks”
- Chapter 11, “IEEE 802.1X-Authenticated Wired Networks”

- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Wireless Networking” (<http://www.microsoft.com/wifi>)
- “Wired Networking with 802.1X Authentication” (<http://technet.microsoft.com/en-us/network/bb545365.aspx>)

For additional information about remote access VPNs, see the following:

- Chapter 12, “Remote Access VPN Connections”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Virtual Private Networks” (<http://www.microsoft.com/vpn>)

For additional information about DHCP, see the following:

- Chapter 3, “Dynamic Host Configuration Protocol”
- Windows Server 2008 Technical Library at <http://technet.microsoft.com/windowsserver/2008>
- Windows Server 2008 Help and Support
- “Dynamic Host Configuration Protocol” (<http://www.microsoft.com/dhcp>)