

CompTIA® Network+

Exam N10-005

Rapid Review



Craig Zacker

Rapid Review

CompTIA Network+ Exam N10-005

Assess your readiness for CompTIA Network+ Exam N10-005—and quickly identify where you need to focus and practice. This practical, streamlined guide walks you through each exam objective, providing “need to know” checklists, review questions, tips, and links to further study—all designed to help bolster your preparation

Reinforce your exam prep with a *Rapid Review* of these objectives:

- Network Concepts
- Network Installation and Configuration
- Network Media and Topologies
- Network Management
- Network Security



This book is an ideal complement to the in-depth training of the Microsoft Press® *Training Kit* and other exam-prep resources for CompTIA Network+ Exam N10-005.

ISBN: 978-0-7356-6683-2



9 0 0 0 0

9 7 8 0 7 3 5 6 6 6 8 3 2

U.S.A. \$29.99

Canada \$31.99

[Recommended]

*Certification/
CompTIA Network+*



ABOUT THE AUTHOR

Craig Zacker, an editor and educator, has written or contributed to dozens of books on networking, operating systems, and PC hardware, including *CompTIA Network+ Training Kit (Exam N10-005)* and *MCITP Self-Paced Training Kit (Exam 70-686): Windows® 7 Desktop Administrator*.



Microsoft®

Microsoft®

CompTIA® Network+®

Rapid Review

(Exam N10-005)

Craig Zacker

Published with the authorization of Microsoft Corporation by:
O'Reilly Media, Inc.
1005 Gravenstein Highway North
Sebastopol, California 95472

Copyright © 2012 by Craig Zacker

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.

ISBN: 978-0-7356-6683-2

1 2 3 4 5 6 7 8 9 LSI 7 6 5 4 3 2

Printed and bound in the United States of America.

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

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

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

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the author, O'Reilly Media, Inc., 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 and Developmental Editor: Kenyon Brown

Production Editor: Rachel Steely

Editorial Production and Illustration: Diane Kohnen, S4Carlisle Publishing Services

Technical Reviewer: Brian Blum

Copyeditor: Andrew Jones

Indexer: Stephen R. Ingle

Cover Design: Best & Company Design

Cover Composition: Karen Montgomery

Contents at a Glance

<i>Introduction</i>	xix
<i>Preparing for the exam</i>	xxiii
Chapter 1 NETWORK CONCEPTS	1
Chapter 2 NETWORK INSTALLATION AND CONFIGURATION	63
Chapter 3 NETWORK MEDIA AND TOPOLOGIES	113
Chapter 4 NETWORK MANAGEMENT	183
Chapter 5 NETWORK SECURITY	241
<i>Appendix</i>	303
<i>Index</i>	311

Contents

<i>Introduction</i>	<i>xix</i>
<i>Preparing for the exam</i>	<i>xxiii</i>
Chapter 1 Network Concepts	1
Objective 1.1: Compare the layers of the OSI and TCP/IP models	1
Exam need to know	2
OSI model	2
TCP/IP model	3
Can you answer these questions?	4
Objective 1.2: Classify how applications, devices, and protocols relate to the OSI model layers	4
Exam need to know	5
MAC addresses	6
IP addresses	7
EUI-64	7
Frames	8
Packets	9
Switches	9
Routers	10
Multilayer switches	11
Hubs	11
Encryption devices	12
Cables	12
NICs	12
Bridges	13
Can you answer these questions?	13

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:

microsoft.com/learning/booksurvey

Objective 1.3: Explain the purpose and properties of IP addressing	14
Exam need to know	14
IP address classes	14
Classless inter-domain routing (CIDR)	16
IPv4 and IPv6 address formatting	17
MAC address formatting	17
IP address subnetting	18
Multicasts, unicasts, and broadcasts	19
Automatic private IP addressing	20
Can you answer these questions?	20
Objective 1.4: Explain the purpose and properties of routing and switching	21
Exam need to know	21
Routing tables	22
Static vs. dynamic routing	23
Routing metrics	23
Next hop	24
Link state vs. distance vector routing	24
RIP	25
EIGRP	26
OSPF	26
Convergence	27
Spanning Tree Protocol	28
Virtual LANs	28
Port mirroring	29
Broadcast domains and collision domains	30
IGP vs. EGP	30
Can you answer these questions?	31
Objective 1.5: Identify common TCP and UDP default ports	32
Exam need to know	32
Ports	33
Can you answer these questions?	34
Objective 1.6: Explain the function of common networking protocols	34
Exam need to know	34
TCP/IP suite	35

TCP	37
UDP	37
DHCP	37
FTP	38
TFTP	38
DNS	39
HTTP	39
HTTPS	40
ARP	40
SIP	40
RTP	41
TELNET	41
SSH	41
NTP	42
POP3	42
IMAP4	43
SMTP	43
SNMP2/3	43
ICMP	44
IGMP	44
TLS	45
Can you answer these questions?	45
Objective 1.7: Summarize DNS concepts and its components . . .	45
Exam need to know	46
DNS servers	46
DNS records	47
Dynamic DNS	48
Can you answer these questions?	49
Objective 1.8: Given a scenario, implement the following network troubleshooting methodology.	49
Exam need to know	49
Identify the problem	50
Establish a theory of probable cause	50
Test the theory to determine the cause	51
Establish a plan of action to resolve the problem and identify potential effects	52
Implement the solution or escalate as necessary	52

Verify full system functionality and, if applicable, implement preventative measures	53
Document findings, actions, and outcomes	53
Can you answer these questions?	54
Objective 1.9: Identify virtual network components.....	54
Exam need to know	54
Virtual desktops	55
Virtual servers	55
Virtual switches	56
Virtual PBX	57
Onsite vs. offsite	57
Network as a Service (NaaS)	58
Can you answer these questions?	58
Answers	58
Objective 1.1: Compare the layers of the OSI and TCP/IP models	58
Objective 1.2: Classify how applications, devices, and protocols relate to the OSI model layers	59
Objective 1.3: Explain the purpose and properties of IP addressing	59
Objective 1.4: Explain the purpose and properties of routing and switching	59
Objective 1.5: Identify common TCP and UDP default ports	60
Objective 1.6: Explain the function of common networking protocols	60
Objective 1.7: Summarize DNS concepts and its components	60
Objective 1.8: Given a scenario, implement the following network troubleshooting methodology	60
Objective 1.9: Identify virtual network components	61
Chapter 2 Network Installation and Configuration	63
Objective 2.1: Given a scenario, install and configure routers and switches	63
Exam need to know	64
Installing and configuring routing tables	64
Installing and configuring NAT	66

Installing and configuring PAT	67
Installing and configuring VLAN trunking	69
Installing and configuring managed vs. un-managed routers and switches	70
Configuring router and switch interfaces	70
Installing and configuring PoE	71
Configuring traffic filtering	71
Installing and configuring diagnostics	72
Configuring VLAN Trunking Protocol (VTP)	72
Installing and configuring QoS	73
Installing and configuring port mirroring	74
Can you answer these questions?	74
 Objective 2.2: Given a scenario, install and configure a wireless network	75
Exam need to know	75
Locating access points	76
Installing antennae	76
Preventing interference	77
Selecting frequencies	78
Configuring channels	78
Understanding wireless standards	80
Disabling SSID broadcasts	81
Selecting compatible standards	82
Can you answer these questions?	83
 Objective 2.3: Explain the purpose and properties of DHCP	83
Exam need to know	83
Static vs. dynamic IP addressing	84
DHCP reservations	85
DHCP scopes	85
DHCP leases	85
DHCP options	87
Can you answer these questions?	88
 Objective 2.4: Given a scenario, troubleshoot common wireless problems	88
Exam need to know	89
Troubleshooting interference problems	89
Troubleshooting signal strength problems	90

Troubleshooting configuration problems	90
Troubleshooting incompatibility problems	91
Troubleshooting incorrect channel problems	92
Troubleshooting latency problems	92
Troubleshooting encryption type problems	93
Troubleshooting bounce problems	93
Troubleshooting SSID mismatch problems	94
Troubleshooting incorrect switch placement problems	94
Can you answer these questions?	95
Objective 2.5: Given a scenario, troubleshoot common router and switch problems	95
Exam need to know	95
Troubleshoot switching loop problems	96
Troubleshoot cable problems	97
Troubleshoot port configuration problems	98
Troubleshoot VLAN assignment problems	99
Troubleshoot MTU problems	99
Troubleshoot power failures	100
Troubleshoot bad or missing routes	101
Troubleshoot bad modules	102
Troubleshoot incorrect subnet masks	102
Troubleshoot incorrect gateways	103
Troubleshoot duplicate IP addresses	103
Troubleshoot incorrect DNS addresses	104
Can you answer these questions?	104
Objective 2.6: Given a set of requirements, plan and implement a basic SOHO network	104
Exam need to know	105
Plan a list of requirements	105
Plan for cable lengths	106
Plan device types	107
Plan for environmental limitations	107
Plan for equipment limitations	108
Plan for compatibility requirements	108
Can you answer these questions?	109
Answers	109

Objective 2.1: Given a scenario, install and configure routers and switches	109
Objective 2.2: Given a scenario, install and configure a wireless network	110
Objective 2.3: Explain the purpose and properties of DHCP	110
Objective 2.4: Given a scenario, troubleshoot common wireless problems	110
Objective 2.5: Given a scenario, troubleshoot common router and switch problems	111
Objective 2.6: Given a set of requirements, plan and implement a basic SOHO network	111
Chapter 3 Network Media and Topologies	113
Objective 3.1: Categorize standard media types and associated properties	113
Exam need to know	113
Categorize fiber media	114
Categorize copper media	115
Categorize plenum vs. non-plenum cables	119
Categorize media converters	119
Categorize cable distance and speed limitations	120
Categorize broadband over powerline networking	121
Can you answer these questions?	122
Objective 3.2: Categorize standard connector types based on network media	123
Exam need to know	123
Categorize fiber connector types	123
Categorize copper connector types	124
Can you answer these questions?	127
Objective 3.3: Compare and contrast different wireless standards	128
Exam need to know	128
Compare and contrast the 802.11a/b/g/n wireless standards	128
Can you answer these questions?	130

Objective 3.4: Categorize WAN technology types and properties	130
Exam need to know	131
Categorize WAN technology types	131
Categorize WAN technology properties	144
Can you answer these questions?	146
Objective 3.5: Describe different network topologies.	146
Exam need to know	147
Describe MPLS	147
Describe the point-to-point topology	148
Describe the point-to-multipoint topology	149
Describe the star topology	152
Describe the mesh topology	154
Describe the bus topology	155
Describe the client/server topology	157
Describe the peer-to-peer topology	157
Describe the hybrid topology	158
Can you answer these questions?	158
Objective 3.6: Given a scenario, troubleshoot common physical connectivity problems	159
Exam need to know	159
Troubleshoot cable problems	159
Can you answer these questions?	161
Objective 3.7: Compare and contrast different LAN technologies	161
Exam need to know	161
Compare and contrast LAN types	162
Compare and contrast LAN properties	167
Can you answer these questions?	173
Objective 3.8: Identify components of wiring distribution.	173
Exam need to know	174
Identify an IDF in a wiring distribution	174
Identify an MDF in a wiring distribution	175
Identify a demarc in a wiring distribution	177
Identify a demarc extension in a wiring distribution	177
Identify a smartjack in a wiring distribution	178

Identify a CSU/DSU in a wiring distribution	178
Can you answer these questions?	178
Answers	179
Objective 3.1: Categorize standard media types and associated properties	179
Objective 3.2: Categorize standard connector types based on network media	179
Objective 3.3: Compare and contrast different wireless standards	179
Objective 3.4: Categorize WAN technology types and properties	180
Objective 3.5: Describe different network topologies	180
Objective 3.6: Given a scenario, troubleshoot common physical connectivity problems	181
Objective 3.7: Compare and contrast different LAN technologies	181
Objective 3.8: Identify components of wiring distribution	181
Chapter 4 Network Management	183
Objective 4.1: Explain the purpose and features of various network appliances	183
Exam need to know	183
Explain the purpose and features of a load balancer	184
Explain the purpose and features of a proxy server	185
Explain the purpose and features of a content filter	186
Explain the purpose and features of a VPN concentrator	187
Can you answer these questions?	187
Objective 4.2: Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues	188
Exam need to know	188
Troubleshoot using a cable tester	188
Troubleshoot using a cable certifier	189
Troubleshoot using a crimper	190
Troubleshoot using a butt set	191
Troubleshoot using a toner probe	192
Troubleshoot using a punch down tool	193
Troubleshoot using a protocol analyzer	194

Troubleshoot using a loopback plug	197
Troubleshoot using a TDR	198
Troubleshoot using an OTDR	198
Troubleshoot using a multimeter	199
Troubleshoot using an environmental monitor	200
Can you answer these questions?	200
Objective 4.3: Given a scenario, use appropriate software tools to troubleshoot connectivity issues	200
Exam need to know	201
Troubleshoot using a protocol analyzer	201
Troubleshoot using a throughput tester	202
Troubleshoot using connectivity software	202
Troubleshoot using Ping	202
Troubleshoot using Traceroute	204
Troubleshoot using dig	206
Troubleshoot using Ipconfig/ifconfig	207
Troubleshoot using nslookup	208
Troubleshoot using ARP	209
Troubleshoot using Nbtstat.exe	210
Troubleshoot using netstat	212
Troubleshoot using Route.exe	215
Can you answer these questions?	217
Objective 4.4: Given a scenario, use the appropriate network monitoring resource to analyze traffic.	217
Exam need to know	218
Monitor a network using SNMP	218
Monitor a network using SNMPv2	220
Monitor a network using SNMPv3	220
Monitor a network using syslog	221
Monitor a network using system logs	222
Monitor a network using history logs	223
Monitor a network using general logs	223
Monitor a network using traffic analysis	224
Monitor a network using a network sniffer	224
Can you answer these questions?	225

Objective 4.5: Describe the purpose of configuration management documentation	225
Exam need to know	225
Describe the purpose of wire schemes	226
Describe the purpose of network maps	227
Describe the purpose of documentation	229
Describe the purpose of cable management	230
Describe the purpose of asset management	230
Describe the purpose of baselines	231
Describe the purpose of change management	232
Can you answer these questions?	233
Objective 4.6: Explain different methods and rationales for network performance optimization	233
Exam need to know	233
Explain different methods of network performance optimization	233
Explain different reasons for network performance optimization	237
Can you answer these questions?	237
Answers	238
Objective 4.1: Explain the purpose and features of various network appliances	238
Objective 4.2: Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues	238
Objective 4.3: Given a scenario, use appropriate software tools to troubleshoot connectivity issues	239
Objective 4.4: Given a scenario, use the appropriate network monitoring resource to analyze traffic	239
Objective 4.5: Describe the purpose of configuration management documentation	239
Objective 4.6: Explain different methods and rationales for network performance optimization	239
Chapter 5 Network Security	241
Objective 5.1: Given a scenario, implement appropriate wireless security measures	241
Exam need to know	241
Implement Encryption Protocols	242

Implement MAC address filtering	244
Implement appropriate device placement	245
Implement appropriate signal strength	246
Can you answer these questions?	246
Objective 5.2: Explain the methods of network access security	246
Exam need to know	247
Explain the use of access control lists	247
Explain the use of tunneling and encryption	248
Explain the use of remote access	258
Can you answer these questions?	263
Objective 5.3: Explain methods of user authentication	264
Exam need to know	264
Explain PKI authentication	265
Explain Kerberos authentication	267
Explain AAA authentication	269
Explain network access control	271
Explain CHAP authentication	273
Explain MS-CHAP authentication	273
Explain EAP authentication	274
Explain multifactor authentication	275
Explain two-factor authentication	275
Explain single sign-on authentication	275
Can you answer these questions?	276
Objective 5.4: Explain common threats, vulnerabilities, and mitigation techniques	276
Exam need to know	276
Explain common threats to wireless networks	276
Explain common attacks	278
Explain common mitigation techniques	281
Can you answer these questions?	282
Objective 5.5: Given a scenario, install and configure a basic firewall	283
Exam need to know	283
Install software and hardware firewalls	284
Configure port security	284

Configure stateful inspection vs. packet filtering	285
Configure firewall rules	288
Configure NAT/PAT	292
Configure a DMZ	295
Can you answer these questions?	296
Objective 5.6: Categorize different types of network security appliances and methods	296
Exam need to know	296
Categorize IDSEs and IPSes	296
Categorize vulnerability scanners	298
Categorize honeypots and honeynets	300
Can you answer these questions?	300
Answers	301
Objective 5.1: Given a scenario, implement appropriate wireless security measures	301
Objective 5.2: Explain the methods of network access security	301
Objective 5.3: Explain methods of user authentication	301
Objective 5.4: Explain common threats, vulnerabilities, and mitigation techniques	301
Objective 5.5: Given a scenario, install and configure a basic firewall	301
Objective 5.6: Categorize different types of network security appliances and methods	302
Appendix	303
Index	311

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:

microsoft.com/learning/booksurvey

Introduction

This Rapid Review is designed to assist you with studying for the CompTIA Network+ exam N10-005. The Rapid Review series is designed for exam candidates who already have a good grasp of the exam objectives through a combination of experience, skills, and study, and could use a concise review guide to help them assess their readiness for the exam.

The N10-005 exam is aimed at an IT networking professional who has:

- CompTIA A+ certification or equivalent knowledge
- A minimum of 9 to 12 months of experience in IT networking

Although this experience would preferably include foundation-level networking skills and knowledge, you might have real-world experience. Most candidates who take this exam have the knowledge and skills that are required to implement a defined network architecture with basic network security. Furthermore, a successful candidate will be able to configure, maintain, and troubleshoot network devices using appropriate network tools and understand the features and purpose of network technologies. Candidates will be able to make basic solution recommendations, analyze network traffic, and be familiar with common protocols and media types. It is important to note that you should have some real-world experience with networking prior to taking the N10-005 exam and that having practical knowledge is a key component to achieving a passing mark.

This book will review every concept described in the following exam objective domains:

- Objective 1.0: Network Concepts
- Objective 2.0: Network Installation and Configuration
- Objective 3.0: Network Media and Topologies
- Objective 4.0: Network Management
- Objective 5.0: Network Security

This is a Rapid Review and not a comprehensive guide such as the *CompTIA Network+ Training Kit*. The book covers every exam objective on the N10-005 exam, but will not necessarily cover every exam question. CompTIA regularly adds new questions to the exam, making it impossible for this (or any) book to provide every answer. Instead, this book is designed to supplement your existing independent study and real-world experience with the product.

If you encounter a topic in this book that you do not feel completely comfortable with, you can visit the links described in the text, in addition to researching the topic further using other websites, as well as consulting support forums. If you review a topic and find that you don't understand it, you should consider consulting the *CompTIA Network+ Training Kit* from Microsoft Press. You can also purchase practice exams, or use the one available with the Training Kit, to further determine if you need further study on particular topics.

NOTE The Rapid Review is designed to assess your readiness for the N10-005 exam. It is not designed as a comprehensive exam preparation guide. If you need that level of training for any or all of the exam objectives covered in this book, we suggest the *CompTIA Network+ Training Kit* (ISBN: 9780735662759). The Training Kit provides comprehensive coverage of each N10-005 exam objective, along with exercises, review questions, and practice tests. The Training Kit also includes a discount voucher for the exam.

CompTIA Professional Certification Program

CompTIA professional certifications cover the technical skills and knowledge needed to succeed in a specific IT career. Certification is a vendor-neutral credential. An exam is an internationally recognized validation of skills and knowledge, and is used by organizations and professionals around the globe. CompTIA certification is ISO 17024 Accredited (Personnel Certification Accreditation) and, as such, undergoes regular reviews and updates to the exam objectives. CompTIA exam objectives reflect the subject areas in an edition of an exam, and result from subject-matter expert workshops and industry-wide survey results regarding the skills and knowledge required of a professional with a number of years of experience.

MORE INFO For a full list of CompTIA certifications, go to
<http://certification.comptia.org/getCertified/certifications.aspx>.

Support and feedback

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

Errata

We've made every effort to ensure the accuracy of this book and its companion content. Any errors that have been reported since this book was published are listed on our Microsoft Press site at oreilly.com:

<http://oreilly.com/catalog/errata.csp?isbn=0790145349712>

<http://go.microsoft.com/fwlink/?Linkid=258823>

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at mspinput@microsoft.com.

Please note that product support for Microsoft software is not offered through the addresses above.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://www.microsoft.com/learning/booksurvey>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>.

Preparing for the exam

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

CHAPTER 1

Network Concepts

The Network Concepts domain accounts for approximately 21% of the CompTIA Network+ exam, but more than that, it covers some of the most fundamental principles of computer networking. These are concepts that you encounter repeatedly, both as you prepare for the exam and as you work in the IT field.

To excel at this objective, you must possess a good grasp of certain organizational concepts, such as the OSI reference model; an understanding of basic networking functions, such as IP addressing; and some memorized facts and figures, such as well-known port numbers.

This chapter covers the following objectives:

- Objective 1.1: Compare the layers of the OSI and TCP/IP models
- Objective 1.2: Classify how applications, devices, and protocols relate to the OSI model layers
- Objective 1.3: Explain the purpose and properties of IP addressing
- Objective 1.4: Explain the purpose and properties of routing and switching
- Objective 1.5: Identify common TCP and UDP default ports
- Objective 1.6: Explain the function of common networking protocols
- Objective 1.7: Summarize DNS concepts and components
- Objective 1.8: Given a scenario, implement proper network troubleshooting methodology
- Objective 1.9: Identify virtual network components

Objective 1.1: Compare the layers of the OSI and TCP/IP models

For this exam objective, you must know the structure of the two basic models defining the networking process: the OSI reference model and the TCP/IP model. The OSI model is designed to be independent of any specific networking implementation, and as a result, it does not conform precisely to the networking stacks in general use today. By contrast, the TCP/IP model was designed with specific protocols in mind, and is pertinent only on networks using those protocols. However, virtually all networks today use TCP/IP, so the TCP/IP model is as viable as the OSI model for demonstration, documentation, and discussion purposes.

Exam need to know

- OSI model

For example: Do you know that the OSI reference model consists of seven layers: Layer 1 – Physical; Layer 2 – Data link; Layer 3 – Network; Layer 4 – Transport; Layer 5 – Session; Layer 6 – Presentation; and Layer 7 – Application?

- TCP/IP model

For example: Do you know that the model of the TCP/IP protocol stack consists of four layers: the network interface layer (also known as the link layer); the internet layer; the transport layer; and the application layer?

OSI model

The Open Systems Interconnection (OSI) reference model is a theoretical example of a network protocol stack, which networking educators and administrators use to categorize and define a computer's various networking functions. The top of the model interacts with the applications running on the computer, which might at times require the services of the network. The bottom of the model connects to the network medium over which the system transmits its signals, as shown in Figure 1-1. There are different protocols operating at the various layers of the model, each of which provides functions needed to complete the network communication process.

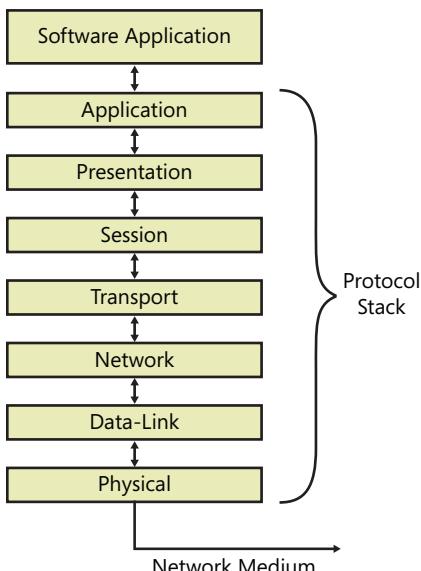


FIGURE 1-1 The seven layers of the OSI reference model.

True or false: The layers of the OSI reference model correspond to the initials of the mnemonic "All People Seem To Need Data Processing."

Answer: *True*. The layers of the OSI model, from top to bottom, are application, presentation, session, transport, network, data-link, and physical.

EXAM TIP While most of the mnemonics that students use to remember the OSI model layers list them from top to bottom, the OSI model layers are traditionally numbered from bottom to top, with the physical layer being Layer 1 and the application layer being Layer 7. One mnemonic for this is “Please Do Not Tell Secret Passwords Anytime.”

MORE INFO The upper layers of the OSI model are seldom referenced by number. The most common use for the layer numbers is in discussions of routing and switching technologies. Switches operate primarily at Layer 2, the data-link layer, and routers at Layer 3, the network layer. However, these devices can have capabilities that span to other layers, resulting in references to technologies such as Layer 3 switching. For more information, see Objectives 1.2 and 1.4.

TCP/IP model

The development of the TCP/IP protocols began years before the documents defining the OSI reference model were published, but the protocols conform to a layered model in much the same way. Instead of the seven layers used by the OSI model, the TCP/IP model—sometimes called the Department of Defense (DoD) model—has four layers. The TCP/IP model layers, in comparison with those of the OSI model, are shown in Figure 1-2.

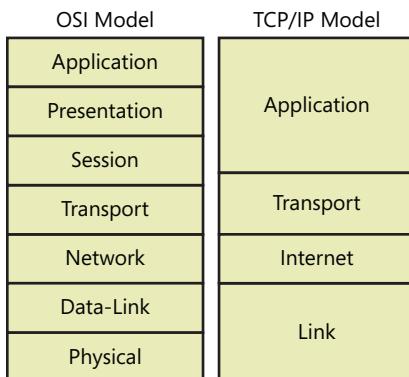


FIGURE 1-2 The four TCP/IP model layers, compared with the seven-layer OSI reference model.

The TCP/IP model layers—even those with the same names—are not exactly analogous to the OSI model layers, nor were the models created with the same intent. The OSI model is intended to be a guide for the creation of networking protocols, whereas the TCP/IP model is a representation of protocols that already exist.

True or false: The link layer of the TCP/IP protocol stack is exactly congruent to the data-link and physical layers of the OSI model.

Answer: *False*. Despite being roughly analogous to the OSI data-link layer, the TCP/IP link layer does not include physical specifications of any kind, nor does it include complex LAN protocols such as Ethernet. Therefore, on many TCP/IP networks, the protocol operating at the link layer might not be part of the TCP/IP suite.

EXAM TIP In the TCP/IP model, the term “internet” is a generic reference to an internetwork and uses a lowercase “i,” as opposed to the public, packet-switching network known as the Internet, with an uppercase “I.” Be careful not to confuse the two.

True or false: The TCP/IP protocol stack was designed to conform to the OSI reference model.

Answer: *False*. Most of the TCP/IP protocols that make up the protocol stack were designed and developed in the 1970s, and therefore predate the OSI reference model. In fact, there is no protocol stack in common use that conforms precisely to the OSI layers. Although originally intended to be a model for an actual networking solution, OSI is now used only as an educational and organizational tool.

EXAM TIP The N10-005 revision of the Network+ exam objectives released in 2011 adds the TCP/IP model and specifically requires students to compare its layers with those of the OSI model. Be careful to distinguish between the two models, and familiarize yourself with the differences between the corresponding layers.

MORE INFO For more information about the structure of the TCP/IP model, see RFC 1122, “Requirements for Internet Hosts – Communication Layers,” available at <http://tools.ietf.org/html/rfc1122>.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is the primary TCP/IP protocol operating at the link layer of the TCP/IP model?
2. Which of the OSI model layers do not have TCP/IP protocols directly associated with them?
3. What are the two protocols operating at the transport layer in both the OSI and TCP/IP models?
4. What organizations were responsible for publishing the original documents defining the OSI reference model and the TCP/IP model?

Objective 1.2: Classify how applications, devices, and protocols relate to the OSI model layers

One of the primary functions of the OSI model is to organize and separate the various elements of the networking process. When defining the function of a network element, such as a protocol, device, or application, it is common to begin

by specifying the OSI model layer at which the element operates. This helps to categorize the function of the element and provides a basic idea of its purpose.

For example, a data-link layer device is generally understood to be involved in local network communications, while the network and transport layers are devoted to end-to-end communications that can span multiple networks. The Network+ exam requires you to understand the functions of many different applications, devices, and protocols, and classifying those elements using the OSI model is the first step to achieving that understanding.

Exam need to know

- Classify how MAC addresses relate to the OSI model layers.

For example: At which layer of the OSI model are network nodes identified using MAC addresses?

- Classify how IP addresses relate to the OSI model layers.

For example: At which layer of the OSI model are network nodes identified using IP addresses?

- Classify how EUI-64 relates to the OSI model layers.

For example: In what layer of the OSI reference model do you find an EUI-64?

- Classify how frames relate to the OSI model layers.

For example: Which layer of the OSI model uses the term "frame" to refer to the protocol data unit generated by a networking protocol?

- Classify how packets relate to the OSI model layers.

For example: At which layer of the OSI reference model are data structures called packets?

- Classify how switches relate to the OSI model layers.

For example: At which layer of the OSI model do switches perform their basic functions?

- Classify how routers relate to the OSI model layers.

For example: At which layer of the OSI model do routers perform their basic functions?

- Classify how multilayer switches relate to the OSI model layers.

For example: At which layers of the OSI model do multilayer switches perform their functions?

- Classify how hubs relate to the OSI model layers.

For example: At which layer of the OSI model do hubs perform their basic functions?

- Classify how encryption devices relate to the OSI model layers.

For example: Which layer of the OSI model is responsible for encrypting data?

- Classify how cables relate to the OSI model layers.

For example: Which layer of the OSI model defines the properties of network cables?

- Classify how NICs relate to the OSI model layers.
For example: At which layer of the OSI model do NICs operate?
- Classify how bridges relate to the OSI model layers.
For example: At which layer of the OSI model do bridges perform their basic functions?

MAC addresses

A media access control (MAC) address, also known as a hardware address, is a six-byte hexadecimal value that local area networking (LAN) protocols at the data-link layer use to identify systems on a shared network medium. Manufacturers of network interface adapters permanently assign a unique MAC address to each adapter, so that there can be no address duplication on any network.

True or false: Ethernet devices use MAC addresses to identify the source and the destination of each frame they transmit over the network.

Answer: *True*. The Ethernet frame format contains a six-byte Source Address field and a six-byte Destination Address field, which together function like the addresses on a postal envelope.

MORE INFO Data-link layer protocols are concerned only with LAN communications, so the values in the Destination Address and Source Address fields must identify systems on the local network. If a computer on the LAN is transmitting to another computer on the same LAN, then its packets contain the address of that target computer in their Destination Address fields. If a computer is transmitting to another computer on a different network, then the value in the Destination Address field must be the address of a router on the LAN.

True or false: Internet transmissions use a MAC address to identify the final recipient of a message.

Answer: *False*. Internet transmissions use an Internet Protocol (IP) address to identify the final recipient of a message, not a MAC address. This is because MAC addresses are local. A packet might pass through multiple LANs on the way to a destination on the Internet, and have different destination MAC addresses as it does so.

EXAM TIP For the Network+ exam, you must be able to distinguish MAC addresses from IP addresses. A MAC address is a six-byte hexadecimal value, such as 60-EB-69-93-5E-E4. An IP version 4 address is 32 bits, generally expressed in four octets of dotted decimal notation, as in 192.168.76.3. An IP version 6 address is 128 bits long, and generally expressed in 8 groups of 16-bit hexadecimal values separated by colons, such as fe80::7441:4473:f204:ec1d.

IP addresses

The Internet Protocol (IP) is the primary end-to-end protocol in the TCP/IP networking stack. Operating at the network layer of the OSI model (and the internet layer of the TCP/IP model), IP has its own addressing system, which it uses to identify systems on the network.

As with Ethernet, IP has header fields that contain the IP addresses of the source and destination systems involved in a network transaction. The difference between the two is that IP uses IP addresses instead of MAC addresses, and the Destination IP Address field identifies the final recipient of the transmission.

True or false: Ethernet systems on a TCP/IP network have a protocol that converts network layer IP addresses to data-link layer MAC addresses.

Answer: *True*. Address Resolution Protocol (ARP) converts IP addresses into MAC addresses by broadcasting request packets containing the IP address on the local network and waiting for the holder of that IP address to respond with a reply containing the equivalent MAC address.

MORE INFO IP is currently in transition from version 4 to version 6, and the two versions have different address formats. For more information, see “Objective 1.3: Explain the purpose and properties of IP addressing.”

True or false: Packets on a TCP/IP network can have two destination addresses pointing to different systems.

Answer: *True*. The IP header at the network layer has a Destination IP Address field that always specifies the ultimate destination of the packet. At the same time, the Ethernet header at the data-link layer will have a changing Destination Address field that points to the next intermediary system on the local network, until it finally reaches the ultimate destination network, at which point both addresses will point to the same ultimate destination.

EUI-64

The extended unique identifier-64 (EUI-64) is a 64-bit value that some TCP/IP systems use to form the second half of a 128-bit IPv6 address. The IPv6 address is a network layer structure, but the EUI-64 value for a system is derived from its data-link layer MAC address.

True or false: No two computers can legitimately have the same EUI-64 value in their IPv6 addresses.

Answer: *True*. The EUI-64 value that makes up the second half of some IPv6 addresses is taken from the system’s MAC address, which, by definition, is unique. Therefore, no two EUI-64 values on different systems can be identical, unless an individual is making a deliberate attempt to spoof the IPv6 address.

MORE INFO To derive the 64-bit EUI-64 value for an interface, an IPv6 implementation splits the 48-bit MAC address into two 24-bit halves, and takes the two 24-bit values and adds a 16-bit value between them: *11111111 11111110* in binary or *ff fe* in hexadecimal. Finally, it changes the seventh bit in the OUI—called the universal/local bit—from a 0 to a 1. This changes the hexadecimal value of the first byte in the address from 00 to 02. Therefore, a computer with a network adapter that has a MAC address of *00-1a-6b-3c-ba-1f* would have an IPv6 global unicast address with the following EUI-64 value: *021a:6bff:fe3c:ba1f*.

True or false: All IPv6 addresses include the system's EUI-64 value.

Answer: *False*. Some IPv6 implementations avoid using the EUI-64 value, for fear that it might be possible to track the physical location of a computer based on its IPv6 address.

Frames

The data structures created by the protocols at the various layers of the OSI reference model have different names. At the data-link layer, the structure that a protocol creates when it encapsulates a network layer message is called a frame. The term frame is not used at any other layer.

Unlike the protocols at the upper layers, a data-link layer frame consists of both a header and a footer, which the protocol adds to the data it receives from the network layer. The resulting frame is the final element added to the data packet, which is then ready for transmission over the network.

True or false: A data-link layer frame includes an error detection mechanism.

Answer: *True*. The frame check sequence (FCS) field in the data-link layer footer contains a checksum calculated by the source computer. Once the frame reaches its destination, the receiving computer performs the same calculation and compares the results to the FCS value. If the two fail to match, then the frame has been corrupted or modified in transit.

EXAM TIP Candidates for the Network+ exam should be familiar with the terminology used for the data units created by the various protocols in the TCP/IP stack. Data-link layer protocols create “frames.” IP, at the network layer, creates “packets.” Connectionless protocols such as UDP, at the transport layer, create units called “datagrams.” Because IP is a connectionless protocol, the term datagram can apply to those data units as well. The TCP protocol at the transport layer creates individual messages called “segments,” which are part of a “sequence.”

True or false: All data-link layer frames include source and destination MAC addresses.

Answer: *False*. Ethernet frames always include source and destination MAC addresses, but there are data-link layer protocols other than Ethernet that do not. The Point-to-Point Protocol (PPP) is designed for use on wide area network (WAN)

connections between systems. Because there are only two systems involved in a WAN connection, there is no need to include addresses in every frame.

EXAM TIP There are several different variants of the Ethernet frame format, the selection of which depends on the version of the Ethernet standard the system is configured to use. The formats are functionally the same, but for systems on the network to communicate, they must all be using the same frame format.

Packets

Although it is often mistakenly used to refer to the entire data structure transmitted over the network, the term packet actually refers to the unit of data carried inside a data-link layer frame. A packet is therefore a network layer structure.

On a packet-switching internetwork, such as the Internet, packets might travel through dozens of networks, with the router for each network stripping off the previous frame and applying its own frame to the data. The packet inside these many different frames remains intact, however.

True or false: Every TCP/IP packet contains a frame.

Answer: *False*. The packet is the network-layer data carried within the data-link layer frame. Therefore, every frame contains a packet.

EXAM TIP The Network+ exam might also refer to the network layer data unit as a datagram. Technically, a datagram is the data unit created by a connectionless protocol. This is why both IP and UDP generate datagrams. However, because there is no connection-oriented protocol at the network layer, the terms datagram and packet are synonymous in TCP/IP networking.

True or false: Every TCP/IP packet must contain a transport layer datagram or segment.

Answer: *False*. Packets carrying transport layer data must contain a UDP datagram or a TCP segment, but there are also packets that carry Internet Control Message Protocol (ICMP) data directly within the IP datagram, which do not use UDP or TCP.

Switches

A switch is a data-link layer device that connects computers and other systems together into a LAN. Basic switches consist of a box or a rack-mounted module with one or more rows of female cable connectors. Plugging devices into the connectors enables them to communicate with each other by transmitting packets.

Unlike hubs, switches have intelligence that enables them to determine the address of the device connected to each port. When a unicast packet arrives through any of the switch's ports, the switch reads its destination addresses and forwards the packet out through the port providing access to the destination system.

True or false: Switches have almost completely replaced hubs on today's local area networks.

Answer: *True*. Switches conserve network bandwidth by delivering packets only to their intended recipients. On a hub-based network, every computer must receive and process every packet received by the hub.

MORE INFO In addition to functioning at the data-link layer, switches can also have network layer capabilities as well. For more information, see “Objective 1.4: Explain the purpose and properties of routing and switching,” and “Objective 2.1: Given a scenario, install and configure routers and switches.”

True or false: All switched networks use a bus topology.

Answer: *False*. A switch functions as the cabling nexus for a LAN. Each computer has its own cable connecting it to the switch. Switched networks can therefore be said to use a star topology.

EXAM TIP The Network+ exam has, at times, referred to the relatively simple switching devices used in home and small-to-medium office networks as “basic switches.” These are strictly data-link layer devices that do not have advanced features, such as VLANs.

Routers

A router is a network layer component that connects two networks together, selectively forwarding only the traffic that is destined for the other network. Because most large networks today are switched internally, the primary function of routers is to connect LANs to WAN connections.

Routers also have tables containing information about other networks, which enable them to direct incoming packets to their ultimate destinations.

True or false: Splitting a network with a router reduces the amount of broadcast traffic on the network.

Answer: *True*. Unlike switches, hubs, and bridges, routers do not forward broadcast traffic.

MORE INFO For more information on routing, see “Objective 1.4: Explain the purpose and properties of routing and switching,” and “Objective 2.1: Given a scenario, install and configure routers and switches.”

True or false: A router must have at least two network interfaces.

Answer: *False*. By the traditional definition, a router must be connected to two or more networks, so it must have at least two network interfaces. These interfaces can be standard LAN adapters, or any type of WAN equipment. However, with the advent of virtual LANs, there are now routers available with a single interface. Called stub routers or one-armed routers, these devices connect to a switch and route traffic between VLANs.

Multilayer switches

A multilayer switch is an advanced networking device that, in addition to functioning as a standard data-link layer switch, also supports functions associated with other OSI model layers, most particularly network layer routing.

True or false: In addition to the data-link layer, switches can also operate at the network layer.

Answer: *True*. Advanced switches have the ability to create virtual LANs (VLANs), which are subnets that exist only in the switch. To enable VLANs to communicate with each other, these switches also support virtual routing, which is a network layer process.

MORE INFO For more information on VLANs and advanced switching techniques, see "Objective 2.1: Given a scenario, install and configure routers and switches."

Hubs

A hub is a cabling nexus for a LAN using a star topology. Unlike a switch, which is often similar in appearance, a hub is a purely physical layer device. The hub amplifies the signals entering through any of its ports and forwards them out through all of the other ports, creating a shared network medium.

True or false: Hubs can read the destination addresses from the frames arriving through its ports.

Answer: *False*. Hubs lack any ability to interpret incoming signals. They are electrical devices that manipulate signals at the physical level, but they cannot interpret them.

EXAM TIP Having been largely replaced by switches, hubs are all but obsolete in the networking world today, and are less likely to appear on the Network+ exam than they have on previous iterations of the test.

True or false: Replacing a hub with a switch increases the efficiency of a LAN.

Answer: *True*. While a hub forwards incoming signals out through all of its ports, switches only forward signals out through the destination port. This conserves bandwidth and provides each pair of computers with what amounts to a dedicated link.

NOTE A repeater is a device that extends the maximum length of a network cable by amplifying the signals passing over it. Because hubs do essentially the same thing for all of their connected devices, they are sometimes referred to as multiport repeaters.

Encryption devices

The term encryption device refers to any mechanism that employs an algorithm to cryptographically encode data. Encryption devices can be as large as a server or as small as a USB flash drive. Whatever the form of the device, however, the encryption process is carried out at the presentation layer of the OSI model.

True or false: On TCP/IP systems, encryption algorithms are standalone protocols that run at the presentation layer of the OSI model.

Answer: *False*. There are no standalone presentation layer protocols in the TCP/IP suite. Presentation layer functions, including encryption, are typically incorporated into application layer protocols.

EXAM TIP Unlike the other hardware components mentioned in this objective, there are no dedicated networking components called encryption devices. Encryption is a function that is incorporated into other hardware and software components. Therefore, while the Network+ exam might refer to encryption devices, this is solely for the purpose of testing your knowledge that encryption is a presentation layer process.

Cables

Cables are the physical layer components that form the network medium on most LANs. Depending on the topology, distance, and environmental requirements for the network, LANs use one of the following three basic cable types: coaxial, twisted pair, or fiber optic.

True or false: Coaxial cables are no longer used to build new Ethernet LANs.

Answer: *True*. Coaxial Ethernet networks require a bus topology, and for various reasons, including cost and ease of installation, this type of cable is no longer used.

NICs

The network interface adapter, also known as a network interface card or NIC, is the hardware implementation of the data-link layer protocol. Virtually all of the NICs sold today are Ethernet, with models available that support various speeds, expansion buses, and cable types.

True or false: Most of the desktop computers manufactured today have an Ethernet network interface adapter integrated into the motherboard.

Answer: *True*. Ethernet network interface adapters are all but ubiquitous on the motherboards manufactured for desktop computers.

EXAM TIP The Network+ exam might persist in using the term NIC (pronounced as “nick”), even when referring to a network interface adapter that is not actually an expansion card.

True or false: Every NIC has a unique MAC address permanently assigned by the manufacturer.

Answer: *True*. It is the network interface adapter that has the MAC address assigned to it by the hardware manufacturer, whether the adapter is a separate card or integrated into the motherboard.

Bridges

A bridge is a data-link layer device that splits a LAN in half and selectively forwards traffic based on its destination address. When a packet arrives through one of the bridge's interfaces, the bridge reads the destination hardware address from the Ethernet header. If the packet is destined for a computer on the other side of the bridge, it forwards the packet out through its other interface. If the packet is destined for a computer on the same side of the bridge from which it was received, the bridge simply discards the packet.

True or false: Installing a bridge on a LAN splits the network into two separate broadcast domains.

Answer: *False*. Bridges forward all broadcasts to the other side of the network. The address-based filtering they perform is limited to unicast transmissions.

EXAM TIP The Network+ exam objectives still mention bridges, even though the devices are rarely used on today's networks.

NOTE Bridges possess a degree of intelligence similar to that of switches. A basic switch is, in essence, nothing more than a multiport bridge.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Only one of the items listed in this objective is associated with an OSI model layer other than the physical, data-link, or network. Which is it?
2. You can build a simple Ethernet LAN with nothing more than a NIC for each computer, a switch, and some cables. Which of those components are associated exclusively with the physical layer of the OSI model?
3. A multilayer switch functions primarily at which two layers of the OSI reference model?
4. Which of the devices listed in this objective can split a network into two broadcast domains?
5. Which layer of the OSI model uses addresses that can be 32 or 128 bits long?

Objective 1.3: Explain the purpose and properties of IP addressing

IP addressing is one of the fundamental functions of the TCP/IP protocol suite and the network layer IP. Every device on an internetwork must have a unique IP address, so that IP can address packets specifically to it. IP addresses specify both the network on which the device is located and the device itself, called a host, on that particular network. Routers use the network identifier to forward packets to the correct network, and the router on the destination network uses the host identifier to forward the packets to the correct device.

Exam need to know

- Explain the intended purpose and properties of now-obsolete IP address classes.
For example: Which IP address class provided the largest number of hosts per subnet?
- Explain the purpose and properties of Classless Inter-Domain Routing (CIDR).
For example: How many bits are allocated for the host identifier in the 10.0.54.0/24 network address?
- Explain the purpose and properties of IPv4 and IPv6 formatting.
For example: What is the largest possible value for each of the four decimal numbers in an IPv4 address?
- Explain the purpose and properties of the MAC address format.
For example: What is the term used for the first three bytes of a MAC address?
- Explain the purpose and properties of subnetting.
For example: How many hosts can you create on a subnet with the mask 255.255.255.240?
- Explain the purpose and properties of multicasts, unicasts, and broadcasts.
For example: What is the standard MAC address value used for a broadcast transmission?
- Explain the purpose and properties of APIPA.
For example: What is the IPv4 network used by default for Automatic Private IP Addressing?

IP address classes

IPv4 addresses contain both a network identifier and a host identifier, which means that some of the 32 bits in the address specify the network on which the host is located and the rest of the bits identify the specific host on that network. However, the division between the network identifier bits and the host identifier bits is not always in the same place. The original IP standard defined three primary classes of

IP addresses: A, B, and C, which provided support for networks of different sizes, as shown in Figure 1-3.

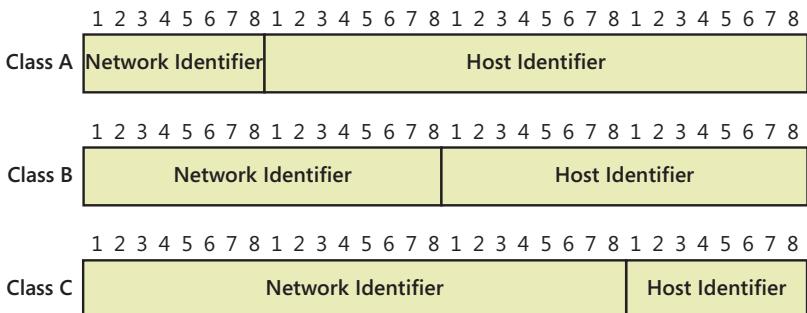


FIGURE 1-3 The three primary classes of IPv4 addresses.

The characteristics of these three address classes are listed in Table 1-1.

TABLE 1-1 IPv4 address classes.

IP ADDRESS CLASS	CLASS A	CLASS B	CLASS C
First bit values (binary)	0	10	110
First byte value (decimal)	0 -127	128-191	192-223
Number of network identifier bits	8	16	24
Number of host identifier bits	24	16	8
Number of possible networks	126	16,384	2,097,152
Number of possible hosts	16,777,214	65,534	254

The “First bit values” row in the table specifies the values that the first one, two, or three bits of an address in each class must have. Early TCP/IP implementations used these bit values to determine the class of an address.

For web servers and other computers to be accessible by clients on the Internet, they must have public IP addresses, that is, addresses registered with an authority, such as an Internet service provider (ISP). For workstations and other computers that do not require an Internet presence, administrators typically use private IP addresses, which are freely available for use on any network and are not registered as belonging to any particular organization.

The private address ranges for each class are as follows:

- **Class A** 10.0.0.0 through 10.255.255.255
- **Class B** 172.16.0.0 through 172.31.255.255
- **Class C** 192.168.0.0 through 192.168.255.255

True or false: You cannot assign all of the possible values in a given address class to network devices.

Answer: *True*. The host identifier values in each address class consisting of all zeroes and all ones are reserved; you cannot assign them to hosts. The all zeroes address identifies the network itself and the all ones address is the broadcast address for the network.

NOTE In addition to classes A, B, and C, the IP standard defines two additional address classes: Class D, which is used for multicast addresses; and Class E, which is experimental. Class D addresses begin with the bit values 1110, and Class E addresses begin with the values 11110.

EXAM TIP The Network+ objectives refer to public and private addresses, but they are also sometimes known as registered and unregistered addresses. Candidates should be familiar with both sets of terms.

True or false: A web server must have a public IP address to be accessible by clients on the Internet.

Answer: *True*. Public, or registered, IP addresses are assigned to particular organization and reserved for use by one host on the Internet.

EXAM TIP Although classful addressing is no longer used on the Internet, CompTIA continues to include it in the Network+ objectives and on the exam, for historical context.

Classless inter-domain routing (CIDR)

There are many networks that have more than the 254 hosts provided by a Class C address, and there are none that have the 16 million provided by a Class A. The classful IP addressing system, therefore, proved to be wasteful as the IP address space grew crowded. CIDR is a subnetting method that enables administrators to place the division between the network bits and the host bits anywhere in the address, not just between octets. This makes it possible to create networks of almost any size.

CIDR also introduced a new notation for network addresses. A standard IPv4 network address is followed by a forward slash and a numeral specifying the size of the network identifier. For example, 192.168.43.0/24 represents an address that uses a 24-bit network identifier, leaving the other 8 bits for up to 254 host identifiers, which would formerly be known as a Class C address.

True or false: Classless IP addresses use the first few binary bits of the network identifier to specify the size of the network.

Answer: *False*. In a classless address, the size of the network is indicated by the suffix, or by the use of a subnet mask.

EXAM TIP The CIDR notation is the most commonly used form of network address used on the Network+ exam. Candidates must know that the number following the slash specifies the size of the network identifier, and that the remaining bits identify the host.

True or false: In the classless address 192.168.76.0/24, the number 24 specifies how many hosts you can create on the network.

Answer: *False*. The number 24 indicates the number of bits in the network identifier. There are therefore 8 host bits, allowing a maximum of 254 hosts on the network.

IPv4 and IPv6 address formatting

The original IP protocol standard calls for 32-bit IP addresses, but the depletion of the IPv4 address space led to the development of IPv6, which uses 128-bit addresses. The IP addresses used in networks around the world are currently in the midst of a lengthy conversion from IPv4 to IPv6.

An IPv4 address is a 32-bit value that contains both a network identifier and a host identifier. The address is notated by using four decimal numbers ranging from 0 to 255, separated by periods, as in 192.168.1.44. This is known as dotted decimal notation.

IPv6 addresses use a notation called colon-hexadecimal format, which consists of eight 16-bit hexadecimal numbers, separated by colons, as in the following example:

21cd:0053:0000:0000:e8bb:04f2:003c:c394

When an IPv6 address has two or more consecutive 8-bit blocks of 0s, you can replace them with a double colon. You can also remove the leading 0s in any block where they appear, as follows:

21cd:53::e8bb:4f2:3c:c394

True or false: The hexadecimal value 21cd:53::e8bb::3c:c394 is a valid IPv6 address.

Answer: *False*. A valid IPv6 address can only have one double colon in it.

EXAM TIP Network+ candidates must be familiar with the rules for contracting IPv6 addresses.

MORE INFO For more information on the formation of IPv6 addresses, see the “EIU-64” section in “Objective 1.2: Classify how applications, devices, and protocols relate to the OSI model layers.”

MAC address formatting

The first three bytes of a MAC address, called the organizationally unique identifier (OUI), consist of a value assigned to the hardware manufacturer by the Institute of Electrical and Electronics Engineers (IEEE). The second three bytes consist of a unique value assigned by the manufacturer to each individual device.

True or false: Two computers can have the same OUI.

Answer: *True*. The OUI is a value assigned to a manufacturer of network interface adapters, and all of the adapters produced by that manufacturer will have MAC addresses with identical OUIs. Only the second three bytes of the MAC address on every adapter must be unique.

EXAM TIP Network+ candidates must be able to differentiate MAC addresses and IPv6 addresses, both of which use hexadecimal (base sixteen) notation.

True or false: The Ipconfig.exe program on a Windows computer displays the MAC address assigned to the network interface adapter.

Answer: *True*. In addition to TCP/IP configuration settings, Ipconfig.exe identifies each of the network interface adapters in the computer and displays their MAC addresses, as in the third line of the following display.

```
Connection-specific DNS Suffix . : zacker.local
Description . . . . . : Realtek PCIe GBE Family Controller
Physical Address . . . . . : 60-EB-69-93-5E-E5
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . . : Yes
Link-local IPv6 Address . . . . . : fe80::7441:4473:f204:ec1d%10(PREFERRED)
IPv4 Address. . . . . : 192.168.2.9(PREFERRED)
Subnet Mask . . . . . : 255.255.255.0
Lease Obtained. . . . . : Sunday, April 15, 2012 1:11:50 PM
Lease Expires . . . . . : Friday, April 27, 2012 1:11:48 PM
Default Gateway . . . . . : 192.168.2.9
DHCP Server . . . . . : 192.168.2.1
DHCPv6 IAID . . . . . : 241232745
DHCPv6 Client DUID. . . . . : 00-01-00-01-14-81-CC-39-60-EB-69-93-5E-E5
DNS Servers . . . . . : 192.168.2.1
Primary WINS Server . . . . . : 192.168.2.1
NetBIOS over Tcpip. . . . . : Enabled
```

MORE INFO For more information on the formation of IPv6 addresses, see the “MAC Address” section in “Objective 1.2: Classify how applications, devices, and protocols relate to the OSI model layers.”

IP address subnetting

When the wastefulness of classful IP addressing was recognized, the designers of the IP protocol developed a system for subdividing network addresses by creating subnets within them. A subnet is simply a subdivision of a network address that administrators can use to represent a part of a larger network, such as one LAN on an internetwork or the client of an ISP. Thus, a large ISP might have a Class A address registered to it, and it might allocate sections of that network address to its clients in the form of subnets.

To understand the process of creating subnets, you must understand the function of the subnet mask. TCP/IP systems at one time recognized the class

of an address simply by examining the values of its first three bits. Today, when you configure the TCP/IP client on a computer, you assign it an IPv4 address and a subnet mask. The subnet mask is a 32-bit value that specifies which bits of the IP address are the network identifier and which bits are the host identifier. For example, the subnet mask 255.255.255.0, in binary form, is 24 ones and eight zeroes. The ones are the network identifier bits and the zeroes are the host identifier bits.

True or false: To create eight-bit subnets on a Class A network address, you would use a subnet mask value of 255.255.0.0.

Answer: *True*. The subnet mask for a Class A address is 255.0.0.0. Borrowing eight bits from the host identifier to create subnets gives you a 16-bit network identifier. The subnet mask, therefore, consists of 16 ones and 16 zeroes, in binary form, or 255.255.0.0 in decimal form.

EXAM TIP Candidates for the Network+ exam should be capable of calculating a subnet mask by converting a 32-bit string of ones and zeroes into a dotted decimal value.

True or false: The boundary between the network identifier and the host identifier in a subnetted IPv4 address must fall between bytes.

Answer: *False*. Subnets can be any size, so the boundary between the network and host identifiers can theoretically fall between any two bits.

Multicasts, unicasts, and broadcasts

IPv4 supports three basic types of addresses, as follows:

- **Unicast** A one-to-one transmission sent to an IP address with a specific host identifier, anywhere on the internetwork.
- **Broadcast** A one-to-many transmission sent to an IP address with a host identifier that consists of all 1s. Broadcast transmissions are received and processed by all of the hosts on the local network.
- **Multicast** A one-to-many transmission sent to a specially-allocated multicast IP address. Multicast addresses are targeted at specific groups of hosts, which can be scattered around the internetwork.

True or false: Registration of hosts in multicast groups is handled by the Internet Control Message Protocol (ICMP).

Answer: *False*. The protocol that hosts use to register themselves in multicast groups is called the Internet Group Management Protocol (IGMP).

EXAM TIP Network+ exam candidates should be able to recognize a broadcast address, and be familiar with the term “broadcast domain,” which refers to the group of network devices that will receive a broadcast transmission generated by a particular computer. The boundaries between broadcast domains are typically set by routers. Switches, bridges, and hubs all forward broadcasts.

True or false: Both MAC addresses and IPv4 addresses support broadcast transmissions, but IPv6 addresses do not.

Answer: *True*. MAC addresses and IPv4 addresses consisting of all ones (ffffffffffff or 255.255.255.255, respectively) cause a transmission to be sent to all of the local network devices. IPv6, however, has no broadcast addresses; it uses multicasts and a new type of transmission called an anycast, instead.

Automatic private IP addressing

Automatic Private IP Addressing (APIPA) is a DHCP failover mechanism used by all of the current Windows operating systems. When a device fails to locate a DHCP server on the network, APIPA takes over and automatically assigns an address on the 169.254.0.0/16 network to the computer. The system then uses the Address Resolution Protocol (ARP) to ensure that no other computer on the local network is using the same address.

For a small network that consists of only a single, unrouted LAN, APIPA is a simple and effective alternative to installing a DHCP server, as it creates and assigns addresses that are all on the same subnet.

True or false: Two computers on the same network that assign themselves addresses using APIPA cannot communicate with each other.

Answer: *False*. APIPA assigns addresses that are all on the same IP subnet, and the systems use ARP to confirm that their addresses are unique. Therefore, two systems on the same network with APIPA addresses can communicate with each other.

EXAM TIP Network+ exam candidates should be able to recognize an IPv4 address assigned by APIPA.

True or false: APIPA is capable of assigning both IPv4 and IPv6 addresses.

Answer: *False*. APIPA can only assign IPv4 addresses. IPv6 has its own mechanism for self-assigning addresses, called stateless address autoconfiguration.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which IP address class provides the largest number of hosts per subnet?
2. What subnet mask must you use for a network with the address 172.16.132.0/21?
3. What subnet does APIPA use when assigning IP addresses?
4. The link local unicast addresses generated by the IPv6 stateless address autoconfiguration process use the network address fe80:0000:0000:0000/64. What is the most compact allowable form of this address?
5. What is the OUI in the following MAC address: 60-EB-69-86-3A-C7?

Objective 1.4: Explain the purpose and properties of routing and switching

Routers and switches are the two basic connectivity devices used to join individual LANs into internetworks. Routing is the process of forwarding data packets from one network to another, until they reach their final destinations. A switch is a multiport bridging device in which each port forms a separate network segment. Similar in appearance to a hub, a switch receives incoming traffic through any of its ports and forwards the traffic out to the single port needed to reach the destination.

Both routing and switching are complex processes that require the additional functionality of many other specialized TCP/IP processes and protocols. This objective covers a good many of these processes and protocols, knowledge of which is essential for the Network+ exam.

Exam need to know

- Explain the purpose and properties of routing tables.
For example: Why does every TCP/IP system need a routing table?
- Explain the differences between static and dynamic routing.
For example: What tools do you use for static routing?
- Explain the function of routing metrics.
For example: Where do routing metric values come from?
- Explain the meaning of next hop routing.
For example: What is a hop and what is its significance to the routing process?
- Explain the differences between link state, distance vector, and hybrid routing protocols.
For example: How does a link state protocol measure route efficiency?
- Explain the purpose and properties of RIP.
For example: What is the difference between RIPv1 and RIPv2?
- Explain the purpose and properties of EIGRP.
For example: How does EIGRP evaluate the efficiency of routes?
- Explain the purpose and properties of OSPF.
For example: How does OSPF offer an improvement over RIP?
- Explain the meaning of convergence.
For example: Why is a network's convergence state significant?
- Explain the purpose of the Spanning-Tree Protocol.
For example: What switching problem does the Spanning Tree Protocol address?
- Explain the purpose and properties of 802.1q VLANs.
For example: Why are VLANs needed on switched networks?

- Explain the purpose of port mirroring.
For example: Why do administrators need mirrored ports?
- Explain the differences between broadcast domains and collision domains.
For example: What effect do switches have on collision domains?
- Explain the differences between IGP and EGP.
For example: What exterior gateway protocol is in common use today?

Routing tables

Every host on a TCP/IP network has a routing table that holds the information the system uses to send packets to their proper destinations. On a LAN, routing is essentially the process of determining what data-link layer protocol address the system should use to reach a particular IP address. In the case of an Ethernet LAN, IP must determine what MAC address the system should use in its Ethernet frames.

If a computer wants to transmit a packet to a destination on the local network, for example, the routing table instructs it to address the packet directly to that system. This is called a direct route. If a packet's destination is on another network, the routing table supplies the address of the router that the system should use to reach that destination.

Remember that data-link layer protocols such as Ethernet can only send frames to the local network. Because the final destination of the packet is on a distant network, the Ethernet destination on the local network must be a router. This is called an indirect route.

True or false: On a TCP/IP network, every router and computer has its own routing table.

Answer: *True*. Every host on a TCP/IP network must have a routing table to determine where to send its packets. This includes routers and computers.

EXAM TIP The Network+ exam typically associates the routing process and routing tables with IP, which runs at the network layer of the OSI model. Dynamic routing protocols, however, which are responsible for populating the routing table, operate at the application layer.

True or false: The default gateway is usually the first entry in a computer's routing table.

Answer: *True*. The default gateway is the router that a system uses for all packets with destinations not listed in the routing table.

MORE INFO For more information on working with routing tables, see "Objective 2.1: Given a Scenario, install and configure routers and switches."

Static vs. dynamic routing

There are two techniques for updating a routing table: static routing and dynamic routing. In static routing, a network administrator manually creates routing table entries, using a program designed for this purpose. In dynamic routing, routers use specialized protocols to create routing table entries automatically.

True or false: Static routing is suitable only for relatively small networks.

Answer: *True*. Static routing requires administrators to type the information for each route, often using a command line program with a cryptic syntax. Therefore, it is a time-consuming process that is prone to errors.

EXAM TIP Network+ exam candidates should be familiar with the software tools used for static routing and the protocols used for dynamic routing.

Routing metrics

Each entry in a routing table contains a metric, which is a value that specifies the efficiency of the route. Metric values are relative; a lower value indicates a more efficient route than a higher value. When a routing table contains multiple routes to the same destination, the system always uses the table entry with the lower metric value.

The term hop count refers to the distance between two networks, based on the number of routers that packets must pass through on the way from the source to the destination. Distance vector routing protocols use hop counts to create metric values in routing table entries. A route with fewer hops is considered to be more efficient than one with more hops.

The size of IP packets depends on the data-link layer protocol the network is using. The transmitting system uses the maximum transmission unit (MTU) of the connected network to determine how large each datagram should be. The MTU is the largest possible frame supported by the data-link layer protocol. Using the largest frame conserves bandwidth by eliminating the overhead involved in transmitting multiple packets instead of one. If, during the journey from source to destination, a packet encounters a network with a smaller MTU, the router for that network fragments the packet into smaller pieces and transmits each one individually.

One of the criteria that link state protocols use to evaluate routes is the route cost. The route cost is a metric assigned by the network administrator used to rate the relative usability of a route. The cost can refer to the literal financial expense incurred by the link, or any other pertinent factor. By using criteria such as this, link state protocols reflect the latency of network routes more precisely. Latency is the time required for data to travel from one location to another.

True or false: The metric values in a routing table must be 15 or less.

Answer: *False*. The Routing Information Protocol (RIP) uses metric values that can be no larger than 15, but that is a limitation of the protocol, not of the routing table.

True or false: On a network that uses static routing, administrators can use any values they wish for the routing table metrics.

Answer: *True*. In static routing, the metric values are relative, and have no statistical meaning. All that matters when there are two routes to the same network is which has the lower metric value.

True or false: IPv4 and IPv6 routers both fragment packets when necessary.

Answer: *False*. In IPv6, intermediate routers do not fragment packets. Instead, end systems use Path MTU Discovery to determine the MTU for an entire route from source to destination.

EXAM TIP Network+ exam candidates should understand the concept of the path MTU and Path MTU Discovery, and how they affect the fragmentation process in IP.

Next hop

The term next hop refers to the next router on a packet's path through an internetwork to its destination. Routing table entries specify only the next hop that a packet should take, not the entire route. RIPv2 routes have a Next Hop field that contains the address of the next router, which in a Windows routing table goes in the Gateway field.

True or false: In distance vector routing, a hop between two LANs in the same building carries the same weight as a transoceanic hop between networks on different continents.

Answer: *True*. The fundamental flaw of distance vector routing is its reliance on hop counts that do not consider the distance or relative speed of the links between routers.

EXAM TIP Network+ exam candidates should associate hop counts with both distance vector routing and the Routing Information Protocol (RIP).

Link state vs. distance vector routing

A routing protocol that uses metrics based on the number of hops to the destination is called a distance vector protocol. The metric value included with each route determines the efficiency of the route, based on the number of hops required to reach the destination. In a distance vector routing protocol, every router on the network advertises its routing table to its neighboring routers. Each router then examines the information supplied by the other routers, chooses the best route to each destination network, and adds it to its own routing table.

Distance vector routing has a fundamental flaw: it bases its routing metrics solely on the number of hops between two networks, which is not always efficient. When an internetwork consists of multiple LANs in the same location, all connected using the same data-link layer protocol, the hop count is a valid indicator. However, when WAN links are involved, a single hop can refer to anything from a high-speed leased

line to a dial-up modem connection. It is therefore possible for traffic moving over a route with fewer hops to take longer than one with more hops.

The alternative to distance vector routing is called link state routing. A link state routing protocol works by flooding the network with messages called link state advertisements. Each router receiving such a message propagates it to its neighbors, incrementing a sequence number value for each entry that indicates its distance from the source. Using these advertisements, each router compiles a map of the network and uses it to construct its own routing table.

True or false: Link state routing protocols are preferable on an internetwork with links running at different speeds.

Answer: *True*. Link state routing evaluates the efficiency of a route based on actual transport times, not hop counts.

EXAM TIP Network+ exam candidates should be able to explain the differences between distance vector and link state routing protocol and provide examples of each.

True or false: Distance vector routing protocols impose a greater processing burden on routers than link state protocols.

Answer: *False*. Link state routing is more complex than RIP and requires more processing by the router.

RIP

The Routing Information Protocol (RIP) is a popular interior gateway protocol in the TCP/IP suite. When a RIP router starts, it generates a RIP request and transmits it as a broadcast over all of its network interfaces. Upon receiving the broadcast, every other router on any network that supports RIP generates a reply message that contains its routing table information. A reply message can contain up to 25 routes. When the router that sent the request receives the replies, it integrates the routing information in the reply messages into its own routing table.

The metric value included with each RIP route determines the efficiency of the route, based on the number of hops required to reach the destination. When routers receive routing table entries from other routers using RIP, they increment the value of the metric for each route to reflect the additional hop required to reach the destination.

RIP version 1 is widely criticized for the large amount of broadcast traffic it produces, and for its lack of a subnet mask field. Version 2 of the protocol adds a subnet mask field and support for the use of multicast transmissions instead of broadcasts.

True or false: Because it lacks a subnet mask field, RIPv1 can only be employed on networks that use classful IP addressing.

Answer: *True*. Without a subnet mask, the only way a router receiving RIPv1 data can identify the size of the network identifier in an address is to read the class from its first few bits. For subnetted classes, or for classless addressing, each RIP route must include a subnet mask.

EXAM TIP Network+ exam candidates should know the differences between the RIP versions, and be aware that RIPv2 is the predominant dynamic routing protocol used on LANs today.

True or false: RIPv1 is a distance vector routing protocol, but RIPv2 is a link state protocol.

Answer: *False*. RIP is a distance vector protocol in both versions, which uses hop counts to generate its metrics.

EIGRP

The Enhanced Interior Gateway Routing Protocol (EIGRP) is a hybrid between a distance vector and a link state protocol, relying on six vector metrics to compare the value of entries in a computer's routing table. These vector metrics are as follows:

- **Bandwidth** The bandwidth of the link between the router and the destination network
- **Load** The relative traffic saturation of the link between the router and the destination network
- **Delay** The total transmission delay between the router and the destination network
- **Reliability** The relative reliability of the link between the router and the destination network
- **MTU** The path maximum transfer unit (MTU) value of the link between the router and the destination network
- **Hop count** The number of intermediate systems between the router and the destination network

True or false: EIGRP was conceived after RIP and before Open Shortest Path First (OSPF).

Answer: *True*. Before OSPF became available, the outcry against RIP grew so loud that Cisco Systems came out with the Interior Gateway Routing Protocol (IGRP), and eventually EIGRP.

EXAM TIP Network+ exam candidates should be able to identify EIGRP as a hybrid routing protocol, combining elements of link state and distance vector protocols.

OSPF

OSPF is a link state routing protocol that, unlike RIP and most other TCP/IP protocols, uses messages that are encapsulated directly in IP datagrams, not in TCP segments or UDP datagrams. Link state routing, as implemented in OSPF, uses a formula called the Dijkstra algorithm to judge the efficiency of a route based on criteria such as the following:

- **Hop count** Though link state routing protocols still use the hop count to judge a route's efficiency, it is only part of the equation.
- **Transmission speed** The speed at which the various links operate is an important part of a route's efficiency. Faster links obviously take precedence over slow ones.
- **Congestion delays** Link state routing protocols consider the network congestion caused by the current traffic pattern when evaluating a route, and bypass links that are overly congested.
- **Route cost** The route cost is a metric assigned by the network administrator used to rate the relative usability of various routes. The cost can refer to the literal financial expense incurred by the link, or any other pertinent factor.

True or false: OSPF evaluates routes by counting the number of hops between the source and the destination.

Answer: *False*. OSPF is a link state protocol, which relies on a combination of factors to evaluate routes, rather than counting hops.

EXAM TIP Network+ exam candidates must be conscious of which routing protocols are distance vector protocols which are link state protocols, and which are those they call hybrids.

True or false: OSPF is a more suitable routing protocol than RIP for an internetwork spanning multiple sites with WAN links running at different speeds.

Answer: *True*. Because OSPF uses actual performance criteria to evaluate routes, rather than hop counts, it is a better choice than RIP for internetworks with links running at various speeds.

Convergence

Convergence is the process of updating the routing tables on all of a network's routers in response to a change in the network, such as the failure or addition of a router. Distance vector protocols such as RIP have a rather slow convergence rate because updates are generated by each router asynchronously, that is, without synchronization or acknowledgment. Link state routing protocols judge the relative efficiency of routes more precisely and have a better convergence rate than RIP.

True or false: The convergence rate of a network is based in part on the routing protocols it uses.

Answer: *True*. Link state routing protocols generally provide a better convergence rate than distance vector protocols, but there are other factors that affect convergence as well, such as the presence of relatively slow WAN links..

EXAM TIP The Network+ exam generally requires candidates to understand nothing more than the meaning of the term convergence, as it applies to dynamic routing.

True or false: Convergence rates are only an issue with networks that use dynamic routing.

Answer: *True*. On a network that uses static routing, there are no dynamic routing protocols, so convergence is only a reflection of how long it takes the administrator to update all of the routing tables on the network.

Spanning Tree Protocol

Installing multiple switches on a network can provide fault tolerance if a switch fails. However, it is also possible for the switches to begin forwarding traffic in an endless cycle, a condition called a switching loop (or a bridge loop, because it can also occur with bridges).

To address the problem of bridge looping, switches (and bridges) use a technique called the Spanning Tree Protocol (STP). STP is a data-link layer protocol that selects a non-redundant subset of switches to form the spanning tree, deactivating the others. Data circulating throughout the network uses only the switches in the tree unless a switch fails, in which case the protocol activates one of the inactive switches to replace it.

True or false: The Spanning Tree Protocol is only needed on networks with multiple switches per segment.

Answer: *True*. Switching loops only occur when there are multiple switches forwarding packets back and forth to each other.

EXAM TIP Network+ exam candidates should be familiar with the purpose of the Spanning Tree Protocol, but they do not need to know the particulars of how it works.

Virtual LANs

A virtual LAN or VLAN is a group of systems on a switched network that functions as a logical network segment. The systems on a VLAN can communicate locally with each other, but not with systems on other VLANs. The physical network is still switched, however; the VLANs exist as a logical overlay to the switching fabric, as shown in Figure 1-4.

The standard that defines the use of virtual LANs on an Ethernet network is IEEE 802.1q. Network administrators create VLANs by using a web-based configuration utility built into the switch. With this utility, administrators can specify the MAC addresses or switch ports of the systems that are to be part of each VLAN. Because VLANs are independent of the physical network, their members can be located anywhere, and a single system can even be a member of more than one VLAN. For systems in different VLANs to communicate, the switch must use routers, either physical or virtual.

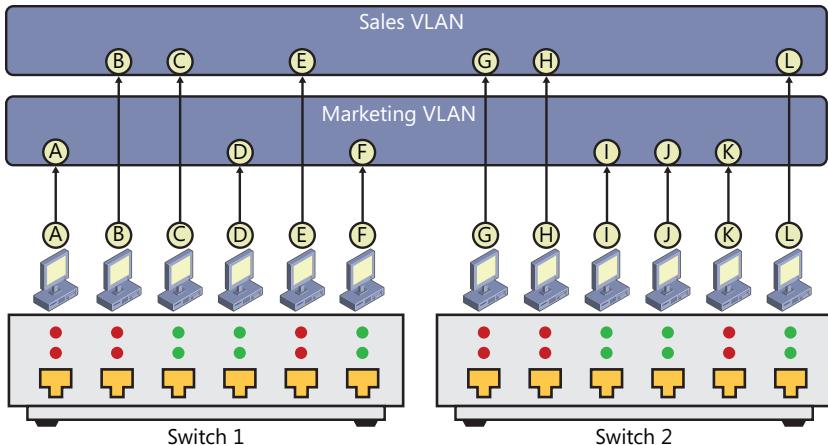


FIGURE 1-4 VLANs on a switched network.

True or false: VLANs are only necessary on networks that use switches instead of routers.

Answer: *True*. On a routed internetwork, the routers create the subnets that divide the network, so there is no need for VLANs.

EXAM TIP Network+ exam candidates must understand the need for VLANs and how they exist solely within switches.

True or false: Virtual LANs cannot communicate with physical LANs.

Answer: *False*. Using routers, VLANs can communicate with each other and with physical LANs.

MORE INFO For more information on VLANs, see “Objective 2.1: Given a scenario, install and configure routers and switches.”

Port mirroring

On a switched network, capturing traffic for monitoring and analysis is difficult, because switches forward incoming unicast traffic only to its intended recipient. A protocol analyzer connected to a standard switch port therefore has access only to one computer’s incoming and outgoing traffic, plus any broadcasts transmitted over the local network segment.

To monitor or capture all of the traffic transmitted on the network, you must plug the computer running the protocol analyzer into a switch that supports port mirroring. Switches that support port mirroring have a special port to which they send all incoming traffic.

True or false: You must employ switches that support port mirroring if you want to connect switches together to create a single network.

Answer: *False*. Port mirroring is only required if you want to use a protocol analyzer or other device to monitor or capture all of the traffic transmitted over the network.

Broadcast domains and collision domains

A broadcast domain is the group of computers that will receive a broadcast message transmitted by any one of its members. A LAN typically forms a single broadcast domain, because hubs, switches, and bridges all propagate broadcast transmissions to every system connected to them. Routers do not propagate broadcasts, however, so connecting two segments with a router creates two broadcast domains.

A collision domain is a group of network devices connected in such a way that if two devices transmit at the same time, a collision occurs. Ethernet LANs that use a shared network medium, such as bus networks or hub-based star networks, form a single collision domain, as do wireless LANs based on IEEE 802.11. Most Ethernet LANs today, however, use switches, which either create a separate collision domain for each pair of devices, in the case of a half-duplex connection; or eliminate collisions entirely, in the case of a full-duplex connection.

True or false: Splitting a hub-based Ethernet network in two by adding a bridge creates two separate collision domains.

Answer: *True*. Bridges wait until they receive an entire packet before they forward it out through the other port. Therefore, if computers on opposite sides of the bridge transmit at once, the packets will be delayed and will not collide.

EXAM TIP Network+ exam candidates must know the difference between a broadcast domain and a collision domain, and how the standard network connectivity devices affect them.

True or false: Switches create a separate broadcast domain for each pair of devices connected to them.

Answer: *False*. Switches forward broadcast packets out through all of their ports, just like hubs, so they maintain a single broadcast domain for all of their connected systems.

IGP vs. EGP

Routing protocols are generally divided into two categories: interior gateway protocols (IGPs) and exterior gateway protocols (EGPs). On the Internet, a collection of networks that fall within the same administrative domain is called an autonomous system (AS). Autonomous systems are the largest and highest-level administrative units on the Internet. Autonomous systems have unique identifiers called autonomous system numbers (ASNs), consisting of two 16-bit decimal numbers, separated by a period.

The routers within an AS use an IGP, such as the RIP or the OSPF protocol, to exchange routing information among themselves. At the edges of an AS are routers that communicate with the other ASes on the Internet, using an exterior gateway protocol (as shown in Figure 1-5) such as the Border Gateway Protocol (BGP) or the Exterior Gateway Protocol (EGP).

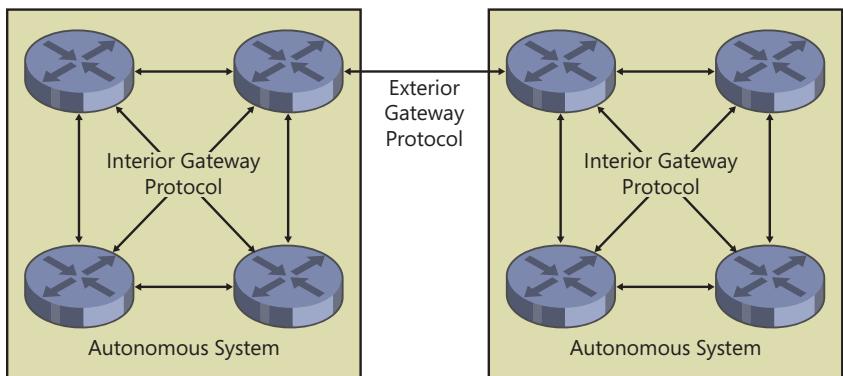


FIGURE 1-5 IGPs and EGPs within and between autonomous systems.

True or false: Link state routing protocols are used for exterior gateway routing, and distance vector protocols are used for interior gateway routing.

Answer: *False*. Both link state and distance vector protocols are used for interior gateway routing.

EXAM TIP The term “exterior gateway protocol” is both a generic name for the routing protocols used between autonomous systems and the name of a specific protocol used between ASes. In the latter, the phrase is capitalized, in the former it is not. The Network+ exam objectives refer to IGP and EGP using only the acronyms, so candidates should be familiar with both usages.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is one of the advantages of creating VLANs on a large switched network?
2. How can switching from RIPv1 to RIPv2 help to conserve bandwidth on a LAN?
3. How does the Spanning Tree Protocol prevent switching loops?
4. What are the main differences between RIPv1 and RIPv2?
5. Why is convergence an important factor in the routing process?

Objective 1.5: Identify common TCP and UDP default ports

One of the important functions of a transport layer protocol is to identify the protocol or process that generated the data it carries so that the receiving system can deliver the data to the correct application. Both TCP and UDP do this by specifying the number of a port that has been assigned to a particular process by the Internet Assigned Numbers Authority (IANA).

When a TCP/IP packet arrives at its destination, the transport layer protocol receiving the IP datagram from the network layer reads the value in the Destination Port field and delivers the information in the Data field to the program or protocol associated with that port.

All of the common Internet applications have particular port numbers associated with them, called well-known ports. The IANA has designated all of the port numbers less than 1024 as well-known ports, but not all of them are assigned to applications. TCP and UDP both maintain their own separate lists of well-known port numbers.

Exam need to know

- SMTP – 25
 - For example:* What well-known port number does SMTP use?
- HTTP – 80
 - For example:* What well-known port number does HTTP use?
- HTTPS – 443
 - For example:* What well-known port number does HTTPS use?
- FTP – 20, 21
 - For example:* What well-known port number does FTP use?
- TELNET – 23
 - For example:* What well-known port number does TELNET use?
- IMAP – 143
 - For example:* What well-known port number does IMAP use?
- RDP – 3389
 - For example:* What well-known port number does RDP use?
- SSH – 22
 - For example:* What well-known port number does SSH use?
- DNS – 53
 - For example:* What well-known port number does DNS use?
- DHCP – 67, 68
 - For example:* What well-known port numbers does DHCP use?

Ports

The well-known port numbers associated with some of the major application layer protocols in the TCP/IP suite are listed in Table 1-2.

TABLE 1-2 Well-known port numbers.

PROTOCOL	ACRONYM	TRANSPORT LAYER PROTOCOL	PORT NUMBER
Simple Mail Transfer Protocol	SMTP	TCP	25
Hypertext Transfer Protocol	HTTP	TCP	80
Hypertext Transfer Protocol Secure	HTTPS	TCP	443
File Transfer Protocol	FTP	TCP	20 (Data), 21 (Control)
TELNET	TELNET	TCP	23
Internet Mail Access Protocol	IMAP	TCP	143
Remote Desktop Protocol	RDP	TCP	3389
Secure Shell	SSH	TCP, UDP	22
Domain Name System	DNS	UDP, TCP	53
Dynamic Host Configuration Protocol	DHCP	UDP, TCP	67 (Server), 68 (Client)

True or false: FTP is an unusual protocol in that it uses two different port numbers on the server for a single transaction.

Answer: *True*. FTP servers use port 21 for control traffic, and port 20 for data. When a client sends a request for a file, it sends it to port 21. The server then opens port 20 and uses it to actually transmit the file.

EXAM TIP This is one of the few Network+ objectives that requires rote memorization. You must know the port numbers associated with the listed protocols for the exam.

True or false: HTTP servers use port 80, but HTTP clients can select their own port numbers.

Answer: *True*. HTTP and many other protocols require clients to select a port number, called an ephemeral port number, for their side of the transaction.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which of the protocols listed in this objective uses well-known ports for both the server and the client?
2. When configuring an email client that will use IMAP and SMTP, what port numbers would you use for incoming and outgoing traffic?
3. What port does a client browser use when establishing an encrypted connection to a web server?
4. What is the number of well-known ports the IANA designates at this time?

Objective 1.6: Explain the function of common networking protocols

Objective 1.6 requires Network+ exam candidates to know the basic functions of the most important protocols in the TCP/IP suite. These protocols are scattered throughout the layers of the OSI model, and many of them are covered in greater detail in other objectives. For those protocols that are not covered elsewhere, you should be familiar with their functions and where they fit into the OSI model, but there is no need to go too deeply into their intricacies.

Exam need to know

- TCP/IP suite
 - For example:* What are the primary protocols of the TCP/IP suite at the network and transport layers?
- TCP
 - For example:* What services does TCP provide that UDP does not?
- UDP
 - For example:* What types of transactions is UDP generally used for?
- DHCP
 - For example:* What is the purpose of DHCP?
- FTP
 - For example:* How does FTP differ from TELNET?
- TFTP
 - For example:* What type of file is TFTP typically used to download?
- DNS
 - For example:* Where does DNS store its information about names and addresses?
- HTTP
 - For example:* What command does an HTTP client use to request a file from a web server?

- HTTPS
 - For example:* How does HTTPS increase the security of web transactions?
- ARP
 - For example:* How does ARP resolve IP addresses into MAC addresses?
- SIP (VoIP)
 - For example:* Why is it necessary for a system to use SIP to establish a session before it sends VoIP data?
- RTP (VoIP)
 - For example:* What function does RTP provide in a Voice over IP session?
- TELNET
 - For example:* What functions can you perform on a remote computer using TELNET?
- SSH
 - For example:* Why is SSH preferable to TELNET?
- NTP
 - For example:* Why is it necessary for servers on a network to synchronize their clocks?
- POP3
 - For example:* What is the primary difference between the POP3 and IMAP protocols?
- IMAP4
 - For example:* Where do IMAP clients store their message data?
- SMTP
 - For example:* How does SMTP determine where to send email message traffic?
- SNMP2/3
 - For example:* A network management console uses SNMP to gather information from what client components?
- ICMP
 - For example:* What TCP/IP utilities use the ICMP protocol?
- IGMP
 - For example:* Why is multicasting a critical function of IPv6?
- TLS
 - For example:* Which protocol does TLS replace?

TCP/IP suite

The TCP/IP suite is a collection of protocols that span layers 2 through 7 of the OSI reference model. Together, the protocols provide a complete networking solution, with the exception of a physical layer implementation. The TCP/IP protocols are defined in documents called Requests for Comments (RFCs), published by the Internet Engineering Task Force. Some of the most important protocols in the TCP/IP suite are listed in Table 1-3.

TABLE 1-3 TCP/IP protocols.

ACRONYM	PROTOCOL	FUNCTION	OSI LAYER
ARP	Address Resolution Protocol	Resolves IP address into MAC addresses	Data-link
FTP	File Transfer Protocol	Transfers files to and from a remote host	Application
HTTP	Hypertext Transfer Protocol	Requests and receives files from web servers	Application
ICMP	Internet Control Message Protocol	Provides error messaging, diagnostic, and routing functions for IP	Network
IGMP	Internet Group Management Protocol	Provides multicast group registration services	Network
IMAP	Internet Message Access Protocol	Retrieves mail from a server and stores it permanently for client access	Application
IP	Internet Protocol	Provides connectionless network services, including addressing, routing, and fragmentation	Network
POP3	Post Office Protocol, version 3	Retrieves mail from a server and stores it temporarily for client download	Application
SMTP	Simple Mail Transfer Protocol	Provides mail transport service	Application
SNMP	Simple Network Management Protocol	Carries operational status information from agents to network management consoles	Application
TCP	Transmission Control Protocol	Provides connection-oriented services, including guaranteed delivery, error correction, and flow control	Transport
UDP	User Datagram Protocol	Provides connectionless transport service	Transport

EXAM TIP The Network+ exam might refer to TCP/IP as a protocol suite or a protocol stack; the two expressions are synonymous.

True or false: A network can conceivably run using only protocols from the TCP/IP suite.

Answer: *False*. The TCP/IP suite does not include physical layer implementations. Therefore a network cannot run without a protocol that provides the physical layer, such as Ethernet.

EXAM TIP The TCP/IP suite includes hundreds of different protocols and specifications, only a few of which are covered on the Network+ exam.

TCP

The TCP/IP suite uses two protocols at the transport layer to provide different levels of service for applications: the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP). Both TCP and UDP generate protocol data units (PDUs) that are carried inside IP datagrams. TCP is a connection-oriented protocol that provides reliable service with guaranteed delivery, packet acknowledgment, flow control, and error correction and detection.

TCP is designed for transmitting data that requires perfect bit accuracy, such as program and data files. Not surprisingly, TCP generates much more control traffic than UDP does, because it provides so many services.

True or false: Before a system can transmit data using TCP, it must exchange connection establishment messages with the destination system.

Answer: *True*. TCP performs a connection establishment procedure called a three-way handshake before sending application data.

EXAM TIP Network+ exam candidates should be familiar with the differences between the TCP and UDP protocols at the transport layer, including the services they provide and the application layer protocols that use them.

UDP

UDP is a connectionless transport layer protocol that provides unreliable service with a minimum of overhead. Many applications use UDP for short transactions that consist only of a single request and reply; others use it for data transmissions that can survive the loss of a few bits, such as audio and video streams.

True or false: The PDUs that UDP and IP create are both called datagrams.

Answer: *True*. The term datagram is used for the PDUs created by any connectionless protocol. UDP and IP are both connectionless, so they can both utilize that term.

DHCP

Dynamic Host Configuration Protocol (DHCP) is a service that automatically configures the TCP/IP client computers on a network by assigning them unique IP addresses and other parameters. Unlike its predecessor, the Bootstrap Protocol

(BOOTP), DHCP leases addresses to clients for a given period of time and reclaims them when they are no longer in use.

MORE INFO For more information on DHCP, see “Objective 2.3: Explain the purpose and properties of DHCP.”

True or false: DHCP can permanently assign IP addresses to clients.

Answer: *True*. DHCP servers can assign specific addresses manually, assign permanent addresses from a pool, and assign addresses dynamically, on a leased basis.

EXAM TIP The Network+ exam nearly always has one or more questions on DHCP, typically involving implementation details, such as the creation of scopes and relay agents.

FTP

FTP, the **File Transfer Protocol**, is an application layer TCP/IP protocol that is used by an authenticated client to connect to a server and transfer files to and from its drives. Using FTP is not the same as sharing a drive with another system on the network, nor is it a terminal emulator like TELNET. Access is limited to a few basic file management commands, and the primary function of the protocol is to copy files to a local system, not to access them in place on the server.

True or false: To use FTP, you must purchase an FTP client application.

Answer: *False*. Virtually all operating systems include a character-based FTP client, so there is no need to purchase one. Most web browsers are also capable of functioning as FTP clients.

EXAM TIP In some cases, the Network+ exam requires candidates to be familiar with basic FTP commands, such as `get`, for downloading a file from the remote system, and `put`, for uploading a file to the remote system.

TFTP

The Trivial File Transfer Protocol (TFTP) is a minimized, low-overhead version of FTP that can transfer files across a network. TFTP uses UDP at the transport layer instead of TCP and does not include FTP’s authentication and user interface features. TFTP was originally designed for use on diskless workstations that have to download an executable system file from a network server in order to boot.

True or false: TFTP can work together with DHCP to provide all the services needed to start a diskless workstation.

Answer: *True*. A diskless workstation can retrieve an IP address and other TCP/IP configuration settings from a DHCP server and then download a boot file using TFTP.

EXAM TIP Network+ exam candidates should know that TFTP uses the connection-less service provided by UDP at the transport layer, while FTP uses the connection-oriented TCP service.

DNS

The Domain Name System (DNS) is a distributed database that contains name and IP address information about the systems on a network. TCP/IP computers can use DNS servers to resolve host names into IP addresses before they initiate communication.

MORE INFO For more information on DNS, see “Objective 1.7: Summarize DNS concepts and its components.”

True or false: Each DNS server contains information about all of the hosts on the network.

Answer: *False*. Each DNS server can only contain information about a part of the network. The system is designed to distribute authoritative data among many servers and forward requests to provide access to any data a client needs.

EXAM TIP DNS is a critical element of TCP/IP communications, both on the Internet and on private networks. The Network+ exam nearly always includes questions that require an understanding of DNS and its operations.

HTTP

Communication between web servers and their browser clients is largely dependent on an application layer protocol called the Hypertext Transfer Protocol (HTTP). HTTP is a relatively simple protocol that takes advantage of the services provided by the TCP protocol at the transport layer to transfer files from servers to clients. When a client connects to a web server by typing a URL in a browser or clicking a hyperlink, the client generates an HTTP request message and transmits it to the server. HTTP consists of only two message types: requests and responses. As with many other application layer protocols, HTTP messages take the form of text commands.

True or false: Displaying a single webpage on a browser can require many HTTP request/response transactions.

Answer: *True*. Each HTTP request and response can retrieve a single file from the web server, but a single webpage can require many text and media files, which the browser must request separately.

EXAM TIP Network+ exam candidates should be aware that HTTP relies on the connection-oriented service provided at the transport layer by TCP.

HTTPS

Hypertext Transfer Protocol Secure (HTTPS) is a variant of HTTP that uses the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) security protocols to provide data encryption and server identification services. HTTPS is the accepted standard for secured Internet transactions such as online banking and e-commerce. An HTTPS connection uses the *https://* prefix in its URL and connects by default to port 443, instead of port 80, which is used by HTTP.

True or false: All websites have support for HTTPS connections. All you have to do is change the prefix in the URL.

Answer: *False*. HTTPS is not an automatic feature provided by all web servers. An administrator must enable and configure it for users to establish an encrypted connection.

ARP

The function of Address Resolution Protocol (ARP) is to reconcile the IP addresses used to identify systems at the upper layers of the protocol stack with the MAC addresses at the data-link layer. When a TCP/IP application requests network resources, it supplies the destination IP address used in the IP protocol header. ARP converts the IP address into the MAC address used in the data-link layer protocol header by broadcasting a request containing the IP address on the local network and waiting for the holder of that IP address to respond with a reply containing the equivalent MAC address.

True or false: ARP can only resolve IP addresses for systems on the local network into MAC addresses.

Answer: *True*. Because it relies on broadcast transmissions, which are limited to the local network, ARP can only resolve local IP addresses into MAC addresses.

EXAM TIP Network+ exam candidates should be careful not to confuse ARP, which resolves IP addresses into MAC addresses, with DNS, which resolves names into addresses.

SIP

The Session Initiation Protocol (SIP) is an application layer request/response protocol that Voice over IP (VoIP) uses to establish a session between two network nodes and terminate the session when the data exchange is completed. SIP does not carry the actual voice traffic; it simply sets up the call between the two parties in preparation for the data exchange.

True or false: Voice over IP relies on TCP to establish a communications session between two callers.

Answer: *False*. VoIP uses a specialized application layer protocol called SIP to establish sessions between callers. At the transport layer, SIP can use either TCP or UDP.

RTP

In a VoIP call, after the SIP protocol establishes a session, the two callers use the Real-Time Transport Protocol (RTP) to transmit the actual audio stream across the network. At the same time, the systems use the RTP Control Protocol (RTCP) to manage and monitor the transmissions.

True or false: VoIP uses application layer protocols to manage call sessions and transmit media streams.

Answer: *True*. RTP and RTCP (and the other protocols that VoIP uses) are all application layer protocols.

EXAM TIP The inclusion of SIP and RTP in this objective is the only mention of Voice over IP in the Network+ exam objectives. While candidates should be familiar with the basic functions of these protocols, there is no need for an in-depth study of VoIP for this exam.

TELNET

TELNET is a terminal emulation program that provides users with access to a text-based interface on a remote system. Unlike FTP, which is designed for file transfers and has only a limited set of file management commands that you can execute on the server, TELNET enables the remote user to execute programs and configure operating system components. As a result, TELNET and FTP tend to complement each other; together, they are known as the DARPA commands and can provide reasonably comprehensive access to a UNIX or Linux system.

True or false: TELNET and FTP provide roughly the same access to a remote system.

Answer: *False*. TELNET provide access to the command line on the remote system, while FTP provides a limited command set used for file management and transfers.

EXAM TIP Because of its lack of security, administrators today typically use a program like SSH rather than TELNET, but it still remains part of the Network+ exam objectives.

SSH

Secure Shell (SSH) is a protocol that provides encrypted command line access to another computer on the network. Used primarily by UNIX/Linux systems, SSH is an improvement over TELNET, which transmits passwords and other data over the network in clear text (that is, unencrypted) form.

True or false: SSH requires that the communicating systems have a client program and a server program.

Answer: *True*. As with TELNET, one of the computers involved in an SSH session must be running a client program and one must be running a server. Most UNIX and Linux distributions include both.

NTP

The Network Time Protocol (NTP) is an application layer protocol designed to synchronize the clocks of computers on packet-switching networks with varying degrees of latency. Because transmissions on a packet-switching network are not precisely predictable, there is no way of knowing exactly how long it will take for a packet to travel from its source to its destination. Therefore, any attempt to transmit a time signal over the network with precise accuracy is likely to be futile. NTP is designed to overcome that network latency and enable systems to synchronize their clocks with a great deal of precision.

True or false: Active Directory requires all of the domain controllers on a network to have synchronized clocks.

Answer: *True*. Because administrators can modify the Active Directory database from any domain controller, properly calibrated time stamps are necessary to ensure that changes are applied in the proper order.

POP3

The Post Office Protocol, version 3 (POP3) is designed to provide mailbox services for client computers that are themselves not capable of performing transactions with SMTP servers. Most of the clients that require a mailbox service are not continuously connected to the Internet and are therefore not capable of receiving messages any time a remote SMTP server wants to send them. A POP3 server is continuously connected and is always available to receive messages for offline users. The server then retains the messages in an electronic mailbox until the user connects to the server and requests them.

POP3 is similar to SMTP in that it communicates with clients using text-based commands and responses. As with SMTP, the client transmits commands to the server, but in POP3, there are only two possible response codes, +OK, indicating the successful completion of the command, and -ERR, indicating that an error has occurred to prevent the command from being executed. In the case of POP3, the server also sends the requested email message data to the client, rather than the client sending outgoing messages to the server as in SMTP.

True or false: POP3 servers must remain connected to the Internet at all times to receive messages destined for clients.

Answer: *True*. SMTP servers forward email traffic based on the MX resource records supplied by DNS servers. The MX records specify the address of the mail server that must be ready to receive message traffic at any time. If the server is offline, mail messages sent to it will bounce.

EXAM TIP The Network+ exam requires candidates to know the various protocols used for email messaging, the ports they use, and the differences between them.

IMAP4

Internet Message Access Protocol (IMAP) version 4 is a mailbox service that is designed to improve upon POP3's capabilities. IMAP functions similarly to POP3 in that it uses text-based commands and responses, but the IMAP server provides considerably more functionality than a POP3 server. The biggest difference between IMAP and POP3 is that IMAP is designed to store email messages on the server permanently and provides a wider selection of commands that enable clients to access and manipulate their messages. Storing the mail on the server enables users to easily access their mail from any computer.

True or false: IMAP clients store email messages in encrypted form on the client computer.

Answer: *False*. IMAP clients permanently store all email messages on the server.

EXAM TIP Network+ exam candidates should know that clients can use email protocols such as IMAP and POP3 to download messages from a mail server, but they cannot use them to send messages. For that, they must use SMTP.

SMTP

Simple Mail Transfer Protocol (SMTP) is an application layer messaging protocol that is responsible for most of the server-to-server mail traffic on the Internet. Like HTTP and FTP messages, SMTP messages are based on text commands. SMTP communications can take place between email clients and servers or between pairs of servers. In each case, the basic communication model is the same. One computer, called the sender-SMTP, initiates communication with the other, the receiver-SMTP, by establishing a TCP connection using the standard three-way handshake.

True or false: Email clients connect to SMTP servers to download their incoming email messages.

Answer: *False*. Email clients use SMTP servers for their outgoing messages, but to download their incoming messages, they must connect to a POP3 or IMAP server.

SNMP2/3

The Simple Network Monitoring Protocol (SNMP) is a TCP/IP application layer protocol and query language that specially equipped networking devices use to communicate with a central console. Many of the networking hardware and software products on the market, including routers, switches, network adapters, operating systems, and applications, are equipped with SNMP agents.

An SNMP agent is a software module that is responsible for gathering information about a device and delivering it to a computer that has been designated as the network management console. The agents gather specific information about the network devices and store them as managed objects in a management information base (MIB). At regular intervals, the agents transmit their MIBs to the console by using SNMP messages, which are carried inside UDP datagrams.

True or false: All versions of SNMP secure the data being collected from agents.

Answer: *False*. SNMPv1 has no security protection other than a community string, which functions as a password, and which systems transmit in clear text. SNMPv2 added a new security system that many people criticized as being overly complex. An interim version, called SNMPv2c, consisted of SNMPv2 without the new security system, and with the old version 1 community string instead. SNMP version 3 has standard security services, including authentication, message integrity, and encryption.

EXAM TIP For the purposes of the Network+ exam, SNMP versions 1 and 2 should be considered as unsecure protocols, while SNMP version 3 is secure.

ICMP

The Internet Control Message Protocol (ICMP) is a network layer protocol that does not carry user data, although its messages are encapsulated in IP datagrams. ICMP fills two roles in the TCP/IP suite; it provides error reporting functions, informing the sending system when a transmission cannot reach its destination, for example, and it carries query and response messages for diagnostic programs. The Ping utility, for instance, which is included in every TCP/IP implementation, uses ICMP echo messages to determine if another system on the network is able to receive and send data.

True or false: ICMP messages are encapsulated in UDP datagrams.

Answer: *False*. Unlike most TCP/IP protocols, ICMP does not use the transport services provided by TCP or UDP. Instead, its messages are carried directly within IP datagrams, with no intervening header.

EXAM TIP ICMP, apart from appearing in the Network+ objectives, is also the basis for some of the most essential TCP/IP troubleshooting tools, including Ping and Traceroute. Candidates for the exam should be familiar with these, as well as other functions of ICMP.

IGMP

Class D IP addresses ranging from 224.0.1.0 to 238.255.255.255 are reserved for multicasting purposes. A multicast transmission is simply a packet transmitted to one of those Class D addresses. However, determining which systems are part of the multicast group that recognizes that address and receives the packets is a process that involves the use of the Internet Group Management Protocol (IGMP).

True or false: Multicasts are preferable to broadcasts because they can be transmitted to systems on other networks.

Answer: *True*. Broadcast transmissions are limited to the local network because routers do not propagate them. However, routers do propagate multicasts, so they can address systems on other networks.

EXAM TIP Network+ exam candidates should be aware that multicasts are a crucial part of IPv6 communications, because IPv6 does not support broadcast transmissions.

TLS

Transport Layer Security (TLS) is the successor to Secure Sockets Layer (SSL) and is now the standard cryptographic protocol for web communications. Virtually all current web servers and browsers support TLS, as do many other Internet applications.

True or false: HTTPS is a combination of HTTP and the TLS security protocol.

Answer: *True*. When you access a secure website on the Internet by using the *https://* prefix on a URL, the web server uses TLS to secure the data it would normally deliver using only HTTP.

EXAM TIP TLS and SSL both are available in several versions providing successively greater degrees of security. However, for the Network+ exam, you need only know that these are both encryption protocols used for web traffic and that TLS is more secure than SSL.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which of the protocols covered by this objective are considered to be transport protocols?
2. Which of the application layer protocols covered by this objective are used by Voice over IP?
3. Which protocols covered by this objective carry email messages?
4. Which transport layer protocol does DNS use for most of its transmissions?
5. Which of the protocols covered in this objective is the only one that operates at the data-link layer?

Objective 1.7: Summarize DNS concepts and its components

Computers are designed to work with numbers, whereas humans are more comfortable working with words. This fundamental dichotomy is the reason why the Domain Name System came to be. Very simply, the DNS is a database service that converts computer names to IP addresses and addresses back into names.

DNS servers are a ubiquitous part of most TCP/IP networks, even if users aren't aware of them. TCP/IP communications are based solely on IP addresses. Before one system can communicate with another, it must know the other system's IP address. Often, the user supplies to a client application a friendly name for a desired server. The application must then resolve that server name into an IP address before it can

transmit a message to it. If the name resolution mechanism fails to function, no communication with the server is possible.

Exam need to know

- Summarize the concept of DNS servers
For example: How many DNS servers contain the entire Internet domain namespace?
- Summarize the concept of DNS records
For example: Which resource record type contains name-to-address mappings for IPv6 addresses?
- Summarize the concept of Dynamic DNS
For example: What network service has made it necessary to develop a mechanism for dynamically updating DNS resource records?

DNS servers

If you connect to the Internet, you use a DNS server each time you enter a server name or URL into a web browser or other application to resolve the name of the system you specified into an IP address. When a standalone computer connects to an Internet Service Provider (ISP), the ISP's server usually supplies the addresses of the DNS servers that the system will use. On a TCP/IP network, administrators configure clients with the addresses of the DNS servers they will use. This can be a manual process performed for each workstation or part of an automated DHCP configuration process.

DNS is a distributed database service; thousands of servers all over the Internet function as the authority for a small piece of the DNS namespace. By forwarding name resolution requests from server to server, it is possible to resolve any DNS name into its equivalent IP address, no matter where on the Internet the authoritative information for that name is stored.

In addition to resolving names into addresses, DNS servers can also resolve addresses into names, when necessary. This is called reverse name lookup. The DNS also plays an essential role in Active Directory Domain Services (AD DS), the Windows directory service.

True or false: Every DNS server contains a small piece of the DNS namespace.

Answer: *False.* Some DNS servers exist only to provide name resolution services to clients. They do not host any part of the DNS namespace. These are called caching-only servers.

EXAM TIP Network+ exam candidates should be familiar with the DNS domain namespace and with the messaging sequence that DNS servers use to resolve a name on the Internet.

True or false: A forwarder is a DNS server that accepts name resolution queries from other DNS servers.

Answer: *False*. All DNS servers accept name resolution queries from other DNS servers. A forwarder is a DNS server that accepts a certain type of query. When a server receives a **recursive query**, it is responsible for trying to resolve the requested name and for transmitting a reply back to the requester. If the server does not possess the required information, it must send its own queries to other DNS servers until it obtains the requested information. The resolvers in client systems nearly always send recursive queries to DNS servers.

When a server receives an iterative query, it can either respond with information from its own database or refer the requester to another DNS server. The recipient of the iterative query responds with the best answer it currently possesses, but it is not responsible for searching for the information, as with a recursive query. DNS servers processing a recursive query from a client typically use iterative queries to request information from other servers. A forwarder is a server that is configured to receive recursive queries from other servers.

DNS records

DNS servers are essentially database servers that store information about the hosts and subdomains for which they are responsible in resource records (RRs). When you run your own DNS server, you create a resource record for the name of each host that you want the rest of the network to be able to access. There are several different types of resource records used by DNS servers, the most important of which are:

- **A (32-bit Address)** Provides a name-to-address mapping that supplies an IPv4 address for a specific DNS name. This record type performs the primary function of the DNS, converting names to addresses.
- **AAAA (128-bit Address)** Provides a name-to-address mapping that supplies an IPv6 address for a specific DNS name. This record type performs the primary function of the DNS, converting names to addresses.
- **MX (Mail Exchanger)** Identifies a system that will direct email traffic sent to an address in the domain to the individual recipient, a mail gateway, or another mail server.
- **CNAME (Canonical Name)** Creates an alias that points to the canonical name (that is, the “real” name) of a host identified by an A or AAAA record. Administrators use CNAME records to provide alternative names by which systems can be identified.
- **PTR (Pointer)** Provides an address-to-name mapping that supplies a DNS name for a specific address in the in-addr.arpa domain. This is the functional opposite of an A record, used for reverse lookups only.

In addition to functioning as the authority for a small section of the DNS namespace, servers process client name resolution requests by either consulting their own resource records or forwarding the requests to another DNS server on the network. The process of forwarding a request is called a referral, and this is how all of the DNS servers on the Internet work together to provide a unified information resource for the entire domain namespace.

True or false: An MX record identifies a mail server that is ready to accept messages sent to recipients in a particular domain.

Answer: *True*. When an SMTP server receives an outgoing message from an email client, it does a DNS MX lookup of the domain in the destination email address, and sends the message to the server specified in the MX record.

EXAM TIP Network+ exam candidates must know the functions of the resource records listed in this objective.

True or false: The standard name resolution process for an Internet web server consists of DNS queries requesting AAAA records from a DNS server.

Answer: *False*. Because the Internet still uses IPv4, the standard name resolution process for an Internet name requests an A record from the server.

True or false: It is possible for a single computer to have multiple names in the DNS.

Answer: *True*. To assign multiple names to a single computer, you can create multiple A or AAAA records, or you can create a single A or AAAA record and one or more CNAME records.

Dynamic DNS

The process of adding resource records to a DNS server is called name registration. Administrators originally registered DNS names manually, by adding resource records to a text file. However, as networks grow larger and more complex, the biggest problem arising from manual name registration stems from the increasing use of DHCP servers to dynamically assign IP addresses to network workstations. Dynamic assignment of IP addresses means that workstations can have different addresses from one day to the next, and the original DNS standard has no way of keeping up with the changes.

To make the use of DNS practical for technologies that require regular updates to resource records, such as AD DS, the IETF published a document that defines a new DNS message type, called an Update, that systems like domain controllers and DHCP servers can generate and transmit to a DNS server. These Update messages can modify or delete existing resource records or create new ones, based on prerequisites specified by the administrator.

True or false: Dynamic updates enable DNS servers to connect to the systems in their resource records and query them for address changes.

Answer: *False*. Dynamic updates originate with DHCP servers and AD DS domain controllers, not with the systems specified in the resource records.

EXAM TIP The IETF standard that defines the Update message refers to the technology as Dynamic Updates, while the Network+ exam objectives refer to Dynamic DNS. There are also Internet-based services that call themselves Dynamic DNS, which enable computers with DHCP-assigned IP addresses to update a DNS

resource record on a public server whenever their addresses change. This enables a user on the Internet to access a remote computer on a home or office network, even when its address changes regularly.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which DNS resource record type can administrators use to create aliases for computers on the network?
2. A DNS client performing a reverse name resolution receives which type of resource record from the DNS server?
3. Apart from name resolution, what other critical function does DNS perform on AD DS networks?
4. In DNS terminology, what is a resolver?
5. Which type of query does a DNS server typically use when querying other DNS servers?

Objective 1.8: Given a scenario, implement the following network troubleshooting methodology

One of the key elements of troubleshooting a network problem is having a plan of action. Many troubleshooting calls are from users who are improperly using software, and these can often be cleared up immediately with some remedial training. When you are faced with what appears to be a real problem, however, you should follow a set troubleshooting procedure, which consists of a series of steps similar to those in this objective.

Exam need to know

- Identify the problem
For example: What questions should the troubleshooter ask the user?
- Establish a theory of probable cause
For example: What are all of the possible causes of the problem?
- Test the theory to determine cause
For example: What can you do to determine whether your theory is correct?
- Establish a plan of action to resolve the problem and identify potential effects
For example: What needs to be done to resolve the problem fully?
- Implement the solution or escalate as necessary
For example: Under what conditions must the problem be escalated?
- Verify full system functionality and if applicable implement preventative measures
For example: Is there anything that can be done to prevent the problem from reoccurring?

- Document findings, actions, and outcomes

For example: What mechanisms does the organization have in place to maintain a history of the problem and its solution?

Identify the problem

The first step in troubleshooting a network problem is to determine exactly what is going wrong and to note how the problem affects the network so that you can assign it a priority. It is sometimes difficult to determine the exact nature of the problem from the description given by a relatively inexperienced user, but part of the process of narrowing down the cause of a problem involves obtaining accurate information about what has occurred. Users are often vague about what they were doing when they experienced the problem, or even what the indications of the problem were.

Begin by asking the user questions like the following:

- What exactly were you doing when the problem occurred?
- Have you had any other problems with your computer lately?
- Was the computer behaving normally just before the problem occurred?
- Has any hardware or software been installed, removed, or reconfigured recently?
- Did you or anyone else do anything to try to resolve the problem?

When a computer or other network component that used to work properly now does not, it stands to reason that some change has occurred. When a user reports a problem, it is important to determine how the computing environment changed immediately before the malfunction. Unfortunately, getting this information from the user can often be difficult. On a network with properly established maintenance and documentation procedures, you should be able to determine whether the user's computer has been upgraded or modified recently.

Major changes, such as the installation of new hardware or software, are obvious possible causes of the problem, but you must be conscious of causes evidenced in more subtle changes as well. For example, an increase in network traffic levels, as disclosed by a protocol analyzer, can contribute to a reduction in network performance.

True or false: The priority you assign to a problem report should, in most cases, be based primarily on the number of users the problem affects.

Answer: *True*. Although there can be political and economic factors that affect your decision, the general rule is that the more users who are affected, the higher the priority of the problem.

Establish a theory of probable cause

After gathering all the information you can, make a list of all the possible problems that fit the circumstances, from the mundane to the extreme. A user's inability to access a website could be caused by a problem in the user's computer, a problem

in the web server, or anywhere in between. When you first begin the troubleshooting process, your list of possibilities might include everything from an unplugged network cable to solar flares. As you gather more information, you should be able to rule out a lot of the possible causes on your list and work your way down to a manageable few.

The final step of this phase is to select the item from your list that seems to be the most probable cause of the problem. Don't be afraid to question the obvious. There's an old doctors' axiom that says, "When you hear hoofbeats, think horses, not zebras." In the context of network troubleshooting, this means that when you look for the probable cause of a problem, start with the obvious cause first.

True or false: The most obvious cause of a problem is usually the correct one.

Answer: *False*. IT troubleshooting is rarely well-guided by simplistic axioms such as these. A problem's cause can be just as easily obvious as obscure.

EXAM TIP Troubleshooting questions on the Network+ exam are often scenario-based, and can contain information that is there only to distract you from the correct answer. Be prepared to use the troubleshooting procedure to eliminate the wrong answers, leaving you with the correct ones.

Test the theory to determine the cause

When you have established your theory of the probable cause of the problem, the next step is to test that theory. If you have isolated the problem to a particular piece of equipment, try to determine whether hardware or software is the culprit. If it is a hardware problem, you might replace the unit that is at fault or use an alternative that you know is functioning properly.

In some cases, the only way to test your theory involves resolving the problem. For example, if you suspect that a computer's inability to access the network is due to a bad patch cable, the only way to test your theory is to replace the patch cable with one you know is good. If that works, then your theory is confirmed.

Confirming your theory might actually resolve the problem, but that is not always so. If the problem affects multiple computers, each of which will require modifications, then you might be able to confirm your theory by modifying one, to see if your procedure works.

If your test concludes that your theory is incorrect, then you have to go back to your list of possible causes and decide which of the remaining ones is the next most probable. Then the whole testing process begins again. It is not unusual for a troubleshooter to disprove several theories before arriving at the correct one.

Depending on the size of your organization and the chain of command, you might have to escalate the problem by bringing it to someone with greater responsibility than yours, someone who can determine when or if you can safely test your theory.

True or false: The easiest way to test if a hardware component has malfunctioned is to replace it with one that you know is working properly.

Answer: *True*. Replacing the suspected component is a sure way of testing it, but it is not always the most practical or most economical way. A component that is vital to the company's operation or extremely expensive might not be easily replaceable, in which case you must find another solution.

EXAM TIP When taking the Network+ exam, do not eliminate answers because you think they are too simple. CompTIA sometimes couches simple concepts in complex language to distract you.

Establish a plan of action to resolve the problem and identify potential effects

If your theory is proven correct and your solution needs to be implemented on a larger scale, the next step of the process is to create a complete plan of what needs to be done to fully resolve the issue. The plan should include all service interruptions that will be needed and all potential effects on the rest of the network. If the plan includes taking critical network components offline, then it should include the ramifications of that downtime and scheduling recommendations for work during off hours.

It is important, throughout the troubleshooting process, to keep an eye on the big network picture and not become too involved in the problems experienced by one user (or application or LAN). While resolving one problem, you could inadvertently create another that is more severe or that affects more users.

True or false: Server troubleshooting takes precedence over user productivity.

Answer: *False*. This is almost never true, especially when user productivity is directly equated with generation of revenue. Server outages should be planned for off hours and coordinated with all of the management personnel involved.

EXAM TIP Network+ exam questions on this objective can be concerned as much with the political realities of network troubleshooting as with the technical challenges.

Implement the solution or escalate as necessary

When you have a solution to the problem mapped out and ready, it is time to implement it. If the solution falls within your area of responsibility, you can go ahead and do what is needed. However, if the solution involves other areas, or if special permission is required for the expenditures needed to execute your plan, then this is the time to escalate the issue to someone higher up in your organization's chain of command.

True or false: Escalation of a problem only occurs when a troubleshooter is unable to arrive at a satisfactory solution.

Answer: *False*. A well-organized IT department has a chain of command that specifies who is responsible for each area of the network. Escalation of a troubleshooting issue should occur whenever it falls under a superior's area of responsibility.

EXAM TIP The Network+ exam questions do not assume the existence of a particular organizational model to which a troubleshooting process must conform. The exam objectives can apply to small, informal businesses as easily as large, enterprise networks.

Verify full system functionality and, if applicable, implement preventative measures

Even if you have already performed small-scale tests to confirm your theory, after your solution is completely implemented, you must test again to confirm its success. To fully test whether the problem is resolved, you should return to the very beginning of the process and repeat the task that originally brought it to light. If the problem no longer occurs, you should test any other functions related to the changes you made, to ensure that fixing one problem has not created another.

At this point, the time you spend documenting the troubleshooting process becomes worthwhile. Repeat the procedures used to duplicate the problem exactly to ensure that the trouble the user originally experienced has been completely eliminated, and not just temporarily masked. If the problem was intermittent to begin with, it might take some time to ascertain whether the solution has been effective. It might be necessary to check with the user several times to make sure that the problem is not recurring.

If the problem ended up being the result of some network condition, or the action of a user administrator, you should consider at this point what must be done to prevent the problem from occurring again. This might involve a change to existing company policy or the creation of a new one.

True or false: Testing a solution to a troubleshooting issue involves recreating the original problem, if possible.

Answer: *True*. Recreate the original steps that caused the problem to appear, or have the original user do so, to determine whether your solution has been successful.

Document findings, actions, and outcomes

Although it is presented here as a separate step, the process of documenting all of the actions you perform should begin as soon as the user calls for help. A well-organized network support organization should have a system in place in which each problem call is registered as a trouble ticket that will eventually contain a complete record of the problem and the steps taken to isolate and resolve it.

The final phase of the troubleshooting process is to explain to the user what happened and why. Of course, the average network user is probably not interested in hearing all the technical details, but it is a good idea to let users know whether their actions caused the problem, exacerbated it, or made it more difficult to resolve. Educating users can lead to a quicker resolution next time or can even prevent a problem from occurring altogether.

True or false: Documentation of a troubleshooting effort should begin as soon as the problem is resolved.

Answer: *False*. Documentation should begin as soon as the problem is reported and continue throughout the troubleshooting process.

EXAM TIP The order of the troubleshooting steps provided in the Network+ exam objective is important. Candidates should be familiar with each step and be able to list them in the proper order.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. A user reports a problem to the help desk; after making a concerted troubleshooting effort for several hours, you are unable to resolve the issue. What should you do next?
2. It is a busy morning at the help desk, and you are currently handling three calls. One appears to be a hard drive failure in a user's workstation, one is a user unable to access a particular website, and the third consists of several calls reporting that the company email server is unavailable. Which should you handle first?
3. A user calls the help desk and reports an inability to access any network resources, whether internally or on the Internet. What should you do to determine the scope of the problem?
4. How do you test whether a network access problem is limited to a single workstation?

Objective 1.9: Identify virtual network components

In networking, virtualization is a process that adds a layer of abstraction between actual, physical hardware and the system making use of it. Virtualization is a relatively recent philosophy in network management. Although virtualization was originally a tool primarily employed for lab testing and pre-production work, administrators are now using virtual components throughout their networks, taking advantage of the flexibility that virtualization provides.

Exam need to know

- Identify the function of a virtual desktop
 - For example:* For what applications are virtual desktop products suitable?
- Identify the function of a virtual server
 - For example:* What are the advantages of virtual servers over physical servers?
- Identify the function of a virtual switch
 - For example:* How is a virtual switch different from a physical switch?
- Identify the function of a virtual PBX

For example: Can a virtual PBX provide the same service as a standard telephone system?

- Identify the difference between onsite vs. offsite virtualization

For example: Why would you want to have virtual machines stored offsite?

- Identify the function of Network as a Service (NaaS)

For example: Is NaaS more cost effective than hosting your own virtual machines?

Virtual desktops

Administrators typically use Type I virtualization products, such as Hyper-V, for server virtualization. This type of virtualization can provide the performance levels needed to run high-volume production servers. Type II virtualization provides an excellent platform for education, laboratory testing, and software evaluation. It also enables desktop users to run an instance of another operating system on a single computer, without the complications of dual booting.

In this practice, often called *desktop virtualization*, a user can run applications that are not compatible with his or her primary operating system. For example, there are several products that enable Apple Macintosh users to run an instance of Windows. Other products run on Windows 7 and enable users to install an earlier version of Windows, allowing them to run an application that has not been updated.

Windows 7 even includes a feature called Windows XP Mode, which is a fully licensed version of Windows XP that you can install on a computer running Windows 7 with Microsoft Windows Virtual PC.

True or false: Desktop virtualization is a low-cost way of deploying multiple production servers as virtual machines on a single workstation.

Answer: *False*. Type II virtualization provides a suitable platform for virtual workstations or for servers in a laboratory or classroom, but not for a production environment.

EXAM TIP Generally speaking, Network+ questions concerning desktop virtualization are referring to Type II hypervisors, while virtual servers are referring to Type I virtualization.

True or false: Hardware virtualization support is required to run any type of hypervisor product.

Answer: *False*. There are hypervisor implementations that do not require hardware virtualization support.

Virtual servers

A virtual server is a separate instance of an installed operating system running on a physical computer. Instead of having the server access the computer's hardware directly, an intervening component called a hypervisor creates a virtual machine (VM) environment, and the server operating system runs in that environment.

The hypervisor is responsible for handling all of the hardware calls that the virtual machine makes and passing them along to the correct physical hardware. When you create a virtual machine, you specify what (virtual) hardware should be in it.

The advantage of this capability is that the hypervisor can create multiple virtual machines on a single computer, sharing the physical hardware among them. Each virtual machine can then have a separate operating system instance installed on it. The instances appear to the network as separate computers, each with its own hardware, its own addresses, and its own applications. If one virtual machine suffers a software malfunction and crashes, the other virtual machines on the same computer are in most cases unaffected.

True or false: Virtual servers enable administrators to run multiple roles on a single computer without them interfering with each other.

Answer: *True*. Multiple applications running on a single operating system instance can interfere with each other, causing the whole system to crash. By running the applications on separate virtual machines, one can crash without affecting the others.

EXAM TIP Questions about Hyper-V on the Network+ exam refer to virtual servers.

True or false: Virtual servers in a production environment typically run on a Type I hypervisor.

Answer: *True*. A Type I hypervisor provides better virtual machine performance than a Type II hypervisor, so production servers usually run in a Type I environment.

Virtual switches

One of the problems that any server or desktop virtualization solution has to solve is that of network access. A physical computer usually has only one network adapter in it, but if there are multiple VMs running on that computer, each one has its own virtual adapter that needs access to the network. One way that a hypervisor can accomplish this is to use virtual switching.

To keep communication within the hypervisor, most virtualization products can create a virtual switch that enables all of the VMs on a computer to communicate with each other, just as if their network adapters were connected to a physical switch. For Type I virtualization solutions, there are also third-party virtual switch products available. These are essentially software switches that provide additional security, management, and wide area networking (WAN) services.

True or false: Virtual switches can enable virtual machines to participate in a physical network.

Answer: *True*. Virtual switches can provide virtual machines with access to the physical network through the physical network adapter in the host computer.

EXAM TIP There are several virtual switch implementations available, both as commercial and open source products. However, any questions on virtual switching on the Network+ exam will be generic, and will not involve the properties or features of any specific product.

Virtual PBX

A private branch exchange (PBX) is essentially a telephone exchange, that is, a switchboard, wholly owned and operated by a business or other private entity, rather than by a telephone company. As its core functionality, the PBX routes incoming calls to the proper extensions and provides outgoing callers with automatic access to a line. The original alternative to a PBX for a business was a key system, which required callers to push buttons to select their own lines.

Deciding on the correct telephone solution was always difficult for relatively small businesses lacking the knowledgeable staff required to maintain a PBX. This eventually led to the appearance of hosted PBX services, sometimes called virtual PBXs, in which a telephone company provided the PBX services to a customer but maintained the actual hardware at their own facility.

Another option is a software-based solution, running on a computer at the customer's site, which provides the same services as a hardware-based PBX.

The recent emphasis on cloud computing has led to the development of several hosted PBX solutions that use VoIP to provide services to customers over the Internet. Because of their decentralized nature, the actual company telephones connected by the virtual PBX service can be located anywhere, whereas a traditional PBX was limited to extensions located in the same facility.

True or false: A virtual PBX provides the same PSTN-based telephone functionality as a physical PBX.

Answer: *False*. A virtual PBX provides telephony services based on VoIP, not the Public Switched Telephone Network (PSTN).

EXAM TIP The Network+ objectives use the term "virtual PBX," which is actually the trademark of a company providing cloud-based VoIP services. However, the term can actually refer to a software-based telephony solution run on a customer's computer, or to PBX services delivered over the Internet.

Onsite vs. offsite

Because virtual machines all interface with the same hypervisor, you can easily copy or move a virtual machine from one physical computer to another. This enables administrators to easily maintain offline copies of virtual machines, so that if a physical computer fails, duplicates of its virtual servers are immediately available. Administrators can also maintain copies offsite, for backups in the event of theft or natural disaster. Some organizations maintain their entire data centers offsite, in a facility belonging to a hosting service that is responsible for its security and environmental maintenance.

True or false: Offsite datacenter hosting can be more economical than hosting the systems yourself.

Answer: *True*. In an area where office space comes at a premium, hosting virtual machines offsite can be cheaper than leasing space locally.

Network as a Service (NaaS)

Some service providers are in the business of selling access to offsite networks of virtual machines to customers; for a monthly fee, you can create a server or a network of servers at another location that runs any applications you need, just as if you were hosting them onsite. Sometimes called Network as a Service (NaaS), this concept is a progenitor of cloud computing.

True or false: NaaS eliminates some of the traditional concerns of the network administrator, such as bandwidth, fault tolerance, and environmental services.

Answer: *True*. NaaS is a pay-as-you-go arrangement that enables you to select the services you want and upgrade them as needed. Part of the arrangement is an agreed quality of service that covers fault tolerance and allowable downtime.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. How does a Type I hypervisor differ from a Type II hypervisor?
2. What relatively new telephony service has made the virtual PBX possible?
3. How do virtual servers provide network administrators with fault tolerance?

Answers

This section contains the answers to the “Can you answer these questions?” sections in this chapter.

Objective 1.1: Compare the layers of the OSI and TCP/IP models

1. The Point-to-Point Protocol (PPP) is the primary TCP/IP protocol operating at the link layer. PPP is designed for use with modems and other direct connections in which there is no need for media access control, as with Ethernet. Because it connects only two systems, PPP is called a point-to-point or end-to-end protocol. On a system using PPP, the TCP/IP protocols define the workings of the entire protocol stack, except for the physical layer itself, which relies on a hardware standard.
2. The presentation and the session layers of the OSI model do not have TCP/IP protocols dedicated exclusively to them. In most cases, application layer protocols include the session and presentation layer functions.
3. At the transport layer, the Transmission Control Protocol (TCP) provides connection-oriented service and the User Datagram Protocol (UDP) provides connectionless service.
4. The OSI reference model is defined in a document published by the International Organization for Standardization (ISO), and the TCP/IP model is defined in a Request For Comments document published by the Internet Engineering Task Force (IETF).

Objective 1.2: Classify how applications, devices, and protocols relate to the OSI model layers

1. Encryption devices function as the presentation layer of the OSI model. All of the other listed components are physical, data-link, or network layer devices.
2. The cables are the only component listed that is exclusively associated with the physical layer; the switch and the NICs are associate with the data-link layer.
3. The data-link layer and the network layer. The basic function of a switch is a data-link layer process, but to accommodate advanced features, such as VLANs, network layer capabilities are required.
4. Routers do not forward broadcast traffic; therefore they split a network into separate broadcast domains.
5. The addresses associated with the Internet Protocol (IP), running at the network layer, are 32 bits long in version 4 and 128 bits long in version 5.

Objective 1.3: Explain the purpose and properties of IP addressing

1. Class A subnets provide over 16 million hosts.
2. The subnet mask for a network with a /21 suffix is, in binary notation: 11111111 11111111 11111000 00000000; or in decimal notation: 255.255.248.0.
3. APIPA uses the 169.254.0.0/16 subnet when assigning IP addresses.
4. The last twelve zeroes in the network address can be compacted as follows: fe80::/64.
5. The OUI in the MAC address is the first three bytes: 60-EB-69.

Objective 1.4: Explain the purpose and properties of routing and switching

1. A large network connected by switches forms a single broadcast domain that can generate a huge amount of traffic. Splitting the network into VLANs enables you to create multiple, smaller broadcast domains.
2. RIPv2 supports the use of multicasts instead of broadcasts. By reducing the amount of broadcast traffic on the network, bandwidth is conserved.
3. On network segments with redundant switches, the Spanning Tree Protocol selects one of the switches to be operative, and leaves the others dormant until they are needed. This prevents the switches from forwarding packets back and forth to each other.
4. RIPv2 includes a subnet mask field that enables the protocol to support networks that use classless addressing. RIPv2 also supports multicasting, which can help to reduce the broadcast traffic on the network.

5. For an internetwork to function efficiently, the routing tables on all of its systems must be current and correct. Convergence is the process by which changes are propagated to all of the routing tables on the network.

Objective 1.5: Identify common TCP and UDP default ports

1. DHCP uses well-known ports for both client and server. This is because DHCP transactions begin before the TCP/IP settings on the client computer are configured.
2. An email client using IMAP and SMTP would use port 25 for outgoing traffic and port 143 for incoming.
3. The client browser connects with the HTTPS protocol, which uses port 443.
4. The port numbers below 1024 are reserved for use as well-known ports, so there are 1023 available.

Objective 1.6: Explain the function of common networking protocols

1. TCP and UDP are transport protocols.
2. SIP and RTP are application layer protocols used by VoIP.
3. SMTP, IMAP, and POP3 are all protocols that carry email messages.
4. DNS typically uses UDP at the transport layer.
5. The ARP protocol operates at the data-link layer.

Objective 1.7: Summarize DNS concepts and its components

1. Administrators can create aliases by using CNAME resource records.
2. A reverse name resolution request causes a DNS server to supply a PTR resource record containing an address-to-name mapping.
3. DNS enables clients to locate AD DS domain controllers on the network.
4. A resolver is a DNS client.
5. DNS servers typically send iterative queries to other servers.

Objective 1.8: Given a scenario, implement the following network troubleshooting methodology

1. The next step would be to escalate the problem to a senior administrator.
2. The email server issue appears to have the potential to affect the most people, so you should address that problem first.
3. To determine the scope of the problem, try to ascertain whether anyone else is having the same experience.
4. You can test whether a problem is limited to a single workstation by trying to reproduce the problem on another workstation.

Objective 1.9: Identify virtual network components

1. A Type I hypervisor addresses the hardware directly, while a Type II runs on top of a host operating system.
2. Voice over IP is the telephony service that has made the virtual PBX possible.
3. By creating identical virtual machines on different host computers, you can leave one VM as an offline backup to the operational one.

CHAPTER 2

Network Installation and Configuration

The Network Installation and Configuration domain accounts for approximately 23% of the CompTIA Network+ exam N10-005. Compared to the first domain, which is mainly theoretical, this domain covers more hands-on processes, such as setting up and troubleshooting the most critical of networking components.

You might know what routers and switches do, and the basics of setting up a wireless network, but the objectives in this domain cover some of the more advanced features found in high-end equipment, as well as the more basic features found in units sold for home and small office users.

These are functions and tasks and problems that you are likely to encounter when working with networks. Knowledge of these elements will benefit you long after you pass the Network+ exam.

This chapter covers the following objectives:

- Objective 2.1: Given a scenario, install and configure routers and switches
- Objective 2.2: Given a scenario, install and configure a wireless network
- Objective 2.3: Explain the purpose and properties of DHCP
- Objective 2.4: Given a scenario, troubleshoot common wireless problems
- Objective 2.5: Given a scenario, troubleshoot common router and switch problems
- Objective 2.6: Given a set of requirements, plan and implement a basic SOHO network

Objective 2.1: Given a scenario, install and configure routers and switches

For this exam objective, you must be familiar with the functions and features of high-end enterprise routers and switches, and how to install and configure them. From a physical standpoint, routers and switches are just boxes with cable connectors in them. Installing one is basically a matter of plugging it in, turning it on, and connecting the cables. It is in the configuration of the device and its features that the process can become complicated, requiring some technical expertise.

Routers and switches each typically have a configuration interface that you access either with a web browser or by plugging in a terminal.

Exam need to know

- Installing and configuring routing tables
For example: How does information get into the routing table?
- Installing and configuring NAT
For example: How does NAT protect a network from intruders on the Internet?
- Installing and configuring PAT
For example: How many registered IP addresses does a PAT router need to protect a private network?
- Installing and configuring VLAN trunking
For example: In what circumstances is trunking necessary?
- Installing and configuring managed vs. unmanaged routers and switches
For example: What are the advantages of managed routers and switches?
- Configuring router and switch interfaces
For example: Are you familiar with the parameters you might have to configure on a router or switch, such as full duplex/half duplex, port speed, IP address, and MAC filtering?
- Installing and configuring PoE
For example: What types of devices can receive power through network cables?
- Configuring traffic filtering
For example: What types of traffic do routers and switches forward and what types do they block?
- Installing and configuring diagnostics
For example: What types of diagnostic tools might you find on a router or switch?
- Configuring VTP
For example: Is there any way to automate the VLAN configuration process on enterprise switches?
- Installing and configuring QoS
For example: Do all of the routers and switches on a network have to support QoS to prioritize packets?
- Installing and configuring port mirroring
For example: Do all switches support port mirroring?

Installing and configuring routing tables

A routing table is essentially a list of network and host addresses, plus the addresses of intermediary routers that the system can use to reach them. Although different operating systems display routing tables in various formats, the information in them is generally the same. In addition, if an operating system supports IPv6 as well as IPv4, it will maintain two sets of routing table entries, possibly in different formats.

The routing table for a Windows Server 2008 R2 workstation is shown in Figure 2-1.

```

Administrator: Command Prompt
C:\Users\Administrator>route print
=====
Interface List
10...00 0c 29 5d ec 17 ....Intel<R> PRO/1000 MT Network Connection
1...........................Software Loopback Interface 1
11...00 00 00 00 00 00 e0 Microsoft ISATAP Adapter
12...00 00 00 00 00 00 e0 Teredo Tunnelling Pseudo-Interface
=====

IPv4 Route Table
=====
Active Routes:
Network Destination      Netmask      Gateway      Interface      Metric
          0.0.0.0          0.0.0.0    10.0.0.1    10.0.0.2      266
          10.0.0.0        255.255.255.0  On-link      10.0.0.2      266
          10.0.0.2        255.255.255.255  On-link      10.0.0.2      266
          10.0.0.255      255.255.255.255  On-link      10.0.0.2      266
          127.0.0.0        255.0.0.0    On-link      127.0.0.1      306
          127.0.0.1        255.255.255.255  On-link      127.0.0.1      306
127.255.255.255        255.255.255.255  On-link      127.0.0.1      306
          224.0.0.0        240.0.0.0    On-link      127.0.0.1      306
          224.0.0.0        240.0.0.0    On-link      10.0.0.2      266
255.255.255.255        255.255.255.255  On-link      127.0.0.1      306
255.255.255.255        255.255.255.255  On-link      10.0.0.2      266
=====
Persistent Routes:
  Network Address      Netmask      Gateway Address  Metric
          0.0.0.0          0.0.0.0    10.0.0.1    Default
=====

IPv6 Route Table
=====
Active Routes:
If Metric Network Destination      Gateway
  1   306 ::1/128    On-link
  10  266 fe80::64    On-link
  10  266 fe80::5058:4e95:a38f:726/128
          1   306 ff00::8    On-link
  10  266 ff00::8    On-link
=====
Persistent Routes:
  None

```

FIGURE 2-1 IPv4 and IPv6 entries in a Windows Server 2008 R2 routing table.

The data in the columns of a Windows routing table's IPv4 entries have the following functions:

- **Network Destination** Specifies the IP addresses of the networks or hosts for which the table is providing routing information.
- **Netmask** Specifies the subnet masks for the values in the Network Destination column. As with any subnet mask, the systems use the Netmask values to determine which parts of the Network Destination value are the network identifier and the host identifier.
- **Gateway** Specifies the IP addresses of the routers that systems should use to send packets to the networks or hosts identified in the Network Destination column. On an Ethernet LAN, the MAC address for the system identified by the Gateway value will become the Destination Address value in the packet's Ethernet header.
- **Interface** Specifies the IP addresses of the network interfaces that the computer should use to transmit packets to the systems identified in the Gateway column.
- **Metric** Contains values that specify the efficiency of the routes. Metric values are relative; a lower value indicates a more efficient route than a higher value. When a routing table contains multiple routes to the same destination, the system always uses the table entry with the lower Metric value.

The Windows IPv6 routing table is arranged differently but contains the same information. The Interface column is abbreviated “If” and uses numbers assigned to the computer’s interfaces, rather than the addresses of the interfaces themselves. An Interface List appears at the top of the routing table display, which in this case assigns the number 10 to the computer’s network interface adapter and the number 1 to the system’s loopback interface. The Metric, Network Destination, and Gateway columns perform the same functions as their IPv4 counterparts, using IPv6 addresses. The Netmask column is omitted, of course, because IPv6 does not use subnet masks.

The routing table on a router is often considerably longer and more complex than the one on a workstation because it contains entries for all of the networks to which it is attached, as well as entries for more distant networks, provided either manually by administrators or dynamically by routing protocols.

A router also makes greater use of the values in the Interface and Metric columns. On a multihomed system, the value in the Interface column is a crucial part of transmitting a packet correctly. Each entry must specify which interface the system should use when transmitting packets to that specific destination.

MORE INFO For basic introductory information on the function of routing tables, see “Objective 1.4: Explain the purpose and properties of routing and switching.” For more information on using the Route utility to create and manage routing table entries, see “Objective 4.3: Given a scenario, use appropriate software tools to troubleshoot connectivity issues.”

True or false: A computer can only send packets to networks that are listed in its routing table.

Answer: *False*. A computer’s routing table typically has a default gateway entry with a network address of 0.0.0.0, which the system uses for all traffic addressed to networks not listed in the table.

EXAM TIP All routing tables contain basically the same information, but they can vary in appearance and format in different operating systems. Candidates for the Network+ exam should therefore be familiar with the actual functions of the various routing table fields, and not just the headings.

True or false: A routing table can have multiple entries for a specific network address.

Answer: *True*. On a complex internetwork with many interconnected routers, there might be multiple paths to the same destination network, so the system uses the Metric value to choose the most efficient one. The type and significance of the Metric values depends on the administrator or routing protocol that created them.

Installing and configuring NAT

Network Address Translation (NAT) is a routing technique that enables computers with private (or unregistered) IP addresses to access the Internet. If you connect a network to the Internet without a translation mechanism, you must use registered

IP addresses for your computers. However, registered IP addresses are visible from the Internet, which makes them vulnerable. NAT protects the computer by enabling you to assign private IP addresses to them.

However, this also means that Internet servers, when they receive requests from the private network computers, cannot send replies to them because they do not have visible addresses. NAT solves this problem by functioning as an intermediary between the Internet and a client computer on an unregistered network. For each packet generated by a client, the NAT router substitutes a registered address for the client's unregistered address before sending it to the server. When the server sends its replies, the NAT router reverses the substitution and forwards the packet to the original client.

MORE INFO For more information on public vs. private IP addressing, see "Objective 1.3: Explain the purpose and properties of IP addressing."

True or false: NAT can only provide protection for private network computers when they are running specific applications, such as web browsers.

Answer: *False*. Because NAT functions at the network layer of the OSI model, it works with any application that communicates by using IP. Client computers on the private network can run Internet email clients, web browsers, FTP clients, or any other Internet application, and NAT provides protection against intruders.

EXAM TIP Candidates for the Network+ exam should be careful not to confuse NAT routers with proxy servers. NAT routers and proxy servers both function as intermediaries between a private network and the Internet, but they are not the same. NAT functions at the network layer, while proxy servers are application layer devices that provide protection only for specific applications.

True or false: When you use a NAT router, you must configure the clients on the Internet to make use of its services.

Answer: *False*. The NAT router's processes are invisible both to the client on the private network and to the server on the Internet. The client generates a request and sends it to a server, and the client eventually receives a reply from that server. The server receives a request from the NAT router and transmits its reply to the same router. Both the client and the server function normally, unaware of the NAT router's intervention. More importantly, the client computer remains invisible to the Internet and is protected from most types of unauthorized access originating from outside the private network.

Installing and configuring PAT

There are several different types of network address translation, which differ primarily in the number of public (registered) addresses they require. The most commonly used type is port address translation (PAT). Also known as masquerading, this method translates all the unregistered IP addresses on a network by using a single registered IP address, as shown in Figure 2-2. The NAT router uses port numbers to differentiate

between packets generated by and destined for different computers, so that multiple clients can access the Internet simultaneously. Masquerading provides the best security of the NAT types because the association between the unregistered client and the registered IP address/port number combination in the NAT router lasts only for a single connection.

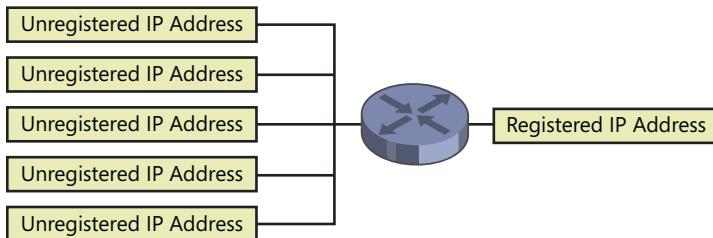


FIGURE 2-2 Port address translation (PAT).

True or false: Most of the NAT implementations in use today rely on masquerading because it uses a minimal number of public IP addresses.

Answer: True. The primary purpose of network address translation is to conserve the IPv4 address space. Without technologies such as NAT, the Internet would have depleted the supply of public IPv4 addresses long ago.

EXAM TIP Previous versions of the Network+ objectives included other NAT types, including static NAT (SNAT), a seldom-used method in which the NAT router substituted a different public address for each private address.

True or false: PAT provides computers on a private network with complete protection from all Internet-based attacks.

Answer: False. NAT by itself, even if it uses masquerading, is not a true firewall, and it does not provide ironclad security for high-risk environments. NAT effectively blocks unsolicited requests and other probes from the Internet, meaning that it prevents intruders from searching for unprotected file shares, open ports, and web or FTP servers on the private network. However, NAT does not prevent users on the Internet from launching directed Denial of Service (DoS) attacks against specific computers or from using other more complex tactics to compromise network security. NAT also does not prevent users from inadvertently running dangerous programs that initiate contact with servers on the Internet. NAT can only prevent unsolicited communication from the Internet to an unregistered computer. If the unregistered computer initiates the communication, intentionally or not, the system is vulnerable. Using private addresses provides a distinct advantage over using public addresses, but it is not a perfect solution.

Installing and configuring VLAN trunking

A large network typically has multiple switches, and configuring virtual LANs (VLANs) on the switches results in ports being dedicated to specific VLANs. You can create VLANs that span multiple switches, but to do so, there must be a means for the switches to exchange traffic outside of the VLAN substructure.

To make this possible, you designate a port on each switch as a trunk port and use it to connect that switch to the others on your network, as shown in Figure 2-3. *Trunking* enables the members of a VLAN on one switch to communicate with members of the same VLAN on another switch, just as if you connected them to the same device.

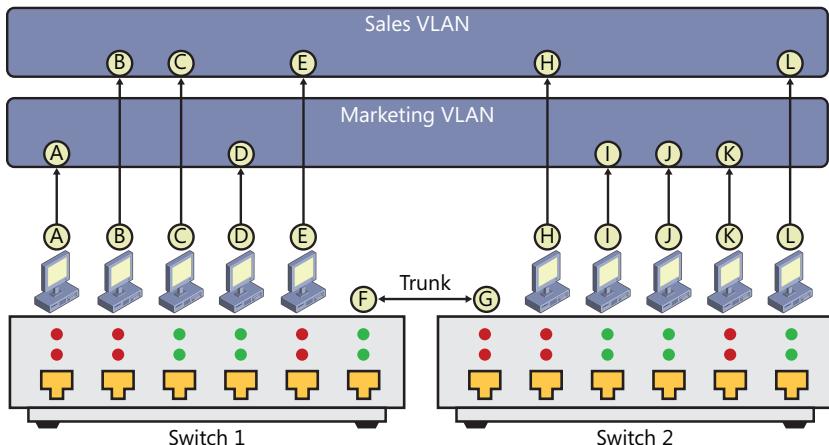


FIGURE 2-3 Trunking between VLAN switches.

True or false: Trunk ports are the functional equivalent of routers connecting the VLANs on different switches.

Answer: *False*. Trunking is a layer 2 process that joins the VLANs on different switches together, forming one larger VLAN. Packets passing through the trunk ports do not have their frames removed, as they would in a router connection, nor are broadcast messages filtered out.

MORE INFO For more information on VLANs, see “Objective 1.4: Explain the purpose and properties of routing and switching.”

True or false: Trunk ports can only carry traffic from a single VLAN.

Answer: *False*. Trunk ports conforming to the IEEE 802.1q specification add tags to the Ethernet frame. When a trunk port receives an untagged frame, it forwards it to the port’s default VLAN. If an incoming frame is tagged, the port forwards it to the VLAN specified in the tag.

EXAM TIP The term **trunking** is sometimes used as a synonym for **link aggregation**, a process in which parallel links are used to increase the throughput of a connection, such as that of a server to a network. In the Network+ exam, the term **trunking** typically refers to **VLAN trunking**.

Installing and configuring managed vs. unmanaged routers and switches

Basic switches need no configuration. You simply plug network cables into the ports and each switch automatically compiles a table containing the MAC addresses of the connected devices. Using that table, the switch can then begin selectively forwarding packets to the appropriate destinations. This is called an “unmanaged” switch.

Switches that have advanced features such as VLANs and layer 3 functionality do require configuration. These are called “managed” switches, because they provide an administrative interface. All routers are managed devices, because, if nothing else, you must be able to configure the IP addresses and other TCP/IP settings for the network interfaces.

At one time, the most common way to access a managed router or switch was to connect a terminal to a serial port in the device. Today, nearly all managed routers and switches have integrated web servers that provide access to their configuration interfaces.

True or false: Despite being primarily layer 2 devices, managed switches typically have IP addresses.

Answer: *True*. To provide client browsers with access to the integrated configuration web server in a managed switch, the switch must have an IP address.

EXAM TIP The term **managed** can also indicate that a device has an integrated SNMP agent, enabling it to send status information to a network management console. When reading Network+ exam questions, use the context to determine which use of the term is implied.

Configuring router and switch interfaces

All routers and some switches require a certain amount of configuration, just to connect to the network. Each network interface in a router requires an IP address, just as every computer on the network does, as well as other TCP/IP configuration parameters, such as a subnet mask.

At the data-link layer, all contemporary router and switch interfaces perform an automatic negotiation with the devices to which they are connected. Ethernet switches, for example, negotiate a port speed with each connected device, as well as whether they should use half duplex or full duplex communications. On wireless routers, it is also possible to use MAC filtering to specify the addresses of the computers that are allowed or denied access to the network.

True or false: MAC filtering uses access control lists to allow or deny access to the router.

Answer: *True*. An access control list (ACL) of this type contains a series of entries specifying the MAC addresses of the computers that can access the router or those that cannot.

MORE INFO For more information on MAC filtering, see “Objective 5.2: Explain the methods of network access security.”

Installing and configuring PoE

Power over Ethernet (PoE) is a technology for delivering electrical power to network devices over standard Ethernet cables, along with regular data signals. The advantages of this are several: supplying power through the network eliminates the need for an electrical socket at every network device location.

PoE also enables administrators to centralize power conditioning and monitoring services. Rather than supply each device with its own surge protector or uninterrupted power supply, the administrator can protect one single power source, and use it to supply power to devices all over the network.

True or false: A switch that supports PoE can supply all of the computers connected to the switch with all the power they need, simultaneously.

Answer: *False*. PoE cannot deliver anywhere near the amount of power that a computer requires to run. The current standard defining PoE calls for a maximum of 25.5 watts, while a typical desktop computer has at least a 250 watt power supply. Devices that can receive power using PoE include wireless access points, remote switches, VoIP telephones, and video cameras.

EXAM TIP While the Network+ exam objectives associate PoE with routers and switches, the technology is not always implemented by these devices. There are also standalone devices that can inject power into an existing cable.

True or false: PoE can deliver high power levels over any kind of cable.

Answer: *False*. PoE cannot deliver power over fiber optic cables (which do not carry electrical signals at all), and cannot deliver high power over any twist pair cable lower than Cat 5.

Configuring traffic filtering

Both routers and switches are selective about the traffic that they forward from one port to another. The process of examining incoming packets and deciding how and if to forward them is called traffic filtering. Routers receive unicast packets through one network interface and, based on their IP addresses, decide whether to transmit them out over another interface or discard them. To modify how a router filters traffic, you can change the information in its routing table.

Switches filter traffic also, but at the data-link layer of the OSI model. A switch receives packets through one of its ports and, after examining their destination MAC addresses, forwards them out through another port that is connected to the destination. The filtering in a switch is based on the information in a table that the switch compiles itself from incoming traffic.

True or false: Routers filter out broadcast traffic, but switches do not.

Answer: *True*. Routers do not forward broadcast traffic between networks, but switches maintain a single broadcast domain and forward incoming broadcasts out through all of their ports.

EXAM TIP Traffic filtering, sometimes called packet filtering, is also a function found in firewalls, which can be integrated into routers. The Network+ exam covers this type of filtering as a security device in "Objective 5.5: Given a scenario, install and configure a basic firewall."

Installing and configuring diagnostics

Troubleshooting router or switch problems typically involves accessing as much information as you have available about the devices and their operational status. Routers and switches at all price levels frequently have some sort of diagnostic capabilities that you can use to monitor their operations.

Routers and switches might have internal diagnostic features that can report on their status. The most basic of these is a log, but there might be others as well, depending on the manufacturer. Many routers and switches also include SNMP-based management agents, which enable them to report performance statistics to a network management console, such as bandwidth utilization, CPU usage levels, and temperature. Finally, some routers have the ability to run diagnostic tools such as Ping and Traceroute from the router interface. These tools can help you to determine whether the router is communicating properly with the network.

True or false: All routers and switches have a standardized set of diagnostic tools.

Answer: *False*. The diagnostic capabilities of a router or switch are wholly dependent on the manufacturer's design. There is no standard defining the types of diagnostic tools that a router or switch should possess.

MORE INFO For more information on SNMP-based network management and monitoring, see "Objective 4.4: Given a scenario, use the appropriate network monitoring resource to analyze traffic."

Configuring VLAN Trunking Protocol (VTP)

In an enterprise network environment, it is common to have multiple switches at various locations, with virtual LANs (VLANs) that are spread among the switches. Configuring the VLANs involves selecting the ports (or conceivably the addresses) that belong to each subnet. It is possible to configure the switches manually, but

this can be time-consuming, and like all manual configuration processes, is prone to input error.

Cisco Systems has developed a proprietary switch configuration mechanism called the VLAN Trunking Protocol (VTP). VTP enables administrators to configure one switch and automatically replicate the settings to the other switches on the enterprise network.

True or false: VTP can only configure switches made by Cisco Systems.

Answer: *True*. VTP is a proprietary protocol created by Cisco for its own products.

EXAM TIP The Network+ exam objectives mention the existence of the VLAN Trunking Protocol (VTP) and its basic function, but candidates for the exam should not expect to need any more detail about the protocol than that.

True or false: Users of switches made by manufacturers other than Cisco have no choice but to configure them manually.

Answer: *False*. There are also open standards defining configuration protocols for switches, including the Multiple VLAN Registration Protocol (MVRP).

Installing and configuring QoS

Quality of Service (QoS) is a means of prioritizing network traffic according to the type of data it contains or the application that generated it. For example, on a network carrying streaming video and VoIP traffic, you might want to ensure that the voice traffic is never interrupted because the video applications are monopolizing the network bandwidth.

There are several different mechanisms that can be considered QoS technologies, and they require varying levels of participation from routers and switches. Integrated services (IntServ) requires devices to communicate using a special protocol, to reserve bandwidth for specific applications. For this system to function properly, all of the devices must support that protocol. Differentiated services (DiffServ) works by adding tags to individual packets, essentially marking them with a requested priority. This too requires participation from routers and switches, which queue packets according to their priorities. However, DiffServ places less of a burden on the intermediate devices and is the preferred technology today.

EXAM TIP In addition to referring to specific standards, the term Quality of Service is also used generically to refer to a class of technologies that regulate bandwidth utilization. The term can also be used descriptively, in which case it is not capitalized.

True or false: All routers and switches include support for some type of QoS mechanism.

Answer: *False*. Routers and switch might include support for one or more QoS mechanisms, or for none at all. It is up to the network administrators to select devices that support the desired QoS standard or select a standard supported by most or all of their devices.

MORE INFO For more information on Quality of Service mechanisms, see “Objective 4.6: Explain different methods and rationales for network performance optimization.”

Installing and configuring port mirroring

Port mirroring is a feature found in some switches that enables administrators to access some or all of the network traffic processed by those devices. A hub forwards its incoming traffic out through all of its ports, so an administrator can plug a protocol analyzer into any port to monitor all of the network traffic. Because switches forward packets selectively, no single port provides access to all of the traffic, unless the switch has a port that is specifically designed to mirror data from the other ports.

True or false: Plugging a computer into a switch’s mirrored port causes it to receive all traffic sent and received by all ports.

Answer: *False*. In some switches, the mirroring feature copies all of the packets from all of the ports; in others, you must select the ports that you want to mirror. In the case of a switch with full duplex connections to network devices, it can be difficult to mirror all of the sent and received data out through a single port, and some packets might be lost.

EXAM TIP The word **port** in the term port mirroring refers to the physical cable connectors in a switch, not to the port numbers that TCP and UDP use to identify application layer protocols, as in Network+ objective 1.5. Be sure not to confuse the two uses of the term in exam questions.

True or false: All mirrored ports must be labeled as such on a switch.

Answer: *False*. Various manufacturers use different names for mirrored ports. Cisco Systems, for example, calls it a Switched Port Analyzer (SPAN).

MORE INFO For more information on monitoring network traffic using a protocol analyzer, see “Objective 4.4: Given a scenario, use the appropriate network monitoring resource to analyze traffic.”

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. How does network address translation protect a privately addressed network from Internet-based intruders?
2. Must you purchase special devices that support PoE to supply them with power through your network cables?
3. How is port mirroring a potential hazard to network security?
4. What capability does trunking provide on enterprise network switches?
5. Why is Quality of Service a more important feature on routers and switches now than it was five years ago?

Objective 2.2: Given a scenario, install and configure a wireless network

In the most recent iterations of the Network+ exam, CompTIA has reduced the coverage of LAN protocols to reflect the almost complete obsolescence of technologies such as Token Ring in the marketplace, in favor of Ethernet, which now dominates. At the same time, however, the exam's coverage of wireless LAN technologies has expanded considerably, to the point at which candidates are now expected to be able to install and troubleshoot wireless networks at least as well as Ethernet LANs.

For this objective, exam candidates must be familiar with the various types of wireless networking products on the market and the circumstances governing their selection for specific applications. Selecting appropriate products is a large part of installing an effective wireless LAN, as is configuring the components to suit the performance and security requirements of the organization.

Exam need to know

- Locating access points

For example: How does moving an access point affect network security?

- Installing antennae

For example: Can replacing an antenna improve the performance of a wireless network?

- Preventing interference

For example: Can wireless LAN signals penetrate concrete walls as easily as drywall?

- Selecting frequencies

For example: Do all 802.11 products use the same frequencies?

- Configuring channels

For example: How does channel selection affect wireless network performance?

- Understanding wireless standards

For example: Are all of the 802.11 wireless LAN standards backwards compatible?

- Disabling SSID broadcasts

For example: Can configuring an access point not to broadcast its SSID increase the security of the network?

- Selecting compatible standards

For example: Can 802.11a, b, g, and n devices work together?

Locating access points

A site survey is the first step in a wireless LAN deployment. Even though you might have already examined the location before rejecting a wired network solution, you should take another look when considering wireless; the criteria governing a wireless installation can be different. Distances are still a concern, because wireless devices have limited range, but the measurements you take will have different implications. For a wired network, you try to estimate the length of cable routes inside walls and ceilings, whereas wireless distances are direct routes through walls and other obstructions. You must have some idea of the maximum distances your devices will have to transmit. If those distances are well beyond the estimated ranges of the equipment you are considering, then you might have to add some form of range enhancement, or consider installing multiple access points.

Using the information from your site survey, you should be able to determine approximate locations for your access points. Obviously, your first priority is to provide network access to all of your clients, but there are also security considerations involved in access point placement.

True or false: Placing an access point too near an outside wall can reduce the security of your network.

Answer: *True*. Assuming no other security measures, an access point near an outside wall can enable individuals with portable computers to access your network from outside the building. These individuals would not only be stealing your bandwidth; they could also engage in illegal activities without your knowledge.

EXAM TIP The Network+ exam is likely to include troubleshooting questions in which relocating an access point is a possible solution for intermittent connection problems. This can be a problematic remedy, however, because moving the access point closer to some systems might move it farther away from others.

Installing antennae

Access points usually come equipped with dipole antennae, which are omnidirectional. Therefore, placing access points near the center of the space you want to service is the usual procedure. However, some access points permit you to connect an external antenna, and there are a variety of other antenna designs that can provide greater or more directional coverage.

Directional antennae aim their signals in a single direction, providing greater range in return for the narrower focus. For example, a beam antenna generates a tightly focused signal to a single remote point, such as a distant room or building. A patch antenna mounts on a wall and directs a hemispherical beam towards the inside of a room only. The strength of an antenna, called the gain, is measured in decibels.

True or false: An omnidirectional antenna should be located somewhere near the center of the basic service set.

Answer: *True*. A basic service set (BSS) is a collection of wireless devices serviced by a single access point. Placing an omnidirectional antenna in the center of the BSS provides the largest possible area for wireless devices.

EXAM TIP Most or all of the wireless networking questions on the Network+ exam concern infrastructure networks, in which an access point provides wireless computers and other devices with access to a wired network. The alternative is an ad hoc network, in which two or more wireless devices communicate with each other directly, without an access point.

True or false: Replacement antennae are nearly always used on infrastructure networks, rather than ad hoc networks.

Answer: *True*. Replacement antennae are typically used on the access points of infrastructure networks. It is comparatively rare for a wireless network interface adapter in a computer to have a replaceable antenna.

Preventing interference

Wireless transmission distances are subject to interference from obstructions. The more walls there are between two wireless stations, the shorter the distance across which those stations can transmit successfully. You must also consider the composition of the walls. A typical office with drywall dividers will generate much less interference than a building with cinderblock walls.

You must consider other potential sources of interference as well. Refrigerators, microwave ovens, and other electrical equipment all generate electromagnetic interference that can affect wireless signals, some of them in maddeningly intermittent ways. A microwave oven, for example, might block all wireless traffic in the vicinity, but only when it's running. This is the sort of situation that frustrates technical support providers to no end.

True or false: Wireless LAN transmissions are line of sight only.

Answer: *False*. Wireless LAN devices can transmit through walls and other obstructions, but those obstructions do interfere with the signals, weakening them and reducing the transmission range of the devices.

EXAM TIP Interference is one factor that Network+ exam questions can use in a wireless network troubleshooting context. Interference, access point location, and antenna gain are all factors that can affect network transmission quality in roughly the same way. Be sure to consider all of these factors when trying to determine the cause of a wireless network problem.

True or false: Locating an access point in a kitchen full of appliances can reduce its transmission range.

Answer: *True*. Appliances such as refrigerators and microwave ovens, aside from being large metal obstructions, are also sources of electromagnetic interference that can disturb wireless network signals.

Selecting frequencies

Most 802.11 wireless networks in operation today are based on the 802.11b/g standards, using the 2.4-GHz frequency band that occupies the 83 MHz of bandwidth between 2.4000 and 2.4835 GHz. These frequencies are unlicensed in most countries/regions, and as a result, the 2.4-GHz band is comparatively crowded with signals from other wireless consumer devices.

The 802.11n standard reintroduces the use of the 5-GHz band from 802.11a as an option. The 5-GHz band is relatively uncrowded, but hardware implementations that support the 5-GHz band are rare and are found at the high end of the price range. The 802.11ac standard, in its current form, will use only the 5-GHz band.

True or false: In a large office building with many separate wireless networks in it, selecting hardware that can use the 5-GHz frequency band can be a good idea.

Answer: *True*. Most or all of the other wireless networks in the building will probably be using the 2.4-GHz band, and it's likely that they will conflict with each other. By using 5-GHz equipment, your network will be less likely to suffer from interference with the others around you.

EXAM TIP Discussion of wireless frequencies in the Network+ exam is typically concerned only with selecting between the 2.4 and 5-GHz bands.

True or false: The 5-GHz band provides more frequency space than the 2.4-GHz band.

Answer: *True*. The 5-GHz band contains many more channels than the 2.4-GHz band, and they are non-overlapping.

Configuring channels

The wireless LAN standards divide the frequency band that a given technology uses into channels, so that multiple networks can coexist in the same area by using different parts of the available bandwidth. The channels defined by the standards up to and including 802.11g are 20 to 22 MHz in width (depending on the type of modulation).

For example, in implementations using DSSS modulation, the channels are 22 MHz wide, and the 2.4-GHz band contains channels that are 5 MHz apart. This enables the standard to define 13 channels in that band, as shown in Figure 2-4.

Of course, spacing 22 MHz channels 5 MHz apart means that the channels are going to overlap, making it possible for networks using different channels to interfere with each other. This can result in the need for retransmissions at the data-link layer, reducing network throughput and increasing latency. Therefore, in the 2.4-GHz band, it has become a common practice to favor channels 1, 6, and 11, because they do not overlap and do not interfere with each other, as shown in Figure 2-5.

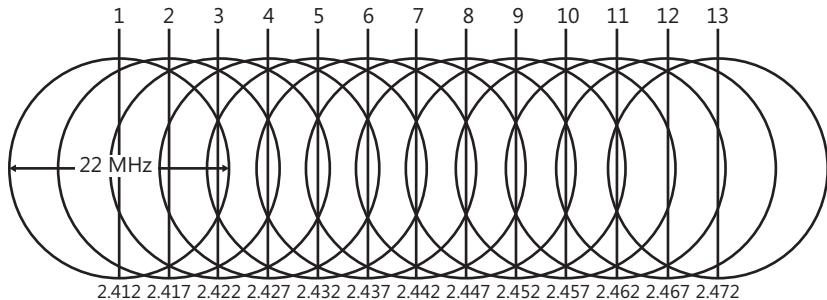


FIGURE 2-4 The 22-MHz channels in the 2.4-GHz band.

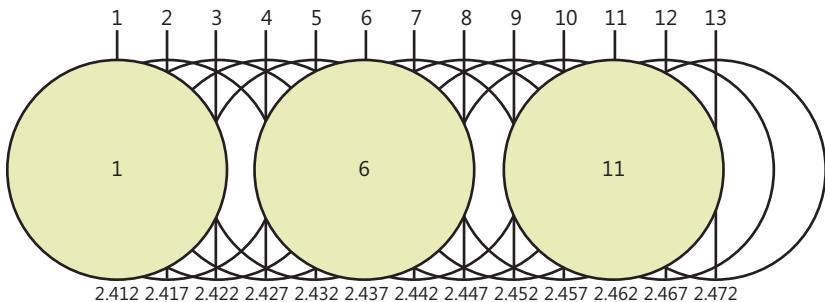


FIGURE 2-5 Non-overlapping 22-MHz channels in the 2.4-GHz band.

This practice has persisted even on 802.11g networks, which is unfortunate, because the OFDM modulation that 802.11g uses creates channels that are 20 MHz wide, not 22 MHz. With 20-MHz widths, the non-overlapping channels are 1, 5, 9, and 13.

True or false: Before wireless devices can connect to the network, you must configure both the access point and each device to use the same channel.

Answer: *False*. You can configure the access point to use a particular channel, but the wireless adapters in computers and other devices automatically configure themselves to use the same channel as the access point.

True or false: In a large wireless LAN installation with multiple access points, you must configure the access points with the same SSID and different channels.

Answer: *True*. The access points must use different channels so that they do not interfere with each other. This is the type of situation in which administrators often apply a multiple-channel architecture, using non-conflicting channels for adjacent access points, as shown in Figure 2-6.

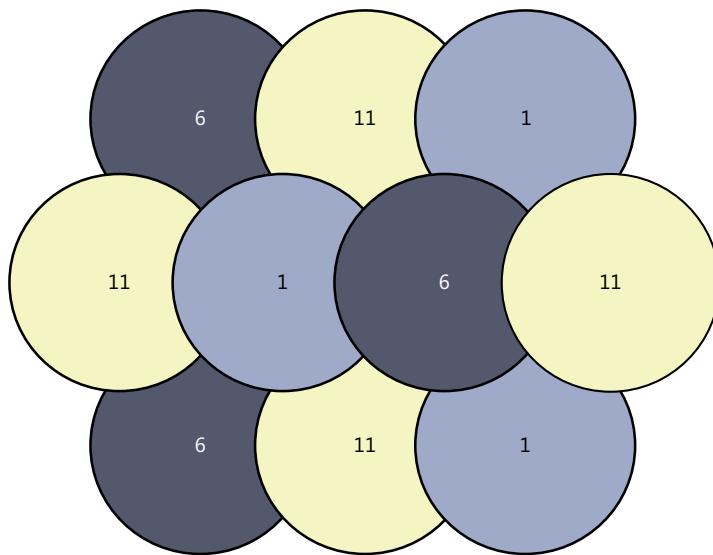


FIGURE 2-6 A multiple-channel architecture.

Understanding wireless standards

The wireless LAN equipment on the market today is based on the 802.11 standards published by the Institute of Electrical and Electronics Engineers (IEEE), from the same LAN/MAN Standards Committee, IEEE 802, that publishes the 802.3 Ethernet standards.

As with the 802.3 Ethernet standard, the IEEE has updated and expanded on the 802.11 specification several times over the years, increasing the maximum transmission speed of the network and altering the frequencies and modulation techniques. The standard publications and their basic specifications are listed in Table 2-1.

TABLE 2-1 IEEE 802.11 standards.

STANDARD	FREQUENCY (GHZ)	TRANSMISSION RATE (MBPS)	MODULATION TYPE	RANGE (INDOOR/OUTDOOR) (METERS)
802.11-1997	2.4	1, 2	DSSS, FHSS	20/100
802.11a-1999	5	6 to 54	OFDM	35/120
802.11b-1999	2.4	5.5 to 11	DSSS	38/140
802.11g-2003	2.4	6 to 54	OFDM, DSSS	38/140

STANDARD	FREQUENCY (GHZ)	TRANSMISSION RATE (MBPS)	MODULATION TYPE	RANGE (INDOOR/OUTDOOR) (METERS)
802.11n-2009	2.4 and 5	7.2 to 288 (at 20 MHz) 15 to 600 (at 40 MHz)	OFDM	70/250
802.11ac (Draft)	5	433 to 867 (at 80 MHz) 867 Mbps to 6.93 Gbps (at 160 MHz)	QAM	Undetermined

True or false: Wireless equipment conforming to the IEEE 802.11n standard can increase transmission speeds compared to previous standards by transmitting multiple signals over separate antennae.

Answer: *True*. *Multiple-Input Multiple-Output (MIMO)* is a physical layer enhancement that enables wireless devices to multiplex signals over a single channel simultaneously, by using a technique called Spatial Division Multiplexing (SDM). Each 802.11n device has an array of transmit and receive antennae and a transceiver capable of sending and receiving separate signals by using separate frequencies.

EXAM TIP Wireless networking introduces several technologies that may be alien to the average network administrator. However, the differences between the various modulation systems and other technical aspects of radio frequency communications are not knowledge required for the CompTIA Network+ exam, nor, for that matter, are they required by the administrator responsible for installing and maintaining a typical wireless LAN.

True or false: An 802.11n device with a 2×2:2 rating is the fastest available under the current standard.

Answer: *False*. 802.11n designations describing the MIMO capabilities of each device use the format $a:b:c$, where a is the number of transmit antennae in the device, b is the number of receive antennae, and c is the number of data streams that the radio in the device supports. The maximum configuration defined by the standard is $4\times4:4$, indicating that a device has four transmit and four receive antennae, and can send or receive on four channels at once.

Disabling SSID broadcasts

The *service set identifier (SSID)* is a 32-bit string that identifies a basic service set and all of its members. The SSIDs are the names of the wireless LANs you see when you use a client to scan for a network to join. In an infrastructure network, the administrator typically assigns an SSID to the access point when configuring it.

If not, the AP uses a default name set at the factory. In an ad hoc network, the first device joining the network sets the SSID.

To connect a device to an access point, you select it by its SSID. In most cases, computers and other devices scan for wireless networks and display a list of the SSIDs they find, from which you can choose. This works because access points broadcast their SSIDs by default. However, some access points enable the administrator to suppress the SSID broadcasts, as a security measure. Clients can still connect to the AP, but they must know the SSID of the network to do so.

True or false: Disabling the SSID broadcasts on an access point eliminates the need for any other security protocols.

Answer: *False*. Disabling SSID broadcasts provides security against only the most rudimentary attacks. There are many tools available that can locate a wireless network, even in the absence of an SSID broadcast. In fact, even when an access point does not broadcast SSIDs, devices still transmit the SSID in clear text as they establish a connection to the access point.

EXAM TIP The Network+ exam requires candidates to be familiar with the distinction between the ad hoc and infrastructure topologies, as well as with terms such as "access point," "service set," and "service set identifier."

Selecting compatible standards

IEEE 802.11a represented a fundamental change from 802.11. The 802.11a standard calls for the use of the relatively vacant 5-GHz band and a different form of modulation called Orthogonal Frequency-Division Multiplexing (OFDM). The data transfer rate can be as high as 54 Mbps, with fallbacks down to 6 Mbps.

Products actually arrived on the market after 802.11b devices had achieved a considerable popularity, and they cost significantly more. Therefore, dedicated 802.11a equipment did not sell well. Later devices supporting both the 802.11a and 802.11b standards eventually came to market, and cross-compatibility between standards soon became a major selling point.

The *IEEE 802.11g* standard built on the 802.11b technology by adopting the OFDM modulation from 802.11a while retaining the 2.4-GHz frequency from 802.11b. The result was a new standard that was fully backward compatible with 802.11b equipment but that pushed the maximum transfer rate up to 54 Mbps.

The *IEEE 802.11n* standard introduced several modifications, including the potential doubling of channel widths and the addition of the 5-GHz frequency band from 802.11a to the standard 2.4-GHz band from 802.11b/g. 802.11n also includes two innovations, MIMO and frame aggregation, which can push network transmission speeds well beyond the 54 Mbps realized by 802.11g to levels as high as 600 Mbps. Wireless devices supporting multiple standards are now the industry norm, with many supporting 802.11b/g/n, and in some cases 802.11a/b/g/n.

True or false: All of the IEEE 802.11 standards support both ad hoc and infrastructure networking.

Answer: *True*. You can use devices conforming to any of the standards to build an infrastructure network, or connect two devices together in an ad hoc network.

EXAM TIP Candidates for the Network+ exam must be familiar with the characteristics of the four main standards upon which wireless LAN hardware devices are based: 802.11a, 802.11b, 802.11g, and 802.11n.

True or false: IEEE 802.11g devices can connect to an IEEE 802.11n access point using the 5-GHz frequency band.

Answer: *False*. Equipment conforming to the 802.11b and 802.11g standards can only use the 2.4-GHz band, and therefore cannot connect to a 5-GHz network, no matter what standard it uses.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Using 802.11n equipment with what frequency is likely to provide the best performance from your wireless LAN?
2. How does MIMO increase transmission speeds on an 802.11n network?
3. How do you connect a computer to a wireless network when the access point is configured not to broadcast its SSID?
4. On a network with multiple access points, why must you configure each one to use a different channel?

Objective 2.3: Explain the purpose and properties of DHCP

The Dynamic Host Configuration Protocol (DHCP) is a service that network administrators configure with ranges of IP addresses and other settings. Computers and other devices configured to run as DHCP clients contact a DHCP server at boot time, and the server assigns each one a set of appropriate TCP/IP parameters, including a unique IP address. The computer uses these parameters to configure its TCP/IP client, and network communication commences with no manual configuration necessary. For this exam objective, you must be familiar with the basic elements of the DHCP transaction and how to use them.

Exam need to know

- Explain the purpose and properties of static vs. dynamic IP addressing
For example: Why would an administrator use DHCP for static addressing?
- Explain the purpose and properties of DHCP reservations
For example: When is it necessary to create DHCP reservations?
- Explain the purpose and properties of DHCP scopes
For example: What is a scope and how does a DHCP server use one?

- Explain the purpose and properties of DHCP leases
For example: What happens when a DHCP lease expires?
- Explain the purpose and properties of DHCP options, including DNS servers and suffixes
For example: Why are options needed for DHCP address assignments?

Static vs. dynamic IP addressing

The primary function of DHCP is to assign IP addresses and to accommodate the needs of all types of client systems. The DHCP standard defines three types of address assignment:

- **Manual allocation** The administrator configures the DHCP server to assign a specific IP address to a given system; the IP address will never change unless it is manually modified.
- **Automatic allocation** The DHCP server assigns permanent IP addresses allocated from a pool; the addresses do not change unless they are manually modified by the user or the administrator.
- **Dynamic allocation** The DHCP server assigns IP addresses from a pool by using a limited-time lease, so that an address can be reassigned if the client system doesn't periodically renew it.

EXAM TIP The terms manual, automatic, and dynamic allocation are those used by the DHCP standards. Most DHCP server implementations support all three types of address allocation, but don't necessarily use these exact names. The Network+ exam objectives use the term static addressing to refer to a manually configured TCP/IP client and dynamic addressing to refer to one configured by using DHCP.

True or false: The only way to create a static IP address assignment is to use DHCP manual allocation.

Answer: *False*. Any unchanging IP address is considered to be a static address, no matter how it is configured. You can therefore use DHCP or simply configure the address manually through the interface provided by the operating system.

True or false: Of the three types of address allocation defined in the DHCP standards, only dynamic allocation is designed to conserve the IP address space.

Answer: *True*. When a DHCP server dynamically allocates IP addresses, clients lease addresses for a limited time period. If a client continues to need its address, it may renew the lease prior to expiration. If the client stops using the address, because the computer is turned off or moved to another subnet, the lease eventually expires and the address returns to the pool to be reassigned. This prevents addresses in the subnet from being wasted on systems that are not using them.

MORE INFO For more information on IP addressing, see "Objective 1.3: Explain the purpose and properties of IP addressing."

DHCP reservations

In the Microsoft DHCP Server service, dynamic allocation is the default. If you want to use automatic allocation, you must change the Lease Duration setting to unlimited. For manual allocation, you create what the Microsoft server and the Network+ objectives call an *address reservation*.

True or false: DHCP reservations deplete the IPv4 address space, whether a device is actually using them or not.

Answer: *True*. A DHCP reservation is a permanent IP address assignment that does not come out of a scope and is not returned to the server for reassignment.

True or false: To create a DHCP reservation, you must know the MAC address of the device that will receive the IP address.

Answer: *True*. DHCP can only assign a reserved IP address to a specific device if it can identify it using a MAC address.

DHCP scopes

In DHCP terminology, a *scope* is a range of IP addresses on a particular subnet that a DHCP server uses as a pool for its lease assignments. When a client requests an address, the server selects the first available one from an appropriate scope and creates a lease for it. If the lease expires, the client must stop using the address and the server returns it to the pool for reassignment. This enables the server to utilize all available addresses and keep track of all the address assignments on the network.

True or false: A DHCP scope must contain addresses from a single subnet.

Answer: *True*. If you have more than one subnet on which you want to assign IP addresses using DHCP, you must create a separate scope for each one.

EXAM TIP The word “scope” is not an official term used in the DHCP standards, but the Network+ objectives mention it specifically, as well as “reservation,” another term from the Microsoft DHCP implementation.

True or false: Administrators can split the addresses of a single subnet among scopes on multiple DHCP servers.

Answer: *True*. For fault tolerance purposes, it is a common practice for administrators to use two or more DHCP servers to host the addresses for a single subnet.

DHCP leases

When a DHCP server uses dynamic allocation, clients initiate the lease negotiation process by broadcasting DHCPDISCOVER messages. DHCP servers then reply with DHCPOFFER messages that contain IP addresses and options that specify the nature of the address lease agreement. These options include the IP Address Lease Time, the Renewal (T1) Time Value, and the Rebinding (T2) Time Value.

MORE INFO As an example, Microsoft DHCP Server has a default IP Address Lease Time of eight days. The Renewal (T1) Time Value defaults to 50 percent of the lease time and the Rebinding (T2) Time Value to 87.5 percent of the lease time. Other implementations might have different values, and administrators can specify values for these intervals in the DHCP server configuration.

When a client system successfully leases an address, it has no further communications with the server until the system restarts or it reaches the T1 or renewal time. At that point, it begins transmitting DHCPREQUEST unicast messages to the server that assigned its IP address, to renew the lease. If the server receives the message, it responds with a DHCPACK message that renews the lease. No further communication is necessary until the next renewal.

If the server cannot renew the lease, it terminates the transaction and the lease. The client must then restart the entire lease negotiation process. If the client receives no response to its DHCPREQUEST, it retransmits the message each time half of the interval between the current time and the T2 time has expired. When the lease time hits the T2 point, the client begins transmitting its DHCPREQUEST messages as broadcasts to solicit an IP address assignment from any available server.

If the lease time expires with no response from any server, the client releases the IP address and must begin the whole lease negotiation procedure again.

True or false: DHCP servers can only lease IP addresses to clients on the same subnet.

Answer: *False*. DHCP relies on broadcast transmissions for the lease negotiation process, which are limited to the local subnet. However, the DHCP standard also defines the functionality of a DHCP relay agent, which is an intermediary on another network that forwards the broadcasts to and from the DHCP server. DHCP relay agents enable a DHCP server to assign addresses to clients anywhere on the network.

EXAM TIP In many cases, when a computer configured to use DHCP fails to obtain an IP address, it reverts to automatic assignment of a link-local address on the 169.254.0.0/16 network. Windows calls this process Automatic Private IP Addressing (APIPA), but it also goes by other names, such as zero-configuration networking. Although APIPA is not directly associated with DHCP in the Network+ exam objectives, it is important for candidates to understand the relationship between the two technologies. For more information about APIPA, see “Objective 1.3: Explain the purpose and properties of IP addressing.”

True or false: When you move a DHCP client computer from one subnet to another, it automatically leases a new IP address from a server on the new subnet.

Answer: *True*. DHCP is designed to conserve the IP address space by releasing addresses that are no longer in use and returning them to the scope. When you move a client to a different subnet, the existing address lease is terminated, and the computer negotiates a new address lease with a DHCP server on the new subnet.

DHCP options

The DHCP message format contains an options field, which is a catchall area designed to carry the various parameters (other than the IP address) used to configure the client system's TCP/IP stack. Because you can configure a DHCP server to deliver many options to clients, defining separate fields for each one would be impractical.

The DHCP Message Type option identifies the overall function of the DHCP message and is required in all DHCP packets. The DHCP message types are as follows:

- **1 – DHCPDISCOVER** Used by a client system to locate DHCP servers and request an IP address
- **2 – DHCPOFFER** Used by a server to offer an IP address to a client
- **3 – DHCPREQUEST** Used by a client to request a specific IP address assignment or to renew a lease
- **4 – DHCPDECLINE** Used by a client to reject an IP address offered by a server
- **5 – DHCPACK** Used by a server to acknowledge a client's acceptance of an offered IP address
- **6 – DHCPNACK** Used by a server to reject a client's acceptance of an offered IP address
- **7 – DHCPRELEASE** Used by a client to terminate a lease
- **8 – DHCPINFORM** Used by a client that has already been assigned an IP address to request additional configuration parameters

The other options defined in the DHCP standard fall into several categories. The most commonly used options contain the basic TCP/IP configuration parameters used by most client systems, such as the following:

- **Subnet Mask (code 1)** Specifies which bits of the IP address identify the host system and which bits identify the network where the host system resides
- **Router (code 3)** Specifies the IP address of the router (or default gateway) on the local network segment that the client should use to transmit to systems on other network segments
- **Domain Name Server (code 6)** Specifies the IP addresses of the servers that the client will use for DNS name resolution
- **Host Name (code 12)** Specifies the DNS host name that the client system will use
- **Domain Name (code 15)** Specifies the name of the DNS domain that the client should use when resolving host names

The DHCP standard defines a great many other options that administrators can use to identify other servers on the network and configure the behavior of the TCP/IP client, the DHCP client, and the data-link layer interface. DHCP also includes the ability to support vendor-specific options provided in custom implementations.

True or false: Without the TCP/IP configuration options, a DHCP client would not be able to function on the network.

Answer: *True*. The DHCPOFFER message provides an IP address, but the TCP/IP configuration options include the subnet mask and other parameters that the TCP/IP client needs to function.

EXAM TIP The Domain Name option in the DHCP standards is the equivalent of what the Network+ objectives and the Windows TCP/IP client refer to as the DNS suffix. The DHCP client appends the domain name specified in this option to all unqualified host names (that is, host names lacking a domain) that it attempts to resolve.

True or false: It is only possible to define a single set of options for all of a DHCP server's address assignments.

Answer: *False*. DHCP server implementations typically distinguish between scope options and server options. Scope options apply only to a particular scope, while server options apply to all of the server's clients.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is the advantage of creating DHCP reservations to assign static IP addresses, when you can simply configure the computers' addresses manually?
2. The DHCP Router option specifies the default gateway address that the client should use. Why would you configure this as a scope option rather than a server option?
3. What changes must you make to the default settings in Microsoft DHCP Server to use automatic address allocation?
4. When creating a reservation for an IP address, how does the DHCP server know which host should receive the IP address?
5. Which DHCP message type contains client configuration settings, but does not contain an IP address?

Objective 2.4: Given a scenario, troubleshoot common wireless problems

Objective 2.2 requires Network+ candidates to be familiar with wireless networking standards and implementations. This objective concentrates on troubleshooting those implementations. Wireless LANs have their own unique problems, many of which are mentioned earlier in this chapter.

The first step in troubleshooting a wireless LAN is to make sure that what you have is actually a wireless problem. If a wired computer experiences the same problem, then you must look at something other than the wireless components. If the problem occurs only on wireless equipment, then the next step is to consider whether there is any connectivity at all. If the problem is that your wireless devices are not connecting to the network at all, then it's time to look at the issues in this objective.

Exam need to know

- Troubleshooting interference problems
 - For example:* What devices can interfere with wireless network transmissions?
- Troubleshooting signal strength problems
 - For example:* How can you increase signal strength without replacing hardware?
- Troubleshooting configuration problems
 - For example:* What configuration settings are most prone to error?
- Troubleshooting incompatibility problems
 - For example:* Which wireless hardware features should you review to confirm device compatibility?
- Troubleshooting incorrect channel problems
 - For example:* When should you modify channel settings on wireless devices?
- Troubleshooting latency problems
 - For example:* How does latency affect wireless network performance?
- Troubleshooting encryption type problems
 - For example:* What encryption settings must you configure in a client and an access point?
- Troubleshooting bounce problems
 - For example:* How does signal bounce affect wireless network efficiency?
- Troubleshooting SSID mismatch problems
 - For example:* What circumstances can cause clients to connect to the wrong access point?
- Troubleshooting incorrect switch placement problems
 - For example:* How much should you move an access point to improve network connectivity?

Troubleshooting interference problems

If the wireless devices are close enough together to be receiving a strong signal but are not, then you must start looking for sources of interference. Large metal objects, such as refrigerators or HVAC ducts, can block wireless radio signals, as can internal walls, especially if they are made of concrete or cinder block instead of drywall. Because radio signals can bounce off walls and other obstacles, moving access points and other devices even just a short distance can sometimes yield better connections.

True or false: Electrical devices, such as air conditioners or microwave ovens, can cause intermittent interference problems that may be difficult to troubleshoot.

Answer: *True.* Some equipment can generate electromagnetic interference that prevents wireless devices from communicating, but they only do so when the equipment is running, which can make the network faults maddeningly intermittent.

Discovering the offending device is a matter of trial and error, but when you find it, try to move the wireless devices and the electrical equipment farther apart.

EXAM TIP The Network+ exam objectives include several wireless network faults that can manifest themselves in roughly the same way. A device that can't connect to the network can be affected by interference, signal strength issues, channel problems, or other causes.

True or false: The more intervening walls, floors, and ceilings there are between a computer and an access point, the shorter the maximum distance between them.

Answer: *True*. Interference reduces the distance that wireless devices can transmit.

Troubleshooting signal strength problems

If a wireless device is receiving a weak signal, causing reduced connection speeds, the device might just be too far away from the access point. You can resolve this by moving the device or the access point so the two are closer together. You can also purchase a higher-gain antenna for the access point to boost the signal, or you can simply add additional (closer) access points to your network.

True or false: Signal strength can be weakened by intervening obstructions.

Answer: *True*. Walls and other barriers can interfere with the signals from wireless LAN devices, causing weakened signal strengths.

EXAM TIP For more information on the troubleshooting procedure recommended in the Network+ exam objectives, see "Objective 1.8: Given a scenario, implement the following network troubleshooting methodology."

True or false: Reduced signal strength can cause wireless LAN devices to operate slower.

Answer: *True*. When the signal strength is reduced, a wireless device reverts to a slower speed to compensate.

Cordless telephones and other devices that use the 2.4-GHz frequency band can interfere with the signals from wireless LAN devices, causing them to revert to slower speeds.

Troubleshooting configuration problems

Configuration errors account for a large number of wireless networking problems, and are usually the first thing the troubleshooter should check as a possible cause. In particular, mistyped text strings are common, such as SSID values or pre-shared keys that have unintended spaces or transposed characters. It is also common to find mismatched security configuration settings, such as systems configured to use different security protocols or encryption algorithms.

True or false: Wireless devices can have different names for the same settings.

Answer: *True*. There is no standard for the names of configuration settings used by wireless equipment manufacturers, and it is common for devices to have configuration settings with completely different names.

EXAM TIP Wireless devices have many configuration settings that can be possible sources of error, but the troubleshooting questions on the Network+ exam do not generally include the names of specific settings that might differ between hardware implementations.

True or false: Some computers can automatically turn their wireless transceivers off, preventing them from connecting to a network.

Answer: *True*. Laptop computers and other mobile devices often have power saving mechanisms that can turn off their wireless radio transceivers with no indication to the user. Some devices also have a physical switch for the radio that can be mistakenly turned off.

Troubleshooting incompatibility problems

Compatibility between wireless networking components is typically a matter of age. The older a component is, the less likely it is to be compatible with the latest equipment. Fortunately, many of the access points on the market support several or all of the IEEE 802.11 standards. It is always a good idea to purchase multi-standard access points, particularly when you are going to be using client computers with built-in wireless adapters, because you might not have a choice of which standards the computers support.

True or false: Within frequency compatibility, wireless network clients can automatically fallback to earlier versions of the IEEE 802.11 standard when connecting to an access point.

Answer: *True*. It is not necessary to configure wireless network clients to match the 802.11 standard of the access point. The two will automatically negotiate a common standard and connect at the best speed they have in common.

EXAM TIP Network+ exam questions can also address incompatibilities at the firmware level. Firmware updates for wireless devices are common, especially for access points. Incompatibilities among devices, even those made by the same manufacturer, can be the result of different firmware levels. Installing the latest firmware updates from the manufacturer's website can address such problems.

True or false: Incompatibility between security protocols can be a more serious problem to address than incompatibility between 802.11 standards.

Answer: *True*. Many wireless devices support multiple security protocols, but unlike 802.11 standards, you must choose one for an access point and configure all of your client devices to use it. If you have one client that only supports an outdated security protocol, then your choices are to use that protocol for all of your clients,

purchase a new network adapter for the outdated client, or add another access point specifically for outmoded clients.

Troubleshooting incorrect channel problems

Most wireless LAN clients select the channel used by the access point automatically. If you choose to configure the channel manually, make sure all of your devices are using the same one. If you have multiple access points, configure them to use non-overlapping channels.

Mismatched channels will prevent any connection at all, but even devices using the same channel can have trouble communicating, if there are too many other networks in the area also using the same channel.

True or false: Multiple wireless networks in the same general area, but with different SSIDs, cannot use the same channel without conflicting.

Answer: *False*. Multiple networks can use the same channel without conflicting, as long as they have different SSIDs, but too many devices in the same area can crowd the channel and slow down transmission speeds.

EXAM TIP Network+ exam questions regarding channel selection typically do not mention specific channels, but rather test the candidate's knowledge of the conditions under which it is necessary to change the channel.

True or false: If you suspect that an overcrowded channel is the reason for poor transmission speeds on your wireless network, all you have to do to test that theory is change the channel on the access point.

Answer: *True*. Because most wireless clients automatically tune to the correct channel for an SSID, you can simply change the channel on the access point, and all of the other devices will follow.

Troubleshooting latency problems

Latency is a measurement of the time it takes for signals to travel from one location to another on a network. Wireless networks are subject to highly-varying degrees of latency, because of all the different factors that can affect the signal transmissions. As a result, latency-sensitive applications, such as Voice over IP (VoIP) and streaming video, are more likely to experience problems on a wireless network than on a wired one. Addressing these problems is a matter of achieving the highest and most consistent transmission speed possible, using all the factors mentioned elsewhere in this objective.

True or false: Adding access points to an existing wireless network can reduce the latency experienced by the clients.

Answer: *True*. Splitting a wireless network's clients between two access points using different channels can reduce latency, because there is less crowding of each channel.

EXAM TIP The Network+ exam typically does not use the term latency in a specifically wireless context. You are more likely to see the term in questions on Quality of Service or traffic shaping, as covered in “Objective 4.6: Explain different methods and rationales for network performance optimization.”

Troubleshooting encryption type problems

One of the most common causes of wireless connection failures is mismatched encryption settings. You can check for this by temporarily turning off all encryption on your access point and your wireless devices. If they connect readily on an open network, then there’s a problem with your security configuration.

Make sure all of the devices are using the same encryption protocol with the same algorithm. (WPA2-AES and WPA2-TKIP are not the same thing.) If you are working with older devices, you might find that they do not support the latest encryption protocols. If the security protocol you have selected requires a pre-shared key, be sure that you spell the key in exactly the same way on each device.

True or false: If you are running WPA2 on your access point, and you have a computer that only supports WEP, your only options are to reconfigure the access point to use the weaker WEP protocol, install another access point for WEP clients, or replace the network adapter on the computer.

Answer: *True*. Wireless devices must be using the same encryption algorithms to communicate.

MORE INFO For more information on wireless security protocols, see “Objective 5.1: Given a scenario, implement appropriate wireless security measures.”

Troubleshooting bounce problems

Wireless network signals are capable of bouncing off of walls and other objects, which can cause them to arrive at a destination multiple times, or in a weaker state, or both. Access points typically have two or more antennae so they can receive multiple copies of the same signals. They then use digital signal processing to compare the copies and reconstruct the most accurate data possible. In some cases, however, signals can bounce off of multiple surfaces, causing them to arrive at widely different times.

True or false: It is difficult or impossible to predict how radio signals will bounce in a particular environment.

Answer: *True*. Radio signals can bounce off of surfaces at different angles and different strengths, resulting in a complex web of transmissions that reach wireless devices at different times and from different directions. Sometimes, achieving the best possible transmissions is a matter of trial and error.

Troubleshooting SSID mismatch problems

If you typed the SSID into a client configuration, make sure that you typed it exactly as it appears in the access point. This is particularly important if you have elected to suppress SSID broadcasting. In addition, if there are other access points in the area that are not broadcasting SSIDs, your systems might be trying to connect to the wrong one. This is called an SSID mismatch.

True or false: To prevent SSID mismatches from occurring, you can temporarily enable SSID broadcasting on your access points while you configure your wireless devices.

Answer: *True*. Enabling SSID broadcasting on your access points enables client devices to select the access point by name, rather than by typing the SSID. Once the devices are connected, you can disable the SSID broadcasts and the client devices will remain connected.

EXAM TIP The Network+ exam uses the acronym SSID without distinguishing between the two types: the Basic Service Set Identifier (BSSID) and the Extended Service Set Identifier (ESSID). Some hardware manufacturers use other terms however, such as network name, when referring to the SSID.

True or false: SSIDs are alphanumeric text strings that are case sensitive and can be up to 32 characters long.

Answer: *True*. When typing SSIDs, you must be sure to use the correct capitalization and distinguish between similar letters and numbers, such as O and 0 (oh and zero).

Troubleshooting incorrect switch placement problems

As mentioned elsewhere in this objective, achieving improved performance from a wireless network can be as easy as moving an access point a short distance away from an obstruction or a source of interference. The problem with this trial and error method is that improving performance for some clients can conceivably worsen performance for others.

Your first goal when adjusting the placement of an access point is to provide coverage to the entire area you need to service. Then, your next goal should be to provide the best possible service to each client.

EXAM TIP The Network+ exam objectives refer to “incorrect switch placement” with respect to wireless troubleshooting. Many wireless access points are actually multifunction devices that include switched ports, but it is not the placement of the switch that is critical to network performance; it is the placement of the access point.

True or false: To make any appreciable change in the performance of a wireless network, you must move the access point at least ten feet in one direction.

Answer: *False*. Because of the many factors that can affect radio signals, including structures, electromagnetic interference, signal bounce, and so on, moving an

access point just a few inches can have a noticeable effect on the performance of the network.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Under what conditions might it be necessary to change the default channel used by an access point?
2. Can weather conditions affect the performance of a wireless LAN?
3. After configuring a wireless access point to WPA2-TKIP security, what must you do on the client devices to make them compatible?
4. What are the two primary ways to increase the signal strength received at a wireless network client?
5. What would be the best way to address an interference problem caused by the presence of many other wireless LANs in the same office building?

Objective 2.5: Given a scenario, troubleshoot common router and switch problems

To some extent, routers and switches are subject to the same sources of error as ordinary computers. Both have complex configuration settings that administrators can inadvertently configure incorrectly, and routers have network-layer interfaces that require IP addresses, subnet masks, default gateways, and DNS addresses that are just as subject to error as those of a computer.

In this objective, Network+ exam candidates are expected to be familiar with some of the most common problems affecting routers and switches, and know how to resolve them as part of their standard troubleshooting methodology.

Exam need to know

- Troubleshoot switching loop problems
For example: How can you prevent switching loops from occurring?
- Troubleshoot bad cables and improper cable types
For example: What is the easiest way to test for a bad cable?
- Troubleshoot port configuration problems
For example: What port configuration settings are found in switches?
- Troubleshoot VLAN assignment problems
For example: What happens when a switch port is in the wrong VLAN?
- Troubleshoot Mismatched MTUs and MTU black holes
For example: What messages do MTU black holes prevent systems from receiving?
- Troubleshoot power failures
For example: Are routers and switches sensitive to power outages?

- Troubleshoot bad or missing routes
For example: What happens when a router contains incorrect routing table information?
- Troubleshoot bad modules
For example: How do you troubleshoot fiber optic modules, such as SFPs and GBICs?
- Troubleshoot wrong subnet masks
For example: How does a wrong subnet mask affect router performance?
- Troubleshoot wrong gateways
For example: Do all routers need default gateway addresses?
- Troubleshoot duplicate IP addresses
For example: How do systems test for duplicate IP addresses?
- Troubleshoot wrong DNS addresses
For example: What happens when the DNS server address supplied by a router is incorrect?

Troubleshoot switching loop problems

One problem common to switches is known as a switching loop. A switching loop is a condition that can occur when you have redundant switches on your network to provide multiple paths between destinations. Administrators often do this to introduce a measure of fault tolerance. If one switch fails, packets can take an alternative route through the network.

The problem results when there are multiple connections between two switches. Each switch receives packets from the other switch and forwards them back, creating a loop that circulates packets endlessly. Switching loops can also be the result of misconfiguration of the switch ports or accidental looping of a cable from one switch port to another.

In the case of broadcasts, the switches forward the packets out through all of their ports, causing the number of looping packets to increase exponentially, flooding the network. This is called a broadcast storm.

When a switching loop occurs, the network is immediately flooded with packets—to the point at which no useful traffic gets through. The solution to a switching loop is to leave the redundant switches in place but use a special protocol called the Spanning Tree Protocol (STP) on the switches. STP allows only one path for packets through the switching infrastructure, activating others only when a switch fails. In the case of a misconfigured switch or a looped cable, STP blocks the offending ports to prevent the looping traffic, but it cannot correct the condition that caused the problem in the first place.

True or false: Switched packets have no way of identifying looping traffic.

Answer: *True.* Routers can identify looping traffic because the network layer protocol has a Time To Live field that specifies the number of routers that have processed each packet. Switches are data-link layer devices and data-link layer protocols have no such field.

EXAM TIP Network+ exam questions typically do not go into great detail about switching loops and their causes. Candidates should, however, be aware that the Spanning Tree Protocol (STP) is the designated remedy for the condition.

True or false: Loops can occur on bridges as well as switches.

Answer: *True*. Switches are essentially multiport bridges, so having multiple bridges on a network can result in the same phenomenon, which is called a bridging loop.

Troubleshoot cable problems

Cables are the fundamental fabric of a data network. When they do not function properly, switches and routers do not work, nor do computers and other devices. Cable failures manifest themselves in two possible ways: either there is no signal at all, or the signal is intermittent, working properly some of the time and then failing again. Intermittent failures are typically due to a faulty connection at one end of the cable. In nearly every case, the solution is the same: replace the cable with one that you know works.

In the case of an internal cable installation, there are three cables that can be at fault: the internal cable run in the walls or ceilings or one of the two patch cables connecting the computer to the wall plate or the patch panel to a switch or router port.

Internal cable runs should be tested at the time of their installation, and should only go bad if someone working in the wall or ceiling spaces disturbs them. The easiest way to determine if the internal run is at fault is to switch the patch cable to a run that you know works properly. If the problem goes away, then you have a bad run that should probably be repaired by the contractor who did the installation.

If one of the patch cables is at fault, replacing them one at a time with cables that you know are good should enable you to isolate the fault. Even new, factory-made patch cables can malfunction, so the replacements should be cables that you know are functional, not just brand new ones.

True or false: Link lights can be lit on both sides of a cable, even when the cable is faulty and not functioning.

Answer: *True*. A cable may be able to carry the link pulse signals that illuminate the link lights on the patch panel and the network interface adapter, while failing to carry data effectively. This can result from a faulty crimp at one end that leaves one or more of the wires disconnected.

EXAM TIP The Network+ exam objectives list bad cables as part of router and switch troubleshooting, but a bad cable will frequently manifest itself first as a computer problem, not a router or switch problem.

True or false: Cable faults can sometimes cause reduced transmission speeds, rather than outright connection failures.

Answer: *True*. Intermittent connections or overlong cable runs can cause the throughput of a cable to be lower than expected. A faulty connection might be cycling so rapidly that it gives the appearance of being slow when it is actually intermittent, or the systems might be dropping packets when the signals weaken beyond recognition.

Troubleshoot port configuration problems

Each port in a switch connects to a different device, so each one has its own independent settings for port speed and duplex operation. The small switches intended for residential and small office use rely on auto-negotiation for each port. Connecting a device causes it to exchange capabilities with the switch, after which the two decide on the fastest possible configuration.

On larger switches, you can opt to configure each port manually. The only rule is that both the connected device and the switch port must be configured to use the same speed and the same duplex mode. If the device and the switch port are set for different speeds, there will be no connection at all. If one device is set for half duplex communication and one for full duplex, there will be a connection, but throughput will suffer.

True or false: When the switch port is configured to use a different duplex mode than the connected device, the transmission rate suffers because of collisions.

Answer: *True*. In full duplex communications, there is no media access control and no collisions; both devices can transmit whenever they want to and receive data at the same time. In half-duplex mode, devices can transmit or receive, but not both simultaneously. Therefore, on a mixed mode connection, the full duplex device can transmit constantly, causing a high collision rate with the half duplex device and forcing the retransmission of many packets and reduced throughput.

EXAM TIP The port configuration issues covered in the Network+ exam can also include VLAN configuration settings, as discussed later in this objective.

True or false: The auto-configuration of port speed and duplex mode is an Ethernet function, and the resulting performance of a connection with mismatched settings can vary depending on the Ethernet technologies involved.

Answer: *True*. The port speed in a switch refers to the speeds of the Ethernet standards the device supports, such as 10, 100, 1,000, and possibly 10,000 Mbps. The duplex mode refers to whether that particular Ethernet connection supports half or full duplex communication. A connection with one device running at 100 Mbps and the other at 1,000 Mbps will transfer no data at all, because the two are using different signaling schemes; they speak different languages. Connections with other combinations of speeds might not suffer the same total failure.

MORE INFO For more information on Ethernet speeds and compatibilities, see “Objective 3.7: Compare and contrast different LAN technologies.”

Troubleshoot VLAN assignment problems

On a large enterprise network, switches typically contain multiple virtual LANs (VLANs) and the switch ports are assigned to specific VLANs. VLANs are logical divisions within the switch that function just like physical subnets on a network. Each VLAN has its own network IP address, and the devices connected to that VLAN's ports must have IP addresses on that network.

If a port on a switch is assigned to the wrong VLAN, then a computer with a static IP address that you connect to that port will not be able to communicate with the network. The effect is the same as if you connected a computer with an address on one network to a subnet using another network address. To troubleshoot these problems, the only solution is to examine the VLAN configuration of the switch and make sure that the port assignments are correct. A properly designed network will have documentation of the VLAN configuration of its switches, making this task easier.

True or false: An incorrectly configured VLAN assignment on a switch port will result in a complete failure of the connected device to communicate on the network.

Answer: *True*. VLAN port assignments are an either/or proposition. Either the connected device works or it does not; there is no middle ground.

EXAM TIP The configuration interfaces for switches are not standardized, so different manufacturers each have their own designs. As a result, the questions on switch configuration in the Network+ exam are general in nature and do not require candidates to know interface controls or specific procedures.

True or false: VLAN assignment problems can also result from a trunking failure.

Answer: *True*. Trunking is the process by which the devices in a VLAN on one switch communicate with the devices in the same VLAN on another switch. If the trunk ports on the two switches are not configured properly, then communication between the switches can be interrupted, conceivably producing the same type of communication failure.

MORE INFO For more information on VLANs, see “Objective 1.4: Explain the purpose and properties of routing and switching.”

Troubleshoot MTU problems

The size of the datagrams that IPv4 creates is based on the maximum transfer unit (MTU) of the network to which the system is attached. If, on the path to its destination, a datagram encounters a network with a smaller MTU, the router providing access to that network must split the datagrams into fragments and transmit them separately. This is called an *MTU mismatch*.

Fragmentation is an inherently inefficient process, particularly for the intermediate routers doing the fragmenting. To avoid it, a technique was created called Path MTU Discovery that enables an end system to determine the MTU for an

entire route through an internetwork, not just the first hop. This enables the system to create datagrams small enough to travel all the way to their destinations without requiring routers to fragment them.

Path MTU Discovery is simply a series of Internet Control Message Protocol (ICMP) Echo Request messages, the same ones used by the Ping utility. These messages have the Do Not Fragment flag set in the IP header, and they specify a message size equal to the MTU of the end system's network. If the messages do not get through to the destination without being fragmented, the system reduces the message size and tries again. This process is repeated until the system determines the largest packet size that can make it to the destination without being fragmented.

One of the problems with this method is that many routers contain firewalls that block ICMP messages by default. This essentially disables the Path MTU Discovery mechanism and creates what is known as an MTU black hole. If your network users experience severe connection problems when accessing the Internet but have no trouble accessing destinations on the local Ethernet network, an MTU black hole is one possible cause. Check the firewall in your Internet access router to see if it is set to block ICMP Echo Request and Echo Reply traffic. Disabling this filter might alleviate the problem.

True or false: MTU mismatches do not occur on IPv6 networks.

Answer: *True*. On IPv6 networks, there is no fragmentation by intermediate systems, such as routers, only by end systems. End systems perform a path MTU discovery and transmit packets of the appropriate size.

EXAM TIP The Network+ exam objective mistakenly refers to MUT black holes; the correct name for the phenomenon is an MTU black hole.

True or false: You can test for yourself whether a router is returning the appropriate ICMP error messages by using the Ping utility.

Answer: *True*. Using Ping with the `-I` parameter, to send packets of a particular size, and the `-f` parameter, to activate the Do Not Fragment bit in the IP header, you should be able to send a packet large enough to force the router to generate a Destination Unreachable error message. If you cannot do so, then the router is likely filtering out the ICMP messages generated by Ping.

Troubleshoot power failures

Routers and switches are as critical to the operation of a network as are servers. Therefore, administrators should take as much care with their power supplies as they would with a server. Routers and switches should have surge protection and uninterruptible power supplies, to ensure that they remain operational at all times.

True or false: Data centers often have centralized power protection.

Answer: *True*. Data centers for large, enterprise networks often have equipment that provides conditioned power and backup power supplies for all of the routers, switches, and other equipment they contain.

EXAM TIP When answering troubleshooting questions, candidates for the Network+ exam should be careful not to neglect the most obvious causes of a problem. Administrators intent on finding a highly technical solution to a router or switch problem can sometimes miss the incredibly obvious: that the power cord has accidentally been unplugged.

Troubleshoot bad or missing routes

Another possible source of router problems is the router's own routing table, particularly if administrators are creating static routes manually. If you connect to the router by using a terminal, you must use a command-line program to create routing table entries. The program is similar to route or Route.exe, with a slightly different, but no less cryptic, syntax.

The slightest typing mistake can result in an incorrect route in the table. Even when you are using a web-based interface, typographical errors are easy to make. Administrators creating multiple routes can also easily skip one, leaving a network unrepresented in the routing table. Convergence problems with dynamic routing protocols that prevent updates from reaching all of the routers on the network can also result in missing routes.

This type of error can be extremely difficult to troubleshoot, because it affects only the traffic going to the destination specified in the incorrect or missing route. The router functions normally except for the traffic going to one network, which ends up being forwarded to the wrong place.

True or false: Networks that use dynamic routing are more prone to having routing table errors than networks that use static routing.

Answer: *False*. Dynamic routing automatically creates routing table entries, thereby eliminating the typing errors and omissions that can easily occur with static routing.

EXAM TIP For the Network+ exam, be prepared to examine routing table entries and spot those that are improperly configured.

True or false: Networks that use link state routing protocols are more prone to having routing table errors than networks that use distance vector routing protocols

Answer: *False*. As far as routing table entries are concerned, the only difference between link state and distance vector routing protocols is the source of the Metric value. The chances of routing table errors occurring are no greater with link state than with distance vector protocols.

Troubleshoot bad modules

The modular nature of most contemporary business-class routers and switches is another possible source of errors that are difficult to detect. The proliferation of different fiber optic cable types has led router and switch manufacturers to create replaceable modules, using standards such as Gigabit Interface Converter (GBIC) and Small Form-Factor Pluggable (SFP). These modules plug into a router or switch and make it easy for administrators to connect networks by using various types of cables and connectors.

With any modular architecture such as this, it can be easy to purchase the wrong module. You need to be sure you have the right form factor for your router or switch, the right module for the fiber optic cable used on your network, and the right connector for your cables. There are many permutations, and the modules can be difficult to tell apart.

Assuming that you have purchased the correct module for your installation, a connection failure could be caused by a faulty module, even while the rest of the router is working. The router or switch should automatically recognize a new module when you insert it. Always check the router's interface to make sure that the module has been recognized.

If the module is not recognized by the router or switch, or it fails to establish a connection, the easiest way to troubleshoot it is to replace the module with one that you know works. If a good module also fails, then it is time to look at the router or switch itself, if the module was not recognized; or the fiber optic cable, if the module seems to work but there is still no connection.

True or false: You must always power down a router or switch when installing a new module.

Answer: *False*. Some routers and switches support hot swapping, which enables you to replace a module without powering down the device. Be sure that your device supports this feature before attempting it, though.

True or false: If you are unsure whether a fiber optic module is faulty or the problem is in the router or switch, you can try plugging the module into a different slot.

Answer: *True*. Replacing hardware is one of the most basic troubleshooting methods. If a router or switch fails to recognize a module, plugging it into a different slot is one of the quickest and easiest ways to isolate the problem.

Troubleshoot incorrect subnet masks

When you configure the network interfaces in a router, you typically type the IPv4 addresses and other parameters manually, which means that it is very easy to make a simple typing mistake.

When you mistype the subnet mask in a TCP/IP configuration, there is in nearly all cases no communication through the interface whatsoever. The subnet specifies which bits of the IP address function as the network identifier and which as the host

identifier. If this value is wrong, then the system is placed on a different subnet all alone, and no communication with the other subnets is possible.

EXAM TIP In Network+ exam questions that require you to troubleshoot a misconfigured interface, you can sometimes tell which parameter is incorrect by the amount of functionality in the interface.

True or false: The subnet mask of a router interface must always be the same as the masks of the devices connected to that interface.

Answer: *True*. All of the devices on an IP subnet must have the same subnet mask and the same network identifier, or they cannot communicate.

Troubleshoot incorrect gateways

An incorrect default gateway address in a TCP/IP configuration prevents the device from forwarding packets to other networks not represented in its routing table. The device can communicate with other systems on the local network, however.

True or false: A router with an incorrect default gateway address can be more difficult to troubleshoot than a regular computer with the same fault.

Answer: *True*. A router interface can have a missing or incorrect default gateway address, preventing it from communicating with networks not found in its routing table, but it can still send traffic to the networks that are found in the table. Depending on how the routing table is populated, the effect of the incorrect gateway might be hardly noticeable.

Troubleshoot duplicate IP addresses

When an IP address is duplicated on a network, the systems using that address are in conflict. They will both respond to Address Resolution Protocol (ARP) messages searching for that IP address, resulting in two MAC addresses associated with one IP address. The ongoing results are unpredictable; traffic might be forwarded to one system using the address or the other.

Because you typically configure the IP addresses of routers manually, it can be relatively easy to duplicate another address on the network. To prevent this from occurring, you should Ping the address before finalizing the configuration on the router.

True or false: DHCP servers cannot assign duplicate IP addresses on a network.

Answer: *True*. Addresses assigned dynamically by DHCP are almost never duplicated, because DHCP servers perform an ARP test looking for duplicates before they finalize an IP address assignment.

EXAM TIP Because so many systems test for IP address duplication, the occurrence of this fault is relatively rare, and so is its appearance on the Network+ exam.

True or false: Computers running Windows check the network for IP address duplication as they start up.

Answer: *True*. Many operating systems, including Windows, perform an ARP test as they start up, and if they detect a system using the same IP address, will disable the interface until the user changes the address.

Troubleshoot incorrect DNS addresses

DNS is a critical service for nearly all networks. Any user accessing the Internet needs DNS to resolve server names into IP addresses. Clients on an Active Directory Domain Services network need DNS to locate domain controllers.

When the DNS name resolution process fails, there are two possible problems: either the DNS server is malfunctioning or the client computer is configured with an incorrect DNS server address. When a router has an incorrect DNS server address, it can prevent all of the computers on the network from resolving names.

Pinging the DNS server address can tell you if the server is running, and the Nslookup utility can tell you if the DNS service is running on the server.

True or false: Broadband routers for the SOHO market provide all of the clients on the network with DNS server addresses.

Answer: *True*. Broadband routers receive an IP address and other TCP/IP configuration settings, including DNS server addresses, from the Internet Service Provider's DHCP server. If the DNS addresses are incorrect, or if the DNS servers are not functioning, Internet access for the entire network is interrupted.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Why are only computers with static IP addresses subject to communication failure when plugged into a switch port with an incorrect VLAN configuration?
2. A newly installed router, on which the administrator has yet to add static routes, can access the computers on the local network but not the computers on other networks. What configuration parameter is likely to have been misconfigured?
3. Your users can't access the Internet, and you suspect that it is only the DNS server address that is wrong. How can you prove that the Internet is accessible, and it is only the DNS server that is the problem?
4. How does ARP determine whether there is a device on the network using a particular IP address?

Objective 2.6: Given a set of requirements, plan and implement a basic SOHO network

A small office/home office (SOHO) network installation is generally characterized as supporting from one to ten users, and might or might not require access to data or applications located at a larger, remote facility. For this objective, candidates must

consider the differences between a SOHO network installation and the larger, more complex networks of larger sites.

Exam need to know

- Plan a list of requirements
 - For example:* What tasks must the SOHO network users perform?
- Plan for cable lengths
 - For example:* How do you determine the proper length patch cable to buy?
- Plan device types
 - For example:* What devices are needed for a SOHO network?
- Plan for environment limitations
 - For example:* What effect does extreme humidity have on a computer?
- Plan for equipment limitations
 - For example:* What are the limitations of a SOHO switch?
- Plan for compatibility requirements
 - For example:* Why must the equipment at a SOHO be compatible with that at the headquarters office?

Plan a list of requirements

The first step in planning and implementing a SOHO network is to develop a list of what the users will require from the network. A SOHO network can range from an actual office space consisting of several rooms to an informal work space in a user's home.

A list of requirements for any network must begin with connectivity, both within the site and with remote sites. The need for a network means, at the very least, that there will be one or more computers connected, probably with access to a printer and an Internet connection. Beyond that, the users might also need access to corporate resources, such as databases and applications, which could be available through a virtual private network (VPN) connection or a web-based application.

True or false: SOHO networks typically have requirements that differ substantially from those of larger networks.

Answer: *False.* SOHO networks often have many of the same requirements as larger networks; they differ only in scale and in the location of resources.

EXAM TIP The Network+ objective for planning and implementing a SOHO network is new to the N10-005 version of the exam, and exists primarily to collect the information on SOHO-specific hardware and procedures in one place.

True or false: When planning requirements for a SOHO network, you don't have to consider many of the standardized requirements for large, corporate networks, such as security, data backup, and disaster recovery.

Answer: *False*. Security, data protection, and disaster recovery are just as important factors on a SOHO network as on a larger installation. The specifications for these requirements might not be as rigidly standardized as they are in the corporate world, but they are no less important to the home or small business user.

Plan for cable lengths

Wireless networks are an increasingly common solution for SOHO networks, because of their ease of installation, but there are still many small networks that use standard Ethernet cabling. Because they are generally located in smaller areas, it is common to find wired SOHO networks that consist of externally-installed, unshielded twisted pair (UTP) patch cables, rather than the internal installations that run cables through walls and ceilings.

All of the Ethernet standards include a UTP cable specification with a 100-meter maximum segment length, which is far more than most SOHO networks need. When planning an external cable installation, you generally choose patch cables of various lengths, with the RJ-45 connectors already on them, rather than pull bulk cable and apply the connectors yourself. When you do this, you must be careful to measure the actual paths your cables will take, so that you can purchase patch cables that are long enough.

For example, an informal SOHO network might consist of several computers in a single room, with cables running around the perimeter. To purchase patch cables of the correct length, you must consider whether you will run the cables over doorways, around or behind obstacles, and so forth. The computers might seem to be only a few feet away from each other, but the path of the cable might be significantly longer.

MORE INFO For more information on network cabling, see “Objective 3.1: Categorize standard media types and associated properties.”

True or false: The 100-meter cable length limitation for UTP Ethernet networks refers to the distance from each computer to the switch, not the distance between computers.

Answer: *True*. The switch on an Ethernet network functions as a repeater, amplifying signals as it forwards them, so each computer can have a cable up to 100 meters long connecting it to the switch. SOHO networks rarely span distances this large, however, and it can be difficult to find prefabricated patch cables any longer than 50 or 100 feet.

EXAM TIP The Network+ objectives list several topics for planning and implementing SOHO networks that are just as applicable to larger networks. Cable lengths, equipment limitations, and compatibility requirements do not change with the size of the network. Therefore, much of the material in this book is applicable to SOHO, as well as larger, enterprise networks.

True or false: SOHO networks often use fiber optic cables as an alternative to UTP.

Answer: *False*. Fiber optic cabling is almost never found on SOHO networks, because of its much greater expense and more difficult installation.

Plan device types

The computers, network interface adapters, and cables that you use to implement a SOHO network are the same as those used on larger networks. However, there is a large market for connectivity devices, such as switches and routers, designed specifically for home and small office installations.

For connectivity purposes, a SOHO network typically requires an Ethernet switch, to connect the computers and other devices together; a broadband modem, to provide Internet access; and a router, to connect the local network to the ISP's network. Manufacturers have capitalized on the almost universal standardization of these requirements, and there are now multifunction devices available that provide some or all of these features and more.

The typical multifunction device, often called a broadband router, contains several switched ports for local device connections, a wide area network (WAN) router to connect the switch to the broadband connection, and either an internal broadband modem or a connection for an external one. These devices may also include other functions, such as wireless access points, DHCP servers, VPNs, and firewalls.

True or false: In addition to multifunction router/switch devices, there are individual SOHO connectivity products on the market as well.

Answer: True. If you need more switch ports than are provided on your broadband router, there are standalone switches available in various sizes that are designed for the SOHO market. Unlike the elaborate rack-mounted switches found in enterprise data centers, these are small, relatively inexpensive devices that you simply plug into your network. This type of switch has none of the advanced features discussed elsewhere in this book, such as VLANs and individual port configuration.

EXAM TIP The Network+ objectives now include coverage of SOHO networks, primarily to distinguish the smaller, less expensive switches and routers on the market from the elaborate, expensive, rack-mounted devices used by large, enterprise networks.

True or false: There are specialized printers designed for SOHO installations.

Answer: *True*. Multifunction printer devices intended for the SOHO market can also provide scanning, copying, and faxing services, as well as network printing.

Plan for environmental limitations

SOHO installations are typically not located in buildings with elaborate environmental support systems designed to maintain networking equipment in optimal conditions. A corporate data center will generally have raised floors for cable access, specialized air conditioning to maintain temperature and humidity levels, and power conditioning to prevent surges—all things that SOHO installations usually lack.

When planning a SOHO installation, you should do your best to maintain adequate environmental conditions at the site. This might include additional air conditioning and individual surge protectors and uninterruptible power supplies for sensitive devices.

True or false: When designing a home office network, installing servers, switches, and routers in a locked closet for security is the recommended practice.

Answer: *False*. While installing the networking equipment in a closet will secure it, the closed environment will also capture heat and humidity that could shorten the life of the devices.

Plan for equipment limitations

The networking devices targeted at the SOHO market are smaller than enterprise networking products, less expensive, and they have fewer advanced features. SOHO products typically lack fault tolerance features you might find in higher-end equipment, such as redundant components, as well as advanced management functions, such as support for the Simple Network Management Protocol (SNMP). In many cases, SOHO devices are also not designed to operate at the speeds or traffic levels that enterprise equipment can handle easily.

For the average SOHO installation, most of these limitations are perfectly acceptable, especially when they help to keep the prices of the products low.

True or false: Switches designed for the SOHO market typically have no configuration interface.

Answer: *True*. SOHO switches typically do not have VLAN or layer 3 functions that require configuration, and the basic function of a switch requires no configuration, either. This is because the device is capable of compiling a table of MAC addresses from the incoming traffic.

EXAM TIP Candidates for the Network+ exam should be familiar with the differences between devices intended for SOHO installations and those for large, enterprise networks.

True or false: SOHO networks rarely use routers for any purpose other than WAN connections.

Answer: *True*. SOHO networks consist of a single subnet, so they have no use for LAN-to-LAN routers. SOHO routers are typically designed to connect an Ethernet LAN to a WAN provided by a broadband ISP.

Plan for compatibility requirements

When discussing compatibility requirements for a SOHO network, there are two distinct contexts. There is compatibility with the other devices at the site, and there is compatibility with the network infrastructure at the home office, if there is one.

When designing a SOHO network from scratch, selecting compatible components is critical. You probably don't have the purchasing infrastructure available in a large corporation, so returning and replacing an incompatible component can be a lengthy process. The compatibility issues of Ethernet components and wireless LAN devices are no different than they are for a large network, however.

When the SOHO network is a satellite or branch office of a large organization, you probably have little or no choice over the selection of equipment. Organizations with offices at many locations usually try to keep the hardware, software, and configurations uniform, to simplify purchasing, asset management, and technical support.

True or false: When planning a SOHO network, the connectivity hardware must be compatible, but you can select any operating system and version for the client computers.

Answer: *False*. Compatibility issues apply to operating systems as well. Most of the major operating systems can interact with each other, but for support purposes, administrators often require users to run specific operating system versions.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. How do the Ethernet cable length limitations for a SOHO network differ from those of a large, enterprise network?
2. How does a typical SOHO router differ from a router intended for a large, enterprise network?
3. Why don't SOHO networks need switches with VLAN capabilities?

Answers

This section contains the answers to the "Can you answer these questions?" sections in this chapter.

Objective 2.1: Given a scenario, install and configure routers and switches

1. Computers on the Internet never communicate directly with the client computers on the private network; they only communicate with the NAT router. This prevents Internet intruders from learning the IP addresses of the client computers and sending packets to them. Also, private IP addresses are unroutable, so Internet-based intruders cannot initiate communication with them directly.
2. No. You can plug the network cable into a splitter that separates the power from the data and provides a standard plug for a device's power cord.
3. An individual with access to a mirrored port in a switch can capture network packets that might contain confidential information.

4. Trunking enables VLANs to have memberships that span multiple switches.
5. QoS is more important now than it was five years ago because networks tend to carry more data types that require a consistent supply of bandwidth, such as voice and video.

Objective 2.2: Given a scenario, install and configure a wireless network

1. Using 802.11n with the 5-GHz band will likely provide the best performance, because the band is less crowded and provides more non-overlapping channels.
2. MIMO increases transmission speed by multiplexing several signals using separate antennae.
3. To connect to an access point that does not broadcast its SSID, you must manually specify the SSID in each device's network configuration.
4. Each access point must use a different channel so that a wireless device cannot be connected to two access points at the same time. By varying the channels among the access points, clients can wander into the area of a different access point, and the wireless adapter will compensate by changing to the appropriate channel.

Objective 2.3: Explain the purpose and properties of DHCP

1. Assigning static addresses using DHCP enables administrators to keep track of all address assignments in one place and prevents the possibility of IP address duplication.
2. You configure the Router option at the scope level because each scope services a different subnet and each subnet will have a different default gateway address.
3. Microsoft DHCP Server uses dynamic allocation by default. To use automatic allocation, you must configure the server to issue leases that never expire.
4. When administrators create a reservation, they must supply the MAC address of the device they want to receive the IP address. This enables the DHCP server to identify the device that should receive the address assignment.
5. DHCP servers use the DHCPINFORM message to carry options to clients that already have an IP address assigned to them.

Objective 2.4: Given a scenario, troubleshoot common wireless problems

1. If you are using multiple access points on a single network, all with the same SSID, then you must change the default channel on each access point to a unique, non-overlapping value. You might also change the channel of an access point when there are other networks in the vicinity using the same channel.

2. Yes, in addition to many other factors, changing weather conditions can have a profound effect on wireless LAN performance.
3. You must configure the client computers to use WPA2-TKIP as well. Unlike the connection speed and 802.11 standard, which clients automatically change to match that of the access point, you must manually configure the security settings.
4. To increase signal strength, you can move the client and the access point closer to each other or you can install a better antenna on the access point.
5. Your best chance of evading the interference from other wireless LANs would be if you installed a network that used the 5-GHz frequency band, instead of the more popular 2.4-GHz equipment.

Objective 2.5: Given a scenario, troubleshoot common router and switch problems

1. Computers with static IP addresses will be on the wrong subnet if they are connected to the incorrect VLAN. A computer that uses dynamic IP addressing will retrieve an IP address from the DHCP server on the subnet to which it is connected, even if it is the wrong subnet.
2. The default gateway address is likely to have been misconfigured.
3. If you can access an Internet server using its IP address instead of its name, then you know that the Internet is accessible and that the DNS server is malfunctioning.
4. ARP generates broadcast messages containing the IP address. Any system on the network possessing that address will reply with a message containing its MAC address.

Objective 2.6: Given a set of requirements, plan and implement a basic SOHO network

1. The cable length limitations for Ethernet networks of any size are all the same.
2. A SOHO router is typically a multifunction device that provides the network with a WAN connection for Internet access, as well as switched ports, a DHCP server, a firewall, and possibly other services as well.
3. SOHO networks have only one to ten users, and therefore require only a single subnet.

CHAPTER 3

Network Media and Topologies

The Network Media and Topologies domain accounts for approximately 17% of the CompTIA Network+ exam, but it contains many of the critical areas that CompTIA uses to generate its questions. All of the objectives in this domain are concerned with the physical and data-link layers of the OSI reference model, and most concentrate on the physical elements of the network, specifically the cables and other media that you use to connect computers together.

This chapter covers the following objectives:

- Objective 3.1: Categorize standard media types and associated properties
- Objective 3.2: Categorize standard connector types based on network media
- Objective 3.3: Compare and contrast different wireless standards
- Objective 3.4: Categorize WAN technology types and properties
- Objective 3.5: Describe different network topologies
- Objective 3.6: Given a scenario, troubleshoot common physical connectivity problems
- Objective 3.7: Compare and contrast different LAN technologies
- Objective 3.8: Identify components of wiring distribution

Objective 3.1: Categorize standard media types and associated properties

For this exam objective, you must be familiar with the various types of cable used to construct local area networks (LANs). Candidates for the Network+ exam must be familiar with the basic cable types, the different grades within each type, and the hardware used to make connections between cables and network devices. The exam covers many different cable types, some of which are almost never used in new network construction, but which you might encounter in a legacy installation.

Exam need to know

- Categorize fiber media

For example: Do you know that there are two types of fiber optic cable: multimode and singlemode?

- Categorize copper media
For example: Do you know that there are several types of copper cable used in networking, including UTP and STP, which are available in CAT3, CAT5, CAT5e, CAT6, and CAT6a varieties; and coaxial; and that twisted pair cables can be wired in crossover, T1 crossover, and straight-through fashion?
- Categorize plenum vs. non-plenum cables
For example: How do cables rated for use in plenums differ from non-plenum cables?
- Categorize media converters
For example: Do you know that there are media converters available for converting singlemode fiber to Ethernet, multimode fiber to Ethernet, fiber to coaxial, and singlemode to multimode fiber?
- Categorize cable distance and speed limitations
For example: What are the distance and speed limitations of the various fiber and copper cable types?
- Categorize broadband over powerline networking
For example: How can you use power lines for data networking?

Categorize fiber media

Fiber optic cable is completely different from copper and avoids nearly all of copper's shortcomings. Instead of transmitting electrical voltages over copper conductors, fiber optic cables transmit pulses of light through a filament of plastic or glass. The core filament is surrounded by a reflective layer called the cladding. The cladding is surrounded by a protective layer of woven fibers, a plastic spacer, and an outer sheath, as shown in Figure 3-1.

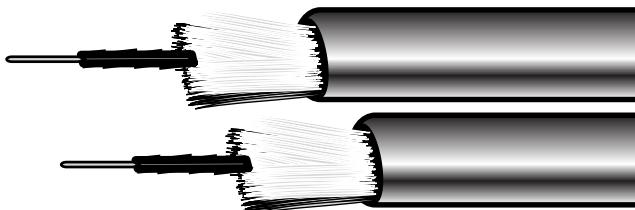


FIGURE 3-1 Fiber optic cables.

Unlike electrical voltages, light pulses are completely unaffected by electromagnetic interference (EMI), which means that you can install fiber optic cables near light fixtures or heavy machinery, or in other environments in which copper would be problematic. Fiber optic cable is also much less prone to attenuation than any copper cable, can span distances as long as 120 kilometers, and is also immune to outdoor conditions, which makes it ideal for installations that span long distances or connect campus buildings together.

There are two primary types of fiber optic cable: singlemode and multimode, which differ in size and in the means used to generate the light pulses they carry, as shown in Table 3-1.

TABLE 3-1 Fiber optic cable characteristics.

CABLE TYPE	CORE/CLADDING DIAMETER	LIGHT SOURCE
Singlemode	8.3/125 μm	1,310 or 1,550 nm Laser
Multimode	62.5/125 μm	850 or 1,300 nm LED

True or false: Fiber optic cables are much more expensive than copper, and their installation is much more difficult.

Answer: *True*. Fiber optic cables are comparatively rare on LANs, because not only are the raw materials far more expensive than those for copper cables, the installation process is completely different, requiring special tools and special skills.

EXAM TIP Candidates for the network+ exam should be aware that the primary advantages of fiber optic cable over copper-based twisted pair are its resistance to electromagnetic interference and its ability to span longer distances.

True or false: Multimode fiber optic cable is more common on LAN installations than singlemode.

Answer: *True*. Multimode is more common, largely because the cable is less expensive and easier to install. Multimode cable has a smaller bend radius, which means that you can bend it more sharply around corners. Because of its expense and relative inflexibility, administrators rarely choose singlemode fiber for LANs, preferring it for long, straight wide area network (WAN) runs instead.

Categorize copper media

There are two basic types of copper cable used in data networking: twisted pair and coaxial.

Unshielded twisted pair cable

A twisted pair cable is a cluster of thin copper wires, with each wire having its own sheath of insulation. Individual pairs of insulated wires are twisted together, usually gathered in groups of four pairs, for a total of eight wires, and are encased in another insulating sheath, as shown in Figure 3-2.

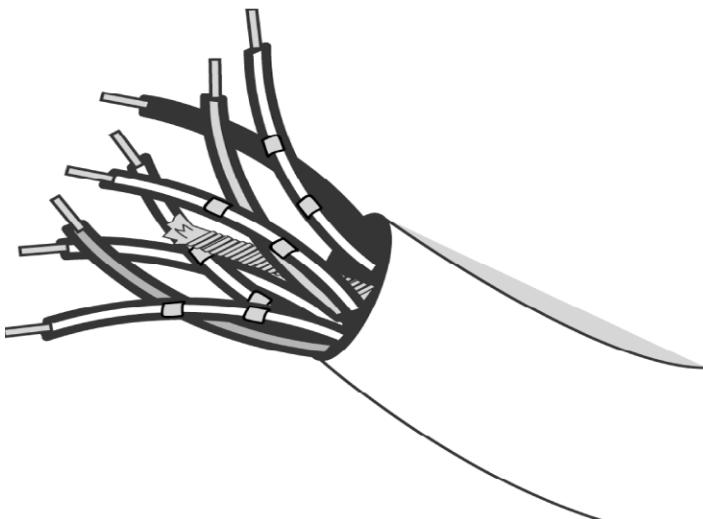


FIGURE 3-2 A twisted pair cable.

Unshielded twisted pair (UTP) cable is the standard medium for copper-based Ethernet LANs. A typical UTP cable contains four wire pairs within a sheath approximately 0.21 to 0.24 inches in outside diameter.

UTP cable comes in a variety of different grades, called categories in the cabling standards published by the Telecommunications Industry Association/Electronic Industries Alliance (TIA/EIA). The categories define the signal frequencies that the various cable types support, along with other characteristics, such as resistance to certain types of interference. The higher the category number, the higher the cable quality and, not surprisingly, the price. The cable categories that administrators are likely to encounter are as follows:

- **Category 3 (CAT3)** Long the standard for telephone communications, CAT3 cables were used by the first UTP-based Ethernet networks (called 10Base-T). CAT3 cables supported frequencies up to 16 megahertz (MHz). Insufficient for any of the faster Ethernet types, Cat3 is no longer supported for new installations.
- **Category 5 (CAT5)** Designed for 100Base-TX Fast Ethernet networks and supporting frequencies up to 100 MHz, CAT5 cabling was dropped from the latest version of the TIA/EIA cabling standards.
- **Category 5e (CAT5e)** Still rated for frequencies up to 100 MHz, CAT5e cable is designed to support full duplex transmissions over all four wire pairs, as on 1000Base-T Gigabit Ethernet networks.
- **Category 6 (CAT6)** Designed to support frequencies of up to 250 MHz, CAT6 cables easily handle 1000Base-T Gigabit Ethernet traffic and, with special installation considerations, 10Gbase-T.

- **Augmented Category 6 (CAT6A)** Created for 10Gbase-T installations with cable segments up to 100 meters long, CAT6a supports frequencies up to 500 MHz and was added to the most recent version of the TIA/EIA standards in 2008.

Shielded twisted pair cable

For environments with greater levels of electromagnetic interference, there are also various types of shielded twisted pair (STP) cables. Some STP variants have metal shielding around each pair of wires, which provides greater protection against crosstalk. Some STP cables have shielding inside the sheath surrounding all of the pairs. This type of shielding is called screening and protects against outside EMI sources. Some cable types have both shielding and screening.

Some now-obsolete data-link layer protocols, such as Token Ring, had physical layer specifications that called for STP cables. Today, however, Ethernet networks only use them in special situations that require additional protection from EMI.

Crossover cables

Twisted pair cables have eight pins in the connector at each end, one for each wire, and connecting the wires to the correct pins is an essential part of the installation or manufacturing process. For internal cable runs, those in walls and ceilings, the process of connecting the cable ends to the jacks is called punching down. The process of attaching connectors to create external patch cables is called crimping.

Twisted pair cable for LANs is wired straight through, meaning that each pin at one end of the cable is connected to the corresponding pin at the other end. A connection between two devices on a network must contain a signal crossover somewhere, so that the signals transmitted over the transmit pin at one end wind up at the receive pin on the other end, but data networks rely on network devices such as switches, routers, and hubs to supply the crossover circuit.

Crossover cables, that is, cables in which the transmit and receive pinouts are reversed on one end, do exist. They have traditionally had two uses: to connect two computers together without a switch or hub and to connect two hubs together.

Standard LAN crossover cables have pins 1 and 3 and pins 2 and 6 transposed. Another type of cable, called a T-1 crossover cable, is used to connect two CSU/DSU devices together for a leased line installation. This cable has pins 1 and 4 and pins 2 and 5 transposed instead.

Coaxial cable

Prior to the introduction of UTP cables, Ethernet networks called for coaxial cable of various types. A coaxial cable consists of a central copper conductor, which carries the signals, surrounded by a layer of insulation. Surrounding the insulation is a shielding, typically made of copper mesh, which functions as a ground, and the whole assembly is encased in a sheath, as shown in Figure 3-3. The mesh shielding in coaxial cables makes them quite resistant to EMI.



FIGURE 3-3 A coaxial cable.

All of the coaxial cable types have been removed from the TIA/EIA network cabling standards, and you are unlikely to encounter any coaxial Ethernet networks in the field, but they are an important part of the history of data networking.

True or false: In a twisted pair cable, the individual wire pairs in the sheath must be twisted at differing rates of twists per foot.

Answer: *True*. The different twist rates are crucial to making the wires in the same cable resistant to crosstalk, that is, the bleeding of signals from one wire pair to another in the same cable.

EXAM TIP CAT3 cable is still included in the Network+ objectives, but in today's networks it is decidedly less prevalent in the installed base, because most CAT3 networks have long since been upgraded to at least Fast Ethernet (at 100 megabits per second), which requires a minimum of CAT5 cable.

True or false: Most UTP cables for Ethernet networks have four pairs (eight wires) in a single sheath, but in some cases the network nodes use only two of the four pairs for actual communication. When this is the case, you can use the remaining pairs for voice traffic, or any other application.

Answer: *False*. Despite the fact that the other two pairs are left unused, it is not permitted to make use of them for other traffic, such as voice telephone communications. Leaving these additional wire pairs unused can also provide an avenue for a future upgrade to a four-pair data networking technology.

True or false: The patch cables used to connect computers to wall plates are crossover cables.

Answer: *False*. All of the patch cables used to construct a LAN are wired straight through. The crossover circuit is provided by the switch or hub.

True or false: One of the main reasons that UTP grew to replace coaxial cable for Ethernet networks is that voice telephone systems also use UTP.

Answer: *True*. Because telephone systems use the same type of cable, telephone cable installers already have the tools and the skills needed to install UTP LANs.

Categorize plenum vs. non-plenum cables

A plenum is an enclosed air space within a building, a passage through which HVAC equipment circulates breathable air, such as the space above a dropped ceiling. When cables are to be installed in an existing building, using the plenums is often the most expedient way to run a cable from one location to another. The only problem with this practice is that in the event of a fire, the outer sheath of many network cables can catch and spread fire and also outgas toxic vapors when it burns. Spreading fire along the entire length of a cable run and releasing these vapors into an air space can obviously be hazardous to the people in the building. A plenum-grade cable is a cable with a sheath that is more fire resistant and produces less toxic smoke when it does burn. Plenum-grade cables are typically much more expensive than non-plenum cables, but building codes in many cities require them when you install data network cable in air spaces.

True or false: Not every drop ceiling or raised floor in a building is necessarily used for air handling.

Answer: *True*. The use of ceiling and floor spaces can vary, depending on local codes and building practices. Some buildings deliver air to the workspace using ducts and remove air through the plenum; others have ducts for both incoming and outgoing air. If you are uncertain whether your intended cabling space is a plenum, look for two sets of air vents, returns, or ducting, or check with the building engineer or an architect.

True or false: Plenum-rated cables are easier to install than cables that are non-plenum-rated.

Answer: *False*. The sheath on plenum-rated cables is actually slightly more rigid than that of standard cables, and as a result, the plenum cables have a smaller bend radius and are a bit harder to install.

Categorize media converters

Ethernet supports physical layer specifications that call for different media, but the other basic elements of the protocol, the frame format and the media access control (MAC) mechanism, remain the same, no matter what medium they use. As a result, you can use a physical layer device called a media converter to connect Ethernet networks that use different media together.

Like a repeater, a media converter is simply a box with two network connectors in it, the only difference being that the connectors are for two different network media. The most common configuration for an Ethernet media converter is one port for UTP cable and one for fiber optic. Thus, if you have a UTP network covering most of your systems, but there are still a few that require the extra-long segment length or EMI resistance that fiber optic provides, you can build a fiber optic network for those computers and connect it to your UTP network by using a media converter.

There are media converters available to connect a UTP segment to either a singlemode or multimode fiber optic cable. For cases where a network uses relatively short multimode segments, and a need for a longer length arises, there

are singlemode to multimode media converters as well. Finally, for legacy networks, there are converters that can connect a coaxial cable segment to a fiber optic one.

True or false: Media converters work by stripping the data-link layer protocol frame off of packets and applying new ones.

Answer: *False*. Media converters are physical layer devices that are not capable of reading data-link layer structures or manipulating them.

EXAM TIP Candidates for the Network+ exam should be familiar with the basic applications of media converters and the cable types they support.

True or false: Media converters can enable administrators to extend the length of a cable run.

Answer: *True*. A UTP Ethernet cable run is limited to a maximum length of 100 meters. However, if you add two media converters and a fiber optic segment in between them, you can exceed that maximum safely.

Categorize cable distance and speed limitations

The various types of cable used to construct networks are characterized by their limitations in transmission speed and in the maximum length of a cable segment. The evolution of cable designs was motivated primarily by the desire to increase network speeds and sizes.

Ethernet speeds have increased by orders of magnitude over the years, and new cables were designed to accommodate those speeds. Maximum segment lengths for UTP cables have remained more static, largely to avoid the need to redesign network installations and move data centers.

Table 3-2 lists the cable types used for Ethernet LANs and their respective transmission speeds and maximum segment lengths.

TABLE 3-2 Cable distance and speed limitations.

CABLE TYPE	CABLE NAME	ETHERNET DESIGNATION	TRANSMISSION SPEED	MAXIMUM SEGMENT LENGTH
Coaxial	RG-8/U	10Base5 (Thick Ethernet)	10 Mbps	500 meters
Coaxial	RG-58A/U	10Base2 (Thin Ethernet)	10 Mbps	185 meters
UTP	CAT3	10Base-T	10 Mbps	100 meters
UTP	CAT5	100Base-TX	100 Mbps	100 meters
UTP	CAT5e	1000Base-T	1,000 Mbps	100 meters
UTP	CAT6	1000Base-T	1,000 Mbps	100 meters
UTP	CAT6	10Gbase-T	10 Gbps	55 meters

CABLE TYPE	CABLE NAME	ETHERNET DESIGNATION	TRANSMISSION SPEED	MAXIMUM SEGMENT LENGTH
UTP	CAT6a	10Gbase-T	10 Gbps	100 meters
Fiber optic	Multimode	100Base-FX	100 Mbps	2 kilometers
Fiber optic	Multimode	1000Base-SX	1,000 Mbps	220–500 meters
Fiber optic	Multimode	1000Base-LX	1,000 Mbps	550 meters
Fiber optic	Multimode	10Gbase-SR	10 Gbps	300 meters
Fiber optic	Singlemode	1000Base-LX	1,000 Mbps	2 kilometers
Fiber optic	Singlemode	10Gbase-LR	10 Gbps	10 kilometers
Fiber optic	Singlemode	10Gbase-ER	10 Gbps	40 kilometers

MORE INFO For more information on the various Ethernet versions, see “Objective 3.7: Compare and contrast different LAN technologies.”

True or false: The maximum segment length of a fiber optic cable depends in part on the nature of the light source generating the signal.

Answer: *True*. Multimode fiber optic cables use an LED light source, providing segment lengths that are shorter than those of singlemode cables, which use a more powerful laser light source.

EXAM TIP This is one of the few Network+ objectives that requires rote memorization. Candidates for the exam must be familiar with all of these cable types, their speeds, and their distance limitations.

True or false: The higher the category of a UTP cable, the more resistant the cable is to various types of crosstalk.

Answer: *True*. The standard for CAT5e cables calls for increased resistance to Near-End Crosstalk (NEXT) and Return Loss (RL) and also adds testing requirements for Power Sum Near-End Crosstalk (PS-NEXT), Equal-Level Far-End Crosstalk (EL-FEXT), and Power Sum Equal-Level Far-End Crosstalk (PS-ELFEXT). The higher categories impose similar restrictions.

Categorize broadband over powerline networking

The technology for transmitting data over power lines has existed for many years. As a low-bandwidth home networking alternative, power line technology runs data signals through residential power lines. These networks typically run at slower speeds than wired Ethernet networks, however, and the increased popularity of wireless LANs has reduced their market share.

Broadband over power lines (BPL) is a similar technology on a larger scale that is designed to supply homes with Internet access by using the public electric

power grid, rather than a dedicated data network, like those of cable television and telephone providers. Although the principle is the same as the home networking products, BPL uses different frequencies to achieve higher communication speeds over longer distances.

The advantages of BPL in principle are obvious: the vast existing power grid infrastructure eliminates the need for the construction of new networks, particularly to remote locations. However, as an incentive to cable and DSL providers to provide service to remote locations, BPL might have had value, but as a practical alternative to traditional broadband technologies it is essentially a failure as a concept.

True or false: Transmitting broadband data over power lines generates a great deal of electromagnetic interference that compromises radio reception

Answer: *True*. Because of this and other shortcomings, a few pilot programs designed to provide citywide subscribers with broadband access have not been economically successful.

EXAM TIP BPL is a marginal technology that is covered in the Network+ exam only in the most general manner. Candidates need only know that it exists and that it uses standard power lines to provide Internet access.

True or false: Variances in the age and utility of power cables have made it difficult to establish performance standards for powerline networking.

Answer: *True*. The physical characteristics of the electrical network are highly various, with some areas using cables that are nearly a century old. These variations make it very difficult to establish a standard for power line communication that would accommodate every instance.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which UTP cable type would be the most economical choice for a new 1000Base-T network that will someday be upgraded to 10Gbase-T?
2. Which cable type is best suited for a network installation in a robotic manufacturing facility that generates high levels of electromagnetic interference?
3. Under what circumstances would you be likely to install singlemode fiber optic cable, rather than multimode?
4. For what reason would you ever connect two computers together with a UTP crossover cable?

Objective 3.2: Categorize standard connector types based on network media

For this exam objective, you must be familiar with the various types of connectors that installers use to create patch cables and to terminate internal cable runs at wall plates and data centers. While the connectors used on fiber optic cables are interchangeable, copper cables have specific requirements for connectors.

Exam need to know

- Categorize fiber connector types

For example: What are the characteristics of the ST, SC, LC, and MTRJ connectors used with fiber optic cables?

- Categorize copper connector types

For example: What are the characteristics of the RJ-45, RJ-11, BNC, F-connector, DB-9 (RS-232) patch panel, and 110 block (T568A and T568B) connectors used with copper cables?

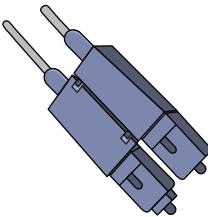
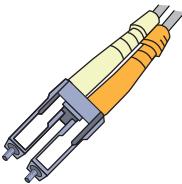
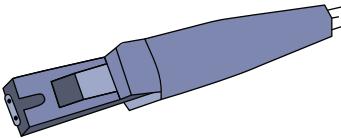
Categorize fiber connector types

Unlike twisted pair and coaxial cables, which give you no choice, fiber optic cables offer a large number of different connectors. The type of connector you elect to use is a matter of compatibility with your existing equipment and personal preference. All fiber optic connectors perform basically the same function: to precisely join two ends of core filament together, face to face, so that light pulses can pass from one cable segment to another. Most are spring loaded, so that the faces of the cores are pressed together firmly.

There are literally dozens of fiber optic connectors on the market, the most common of which are identified in Table 3-3.

TABLE 3-3 Fiber optic cable connector types.

CONNECTOR TYPE	CONNECTOR IMAGE
ST (Straight Tip)	

CONNECTOR TYPE	CONNECTOR IMAGE
SC (Subscriber Connector)	
LC (Local Connector or Lucent Connector)	
MT-RJ (Mechanical transfer – Registered Jack)	

True or false: Many of the traditional fiber optic connectors require the use of a curing oven for installation.

Answer: *True*. Traditional fiber optic connectors must be glued to the cable and the glue cured in an oven. There are now some specialized connectors that use a quick drying epoxy that hardens at room temperature.

EXAM TIP The Network+ exam objectives refer to the fiber optic connectors by their abbreviations only, due in part to some disagreement about the correct spelling out of the acronyms. Candidates for the exam should concentrate on knowing the abbreviations ST, SC, LC, and MT-RJ.

True or false: The newer fiber optic connectors, such as MT-RJ, are generally smaller than the older models, such as the ST connector.

Answer: *True*. The latest fiber optic connector designs are relatively small, so as to fit more connections into the same space on a rack mounted switch.

Categorize copper connector types

Twisted pair cables use modular connectors that are most commonly referred to by the telephone networking designation RJ-45. However, the actual name for the UTP connector that Ethernet networks use is 8P8C, indicating that there are eight positions in the connector and eight electrical contacts in those positions.

Network interface adapters, wall plates, patch panels, and other networking components, such as switches and hubs, all have female connectors. The patch cables used to connect everything together have male connectors, as shown in Figure 3-4.

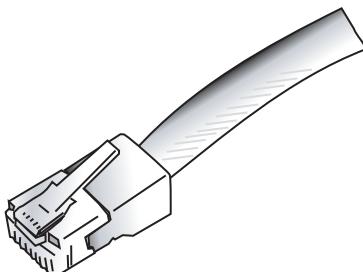


FIGURE 3-4 A UTP patch cable with an RJ-45 connector.

NOTE Connectors for UTP cables are categorized, just as the cables are. When purchasing components for a modular cable installation, be sure that you select connectors of the same category rating as your cable. A cable installation must be rated according to its lowest-rated component. You might use all CAT6 cable for your network, but if you use CAT5 connectors, it is a CAT5 installation.

Although designed for telephones, the RJ-45 connector was actually seldom used for that purpose. The smaller, 6P4C RJ-14 connector became the standard for telephone connections, and remains so to this day. When UTP came into use for data networks, the connectors were so similar in appearance to the RJ-45 telephone connector that network administrators adopted the designation for their own use.

NOTE 6P4C and 6P6C connectors are frequently, incorrectly referred to as RJ-11 connectors, but the standard describes an RJ-11 as a 6P2C connector; the 6P6C is an RJ-25.

True or false: RJ-11 connectors are not used for local area networking.

Answer: *True*. RJ-11 connectors are used for modular telephone connections only, not for local area networking.

In the data center, where you find the other end of the cables, installers typically connect the wires to 110 blocks. A 110 block is a connector with an open framework, into which the installer lays the eight individual wires found in the UTP cable, as shown in Figure 3-5. The layout of the wires in the block can use the T568A or T568B pattern, as long as the connections on both ends are the same. The installer then uses a special tool to press each wire into the block, simultaneously creating both mechanical and electrical connections. The process is called punching down.

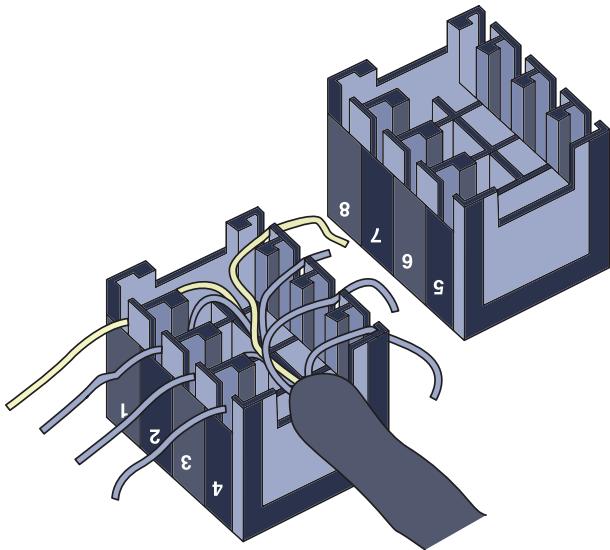


FIGURE 3-5 A 110 block.

110 blocks are found on the back side of rack-mounted patch panels and individual jacks. Like a collection of individual jacks, the front of a patch panel consists of a row of female RJ-45 connectors, which provide modular access to all of the cable connections.

True or false: You can use the same patch cables to connect patch panel ports to switches that you use to connect computers to wall plates.

Answer: *True*. Patch panels, switches, computers, and wall plates all have the same female RJ-45 connectors, enabling administrators to complete connections using patch cables with male RJ-45 connectors on both ends.

EXAM TIP Although you would be technically correct in referring to an Ethernet cable as having 8P8C connectors, few (if any) other people will know what you are talking about. RJ-45 is the colloquial term which the Network+ exam and much of the networking literature (as well as the rest of this book) have adopted to refer to this connector.

The coaxial cables used on Thin Ethernet networks run all the way to the individual computers, using a T configuration with three BNC connectors, as shown in Figure 3-6.

LANs that use a cable television provider for Internet access have a modem to which you connect the incoming coaxial cable, using a screw-on F connector. The modem also usually has an RJ-45 connector, which you use to attach it to a router or a computer.

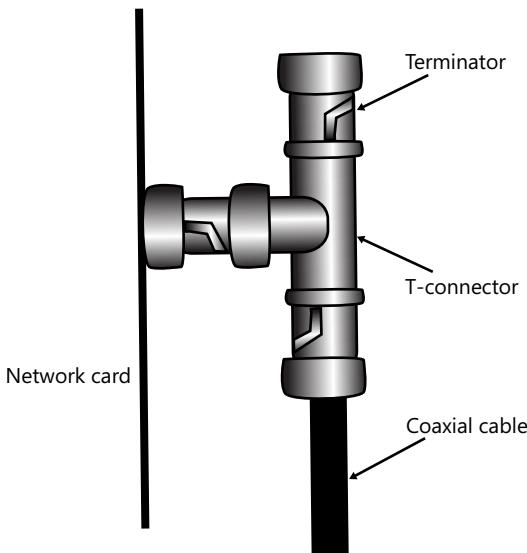


FIGURE 3-6 A Thin Ethernet cable with a T fitting and three BNC connectors.

Although they are not present on many new computers, serial ports have also provided an avenue for connecting a system to a network in the past. The original design for the IBM PC called for the inclusion of one or more serial ports with a 9-pin male D-subminiature connector (also known as a DE9, although incorrectly referred to as a DB9). The serial connection is based on the Electronic Industries Alliance (EIA's) RS-232 standard, which predates the PC by many years and defines the functions of the nine pins and the speed of the interface.

True or false: DB-9 connectors are often found on routers and switches.

Answer: *True*. DB-9 connectors on routers and switches enable administrators to connect a terminal for access to the device's configuration interface.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What tools are needed to attach RJ-45 connectors to patch cables?
2. How many BNC connectors are required to attach a Thin Ethernet coaxial cable to a computer?
3. What is the significance of the 8P8C designation used to refer to RJ-45 connectors?
4. How are fiber optic connectors attached to the cables?
5. What type of connector is found in the wall plates of an internal UTP Ethernet installation?

Objective 3.3: Compare and contrast different wireless standards

This exam objective covers material that was already examined in “Objective 2.2: Given a scenario, install and configure a wireless network.” Selecting a wireless standard for installation is essentially a matter of comparing them, so exam candidates should consult the material for that objective, in addition to the coverage for this one.

Exam need to know

- Compare and contrast the 802.11 a/b/g/n wireless standards
For example: How do the wireless standards compare in terms of distance, speed, latency, frequency, channels, MIMO, and channel bonding?

Compare and contrast the 802.11a/b/g/n wireless standards

The first version of the IEEE 802.11 standard was published in 1997, and was not widely implemented. The standard defined three physical layer specifications: an infrared medium running at 1 Mbps and two types of radio signal modulation using the 2.4-GHz band: Direct-Sequence Spread Spectrum (DSSS) and Frequency-Hopping Spread Spectrum (FHSS), both running at 1 or 2 Mbps. No one has ever marketed an implementation of the infrared option, and the slow transmission speeds of these early spread spectrum radio specifications made them an unattractive solution, even for users of the original 10 Mbps Ethernet networks.

IEEE 802.11a

Although it might appear to be an interim step, the IEEE 802.11a amendment actually represented a fundamental change in the technology. One of the ongoing problems with 802.11 networks is the heavy use of the 2.4-GHz band by a variety of consumer products, including cordless telephones and Bluetooth devices. Wireless LAN performance can degrade in such a crowded environment, causing speed reductions or even service interruptions.

The 802.11a amendment calls for the use of the relatively vacant 5-GHz band and a different form of modulation called Orthogonal Frequency-Division Multiplexing (OFDM). The data transfer rate can be as high as 54 Mbps, with fallbacks to 48, 36, 24, 18, 12, 9, and 6 Mbps.

Because of complications in manufacturing, 802.11a products arrived on the market after 802.11b devices, and they cost significantly more. The 802.11a technology also developed a reputation for having a shorter range than 802.11b and for being more susceptible to signal loss from attenuation. Whatever the reason, dedicated 802.11a equipment did not sell well. Later devices supporting both the 802.11a and 802.11b standards eventually came to market, and cross-compatibility between standards soon became a major selling point.

IEEE 802.11b

Despite its later designation, the IEEE 802.11b standard represents the next step in the evolution of the original 802.11 document. The 802.11b document retains the 2.4-GHz frequency and the DSSS modulation from the original standard, but increases the transmission speed to as much as 11 Mbps. This was, for the first time, a wireless LAN that could run at an acceptable speed for the typical network user. Adoption of the 802.11b standard was quick; many manufacturers released products, and prices fell rapidly.

IEEE 802.11g

The IEEE 802.11g amendment built on the 802.11b technology by adopting the OFDM modulation from 802.11a while retaining the 2.4-GHz frequency from 802.11b. The result was a new standard that was fully backward compatible with 802.11b equipment but that pushed the maximum transfer rate up to 54 Mbps.

Even before the 802.11g amendment was officially ratified by the IEEE, the public began buying new wireless LAN products that supported both the "draft" 802.11g and 802.11b standards. Some added support for the 802.11a standard as well.

IEEE 802.11n

The development of 802.11b/802.11g technology eventually reached a point at which a fundamental change was needed to improve performance. The IEEE 802.11n amendment introduced several modifications to the technology, including the potential doubling of channel widths and the addition of the 5-GHz frequency band from 802.11a to the standard 2.4-GHz band from 802.11b/g. 802.11n also includes two innovations, MIMO and frame aggregation, which, combined with these other improvements, pushed network transmission speeds well beyond the 54 Mbps realized by 802.11g to levels as high as 600 Mbps.

Multiple-Input Multiple-Output (MIMO) is a physical layer enhancement that enables wireless devices to multiplex signals over a single channel simultaneously, by using a technique called Spatial Division Multiplexing (SDM). Each 802.11n device has an array of transmit and receive antennae and a transceiver capable of sending and receiving separate signals by using separate frequencies.

In addition to MIMO, the 802.11n standard also enabled devices to double their channel widths from 20 MHz to 40 MHz, nearly doubling the data transfer rate in the process.

In wireless networking, physical layer improvements can only do so much, because the control overhead is so high. In addition to the data-link layer frame, there are acknowledgment messages, spaces between frames, and radio communication transmissions, which in some circumstances can add up to more data than is carried in the payload. Frame aggregation is a technique that combines the payload data from several frames into one large frame, thus reducing the amount of overhead and increasing the information throughput of the network.

The 802.11n standard retained the fallback capabilities of the previous documents. Depending on a multitude of conditions, devices operating at peak

speeds of 600 Mbps (using 4×4:4 MIMO and a 40-MHz channel) can drop down to successively lower rates and then speed up again when environmental conditions change. In addition, most 802.11n devices also support 802.11g and 802.11b, with some adding 802.11a as well.

Transfer rates for 802.11n networks are dependent on the capabilities of the equipment, the configuration settings selected by the administrator, and the usual environmental conditions that affect all wireless LANs.

True or false: All of the 802.11 standards provide fallbacks from their top speeds.

Answer: *True*. As with all speed measurements in the networking industry, the maximum transmission speeds touted by makers of wireless networking equipment have to be taken with a degree of dubiety. All wireless networking technologies, including the latest and greatest, are subject to variations in transfer speed caused by any number of factors, from physical obstructions to electromagnetic interference to weather conditions. All of the 802.11 standards include a mechanism by which network devices can drop to a lower speed when conditions prevent peak performance.

MORE INFO For more information comparing and contrasting the 802.11 wireless LAN standards, see “Objective 2.2: Given a scenario, install and configure a wireless network.”

True or false: MIMO is only available on networks running IEEE 802.3n using the 5 GHz frequency band.

Answer: *False*. While it is true that only 802.11n networks can use MIMO, they do not have to be using the 5 GHz band. Networks using the 2.4 GHz band can use MIMO as well.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What does it mean when an 802.11n product has the designation 4×4:4?
2. Which 802.11 standards allow for the use of the 5 GHz frequency band?
3. How does frame aggregation improve the efficiency of an 802.11n network?
4. Which of the IEEE 802.11 wireless networking standards is capable of supporting the 54 Mbps transmission speed and is backward compatible with IEEE 802.11b?

Objective 3.4: Categorize WAN technology types and properties

For this exam objective, you must be familiar with a large number of wide area networking technologies. A wide area network (WAN) is a communications link that spans a relatively long distance and connects two or more computers or LANs together. Most WAN connections are point to point, involving two systems only.

There are three main applications for WANs, as follows:

- To connect a LAN to the Internet
- To connect two LANs at remote locations
- To connect a remote computer to a distant LAN

Each of these applications has different requirements in terms of transmission rates, availability, security, and of course, cost, and there are a variety of WAN solutions available for each one.

Exam need to know

- Categorize WAN technology types

For example: Do you know the properties of the following WAN types: T1/E1, T3/E3, DS3, OCx, SONET, SDH, DWDM, satellite, ISDN, cable, DSL, cellular, WiMAX, LTE, HSPA+, fiber, dialup, PON, frame relay, ATM?

- Categorize WAN technology properties

For example: What is the significance of the following WAN properties: circuit switch, packet switch, speed, transmission media, and distance?

Categorize WAN technology types

The following sections examine some of the most important WAN technologies and their properties.

Leased lines

A leased line is a permanent telephone connection between two locations that provides a predetermined amount of bandwidth at all times. Leased lines can be analog or digital, although most of the leased lines installed today are digital.

The most common leased line configuration in the United States is called a T-1, which runs at 1.544 Mbps. In Japan, the same type of connection uses the name J-1. The European equivalent of a T-1 is called an E-1, which runs at 2.048 Mbps. Many organizations use T-1 lines to connect their networks to the Internet or to connect remote networks together. For applications requiring more bandwidth, a T-3 connection runs at 44.736 Mbps, an E-3 runs at 34.368 Mbps, and a J-3 runs at 32.064 Mbps. These designations are collectively known as T-carrier, E-carrier, and J-carrier services, respectively.

Leased line services use a framing method called Digital Signal 1 (DS1). A DS1 frame consists of 8-bit channels called Digital Signal 0s (DS0s), plus a framing bit used for synchronization. A T-1 has 24 such channels, as does a J-1; an E-1 has 32 channels. Each DS0 is sampled 8,000 times per second, resulting in discrete 64-Kbps channels. In a T-1, 24 64-Kbps channels with framing bits equals 1.544 Mbps. This transmission technique is called time division multiplexing.

Faster connections use successive digital signal levels, the definitions of which vary according to region. The designations, number of channels, and transmission rates for the most commonly used levels of the T-carrier, E-carrier, and J-carrier systems are summarized in Table 3-4.

TABLE 3-4 T-Carrier, E-Carrier, and J-Carrier configuration data.

	NORTH AMERICA	EUROPE	JAPAN
DS0	64 Kbps	64 Kbps	64 Kbps
DS1	T-1	E-1	J-1
	24 channels	32 channels	24 channels
	1.544 Mbps	2.048 Mbps	1.544 Mbps
DS2	T-2	E-2	J-2
	96 channels	128 channels	96 channels
	6.312 Mbps	8.448 Mbps	6.312 Mbps
DS3	T-3	E-3	J-3
	672 channels	512 channels	480 channels
	44.736 Mbps	34.368 Mbps	32.064 Mbps
DS4	T-4	E-4	J-4
	4,032 channels	2,048 channels	480 channels
	274.176 Mbps	139.264 Mbps	97.728 Mbps

Because a leased line is a point-to-point connection between two systems only, there is no need for any sort of addressing, as on an Ethernet network. Because the connection is permanent, there is no need for a connection establishment process or a negotiation of protocols, as in a Point-to-Point Protocol connection.

When you use a leased line to connect networks together, you can combine the channels into a single data pipe, or configure them in any combination of voice and data. You can also install a leased line that uses some, but not all, of the channels in a T-1. This is called a *fractional T-1* service, and you can use it to specify exactly the amount of bandwidth you need.

At each end of a leased line connection, you must have a device called a channel service unit/data service unit (CSU/DSU), which you connect to your data network by using a router, as shown in Figure 3-7.

SONET/SDH

As fiber optic cables began replacing copper for long distance runs, various competing telecommunications companies began creating their own proprietary standards for fiber optic communications. Eventually, these companies created a common standard for fiber optic interoperability, called the Synchronous Optical Network (SONET) in the United States and Canada, and the Synchronous Digital Hierarchy (SDH), in the rest of the world.

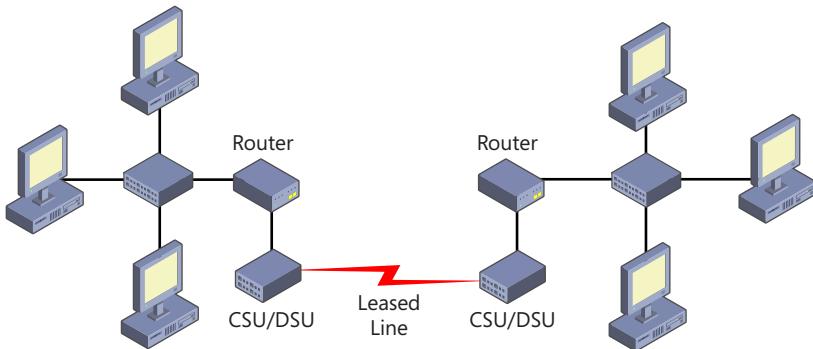


FIGURE 3-7 A leased line connection.

SONET is a physical and data-link layer standard that defines a method for building a synchronous telecommunications network based on fiber optic cables. First ratified by the American National Standards Institute (ANSI), SONET was then adapted by the International Telecommunications Union (ITU), which called it the *Synchronous Digital Hierarchy* (SDH). Intended as a replacement for the T-carrier and E-carrier services used in the United States and Europe, SONET provides connections at various *optical carrier* (OC) levels running at different speeds. The idea behind SONET is to create a standardized series of transmission rates and formats, eliminating the problems that can affect connections between different types of carrier networks. The OC levels are listed in Table 3-5.

TABLE 3-5 SONET OC levels.

OC LEVEL	DATA TRANSMISSION RATE (IN MBPS)
OC-1	51.84
OC-3	155.52
OC-6	311.04
OC-9	466.56
OC-12	622.08
OC-18	933.12
OC-24	1244.16
OC-36	1866.24
OC-48	2488.32
OC-96	4976.640
OC-192	9953.280
OC-768	39813.120

In addition to the high speeds available on fiber optic media, these links also use various types of multiplexing to carry multiple signals on a single link. One of these techniques, wavelength division multiplexing (WDM), uses different light wavelengths to encode different signals.

One WDM variant, called dense wave division multiplexing (DWDM), calls for the use of devices called erbium-doped fiber amplifiers (EDFAs), which are designed to amplify wavelengths in the 1525-nanometer to 1565-nanometer or 1570-nanometer to 1610-nanometer bands without converting them to electrical signals. This is a cheaper and more efficient method of transmitting long-range optical signals than the optical-electrical-optical (OEO) regenerators they replace.

Packet switching

T-carrier and SONET links are both point-to-point technologies that telecommunications companies have used to create the mesh of networks that form the global telecommunications system and the Internet. When you access a website on the Internet, your packets very likely traverse several of these links.

However, the relatively simple frame formats that these technologies use to send signals from one end of a link to the other have nothing to do with the addressing needed to get a request from your browser to a web server in another city. To do that, you must add a packet-switching capability that joins all of these separate links. The two packet-switching protocols most commonly used on WAN connections today are frame relay and asynchronous transfer mode (ATM).

FRAME RELAY

Frame relay is a packet-switching WAN solution that can provide bandwidth similar to that of a leased line, but with greater flexibility. Frame relay services range from 56 Kbps all the way up to T-3 speeds, but the subscriber is not permanently locked into a specific transmission rate, as with a leased line. When you enter into a contract with a frame relay provider, you agree on a specific amount of bandwidth, called the *committed information rate* (CIR), which is the base speed of your link.

However, the frame relay service can provide additional bandwidth (called *bursts*) during your high-traffic periods by borrowing it from other circuits that are not operating at full capacity. In addition to the CIR, you also negotiate a *committed burst information rate* (CBIR), which is the maximum amount of bandwidth that the provider agrees to furnish during burst periods. Your contract specifies the duration of the bursts you are permitted. If you exceed the bandwidth specified in the agreement, you must pay an extra charge.

A frame relay connection is not a permanent link between two points, as is a leased line. Instead, each of the two sites is connected to the service provider's nearest point of presence, usually by using a standard leased line or, formerly, an ISDN connection. The provider's network takes the form of a frame relay cloud, which enables the leased line at one site to be connected to the line at the other site, as shown in Figure 3-8. This connection through the cloud from one point of presence to another is called a *permanent virtual circuit* (PVC).

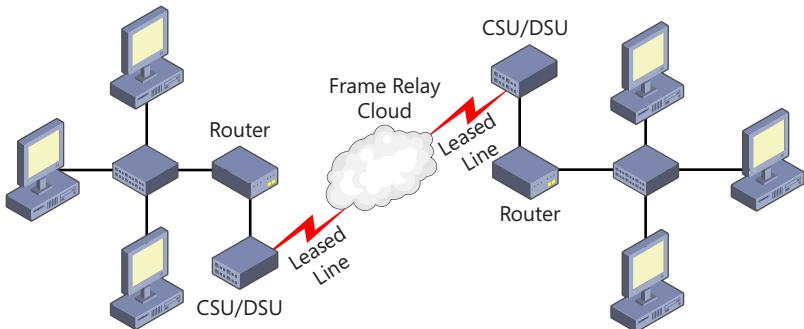


FIGURE 3-8 A frame relay connection.

Because each site uses a local telephone provider for its leased line to the cloud, the cost is generally less than it would be to have a single long-distance leased line connecting the two different sites.

The hardware device that provides the interface between the LAN at each site and the connection to the cloud is called a frame relay assembler/disassembler (FRAD). A FRAD is a network layer device that strips off the LAN's data-link layer protocol header from each packet and repackages it for transmission through the cloud. One of the main advantages of frame relay is that you can use a single connection to a frame relay provider to replace several dedicated leased lines.

Other, newer packet-switching solutions, such as Asynchronous Transfer Mode (ATM), have begun to replace frame relay in the marketplace, and many businesses have turned to using broadband Internet services and virtual private network (VPN) connections for relatively low-speed LAN-to-LAN solutions. However, in remote areas not serviced by DSL and CATV networks, a combination of fractional T-1 services and frame relay can be the most economical solution available.

ASYNCHRONOUS TRANSFER MODE (ATM)

Asynchronous Transfer Mode (ATM) is a protocol that was originally designed to carry voice, data, and video traffic on both LANs and WANs. Today, ATM is most commonly used in WAN connections. Unlike most data-link layer protocols, ATM uses fixed-length, 53-byte frames (called *cells*) and provides a connection-oriented, full-duplex, point-to-point service between devices. Because the cells are a uniform size, unlike the variable-sized packets used by most networking protocols, ATM can provide a guaranteed, predefined quality of service.

This makes it easier to regulate and meter the bandwidth passing over a connection, because when the data structures are of a predetermined size, network traffic becomes more readily quantifiable, predictable, and manageable. With ATM, it's possible to guarantee that a certain quantity of data will be delivered within a given time. This makes the technology more suitable for a unified voice/data/video network than a nondeterministic protocol such as Ethernet, no matter how fast it runs. In addition, ATM has quality of service (QoS) features built into the protocol

that enable administrators to reserve a certain amount of bandwidth for a specific application.

There are no broadcast transmissions in ATM, and data is relayed between networks by switches, not routers. ATM speeds range from a 25.6-Mbps service, originally intended for desktop LAN connections, to a 2.46-Gbps service. Physical media range from standard multimode fiber optic and UTP cables on LANs to SONET or T-carrier services for WAN connections.

You can use an ATM packet-switching service for your WAN links in roughly the same way as you would use frame relay, by installing routers at your sites and connecting them to the carrier's points of presence by using leased lines. This process transmits the LAN data to the point of presence first and then repackages it into cells. It's also possible, however, to install an ATM switch at each remote site, either as part of an ATM backbone or as a separate device providing an interface to the carrier's network. In this case, the switch converts the LAN data to ATM cells at each site before transmitting it over the WAN.

Internet access

There are many different WAN technologies that users and administrators can use to connect computers and networks to the Internet, and they provide varying levels of speed, security, and flexibility. The following sections examine some of the physical layer options that you can use for Internet access connections in the home, the small business, or the enterprise network.

PUBLIC SWITCHED TELEPHONE NETWORK

The Public Switched Telephone Network (PSTN) is just a technical name for the standard analog telephone system that has existed in some form since the late nineteenth century, also commonly known as the Plain Old Telephone Service (POTS). This voice-based system, found all over the world, can use asynchronous modems to transmit data between computers at virtually any location. The PSTN service in your home or office probably uses copper-based twisted-pair cable and modular RJ-11 jacks.

The PSTN connection leads to a central office belonging to a telephone service provider, which can route calls from there to any other telephone in the world. Unlike a LAN, which is digital and uses packet switching, the PSTN is an analog, circuit-switched network.

Before computer data can be transmitted over the PSTN, the digital signals generated by a computer must be converted to analog signals that the telephone network can carry. A device called a *modulator/demodulator*, more commonly known as a modem, handles this conversion. A modem takes the digital signals fed to it by a computer, converts them to analog signals, and then transmits them over the PSTN, as shown in Figure 3-9. At the other end of the PSTN connection, another modem performs the same process in reverse, converting the analog data back into its digital form and passing it to another computer. The combination of the interfaces to the two computers, the two modems, and the PSTN connection forms the physical layer of the networking stack.

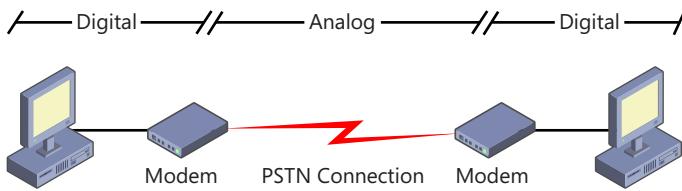


FIGURE 3-9 A PSTN connection between two computers.

The PSTN was designed for voice-grade transmissions, not data-grade transmissions. As a result, connections are relatively slow, with a maximum speed of only 33.6 Kbps when both communicating devices use analog PSTN connections. A 56-Kbps connection requires that one of the connected devices have a digital connection to the PSTN.

PSTN/modem connections were, in the past, the primary means for connecting computers to the Internet. However, the emergence of various broadband technologies running at much higher speeds has rendered dial-up connections nearly obsolete. Because the PSTN can connect virtually anyone to anywhere, it is the most flexible WAN connection available.

INTEGRATED SERVICES DIGITAL NETWORK (ISDN)

The Integrated Services Digital Network (ISDN) is a digital communications service that uses the same network infrastructure as the PSTN. It was designed as a complete digital replacement for the analog telephone system, but it had few supporters in the United States until the need for faster Internet connections led people to explore its capabilities. However, after other high-speed Internet access solutions became available, ISDN was all but forgotten.

ISDN is a dial-up service, like the PSTN, but its connections are digital, so no analog/digital conversions are required. Because it is a dial-up service, you can use ISDN to connect to different networks. For example, if you have an ISDN connection to the Internet, you can change ISPs simply by dialing a different number. However, because ISDN needs special equipment, it cannot be used in mobile devices, such as laptop computers, when traveling away from the service location.

ISDN also delivers greater transmission speeds than the PSTN. The ISDN Basic Rate Interface (BRI) service consists of two 64-Kbps channels (called *B channels*) that carry the actual user data, plus one 16-Kbps channel (called a *D channel*) that carries only control traffic. Because of these channel names, the BRI service is sometimes called *2B+D*. The B channels can function separately, or the subscriber can choose to combine them into a single 128-Kbps connection.

A higher grade of service, called Primary Rate Interface (PRI), consists of 23 B channels and one 64-Kbps D channel. The total bandwidth is the same as that of a T-1 leased line. PRI is used primarily to terminate a large number of ISDN BRI connections from remote sites into a single site such as a data center or headquarters location. Some organizations also use ISDN as a redundant solution to leased-line connectivity.

NOTE Because of the increased speed at which ISDN operates, the length of the connection is limited. The location of the terminal equipment must be within 18,000 feet of the telephone company's nearest central office (or a suitable repeating device).

ISDN can be a viable solution for customers who are a long distance from the nearest telephone company point of presence. Users can also disconnect ISDN links when they are not in use. This allows ISDN customers to avoid paying for bandwidth they are not using and eliminates a potential window through which intruders can access the network.

DIGITAL SUBSCRIBER LINE (DSL)

Digital Subscriber Line (DSL) is a blanket term for a variety of digital communication services that use standard telephone lines and provide data transfer speeds much faster than the PSTN or even ISDN. Each of the various DSL service types has a different descriptive word or phrase added to its name, which is why some sources use the generic abbreviation xDSL. Some of the many DSL services are shown in Table 3-6.

TABLE 3-6 DSL services and their properties.

SERVICE	TRANSMISSION RATE	LINK LENGTH	APPLICATIONS
High-bit-rate Digital Subscriber Line (HDSL)	Up to 1.544/2.048 Mbps	12,000 feet	Used by large networks as a substitute for T-1 or E-1 leased line connections, LAN and Private Branch Exchange (PBX) interconnections, or frame relay traffic aggregation
Symmetric Digital Subscriber Line (SDSL)	Up to 1.544/2.048 Mbps	10,000 feet	Same as HDSL
Multirate Symmetric Digital Subscriber Line (MSDSL)	Up to 2 Mbps	29,000 feet	A variant of SDSL that can use more than one transfer rate, set by the service provider
Asymmetric Digital Subscriber Line (ADSL)	Up to 8 Mbps downstream; up to 1.3 Mbps upstream	18,000 feet	Internet/intranet access, remote LAN access, virtual private networking, video on demand, Voice over IP (VoIP)

Service	Transmission Rate	Link Length	Applications
ADSL2	Up to 12 Mbps downstream; up to 1.4 Mbps upstream	18,000 feet	Same as ADSL
ADSL2+	Up to 24 Mbps downstream; up to 3.5 Mbps upstream	18,000 feet	Same as ADSL
Rate-Adaptive Digital Subscriber Line (RADSL)	Up to 7 Mbps downstream; up to 1.088 Mbps upstream	18,000 feet	Same as ADSL, except that the transmission speed is dynamically adjusted to accommodate the link length and signal quality
ADSL Lite	Up to 1.5 Mbps downstream; up to 512 Kbps upstream	18,000 feet	Internet/intranet access, remote LAN access, IP telephony, videoconferencing
Very High-rate Digital Subscriber Line (VDSL)	Up to 51.84 Mbps downstream; up to 16 Mbps upstream	4500 feet	Multimedia Internet access, high-definition television delivery
Internet Digital Subscriber Line (IDSL)	Up to 144 Kbps	18,000 feet	Internet/intranet access, remote LAN access, IP telephony, videoconferencing

Many DSL services run at different upstream and downstream speeds. These types of services are described as *asymmetrical*. The different speeds occur because some DSL signals cause greater levels of crosstalk in the data traveling from the customer site to the central office than in the other direction. For end-user Internet access, this asymmetrical behavior is usually not a problem, because web surfing and other common activities generate far more downstream traffic than upstream traffic. However, if you plan to use DSL to connect your own servers to the Internet, make sure that you obtain a service that is symmetrical or that offers sufficient upstream bandwidth for your needs. DSL services are also subject to distance restrictions.

An ADSL connection requires an ADSL Termination Unit-Remote (ATU-R), which is sometimes called a *DSL transceiver* or a *DSL modem*. You will also need a *line splitter* if you will also use the line for voice traffic. The ATU-R connects to your computer by using either a standard Ethernet network interface adapter or a USB port. At the other end of the link at the ISP's site is a more complicated device called a Digital Subscriber Line Access Multiplexer (DSLAM), shown in Figure 3-10.



FIGURE 3-10 An ADSL connection.

Unlike ISDN connections, DSL connections are direct, permanent links between two sites that remain connected at all times. This means that if you use DSL to connect to the Internet, the telephone company activates the DSL connection between your home or office and the ISP's site. In many cases, however, telephone companies are themselves offering DSL Internet access, which eliminates one party from the chain.

CABLE TELEVISION (CATV) NETWORKS

All of the remote connection technologies described up to this point rely on cables installed and maintained by telephone companies. However, the cable television (CATV) industry has also been installing a vast network infrastructure throughout most of the United States over the past few decades. In many cases, the long-distance cable network runs use fiber optic cable, but the connections to individual premises—that is, the cable that enters your home or office—is copper-based coaxial. This combination is called a hybrid fiber coaxial (HFC) network.

In recent years, many CATV systems have started providing Internet access to their customers through the same cable used for the TV service. CATV Internet access can be very fast, sometimes as fast as 50 Mbps downstream, and is typically available at multiple performance levels, with varying prices. CATV networks use broadband transmissions, meaning that the one network medium carries many discrete signals at the same time.

Each TV channel you receive over cable is a separate signal, and all the signals arrive over the cable simultaneously. By devoting some of this bandwidth to data transmissions, CATV providers can deliver Internet data at the same time as the television signals. If you already have cable TV, installing the Internet service is simply a matter of connecting a splitter to the cable and running it to a *cable modem*, which is connected to an Ethernet network interface adapter in a computer or router, as shown in Figure 3-11.

Like most residential DSL services, CATV data connections are asymmetrical. CATV networks carry signals primarily in one direction, from the provider to the customer. There is a small amount of upstream bandwidth, part of which is allocated for Internet traffic. In most cases, the upstream speed of a CATV connection is far less than the downstream speed. This makes the service unsuitable for hosting Internet servers, but it is still faster than a PSTN connection.

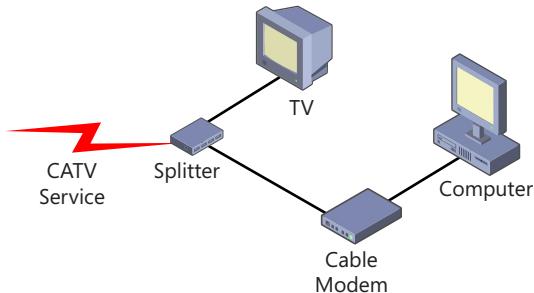


FIGURE 3-11 A CATV connection.

SATELLITE-BASED SERVICES

There are still some Internet users who are not located near enough to a telephone provider's point of presence for DSL or ISDN, and who cannot get Internet access through a cable television network. For these users, the only economical high-bandwidth alternative is satellite-based Internet access. As with satellite television, a user of this service must have a satellite dish pointed at one of the geosynchronous communications satellites orbiting the earth.

Unlike early satellite Internet services, which could only download from the satellite and required a separate dial-up PSTN connection for upstream traffic, today's satellite ISPs provide bidirectional communications. Like DSL and CATV networks, the satellite connection is asymmetrical, with a relatively limited amount of upstream bandwidth available. Satellite Internet access is more expensive than the other popular alternatives, especially when you factor in the hardware costs, and it is usually not as fast as DSL or CATV. But for some users, it is the only option.

LAST MILE FIBER

Telecommunications companies typically use fiber optic cables for long distance runs, because of their resistance to attenuation and electromagnetic interference. However, the so-called last mile, the connection from the nearest switch to the end user's premises, has traditionally been a copper-based medium. However, some of the larger communications companies have launched high-speed broadband services that run fiber optic cables through some or all of the last mile.

Generic names for these services are various, depending on where the fiber optic cable terminates, but they all follow the format FTTx, where x is the terminus of the fiber optic cable. Fiber-to-the-node (FTTN) and Fiber-to-the-curb (FTTC) refer to services that run fiber to within a few kilometers or less than a kilometer from the subscriber's location, respectively. Fiber-to-the-premises (FTTP), sometimes called Fiber-to-the-home (FTTH), runs the fiber optic cable all the way to a demarcation point on the outside wall or in the basement of the subscriber's building.

Fiber-to-the-desk (FTTD) refers to a fiber optic run all the way to a terminal or media converter inside the subscriber's premises.

The performance of these services depends on the proximity of the fiber optic run to the user's premises. Running the fiber optic cable all the way to the building means that the high-speed services can connect directly to a home or office network running a Gigabit Ethernet or 802.11 wireless LAN. If the fiber terminus is more than 100 meters away, then some sort of interim medium is needed, such as VDSL.

Fiber optic cable is much more difficult and expensive to install than any of the standard copper cable types. Deploying fiber in the last mile is an expensive proposition, and one of the methods providers use to control costs is the passive optical network.

A passive optical network (PON) is an arrangement in which data from an optical line terminal (OLT) at the provider's central office runs through a single fiber optic cable to a series of optical splitters near the subscribers' premises. The splitters duplicate all of the incoming signals and send them out through separate fiber runs to optical network terminals (ONTs) at the individual users' locations, as shown in Figure 3-12.

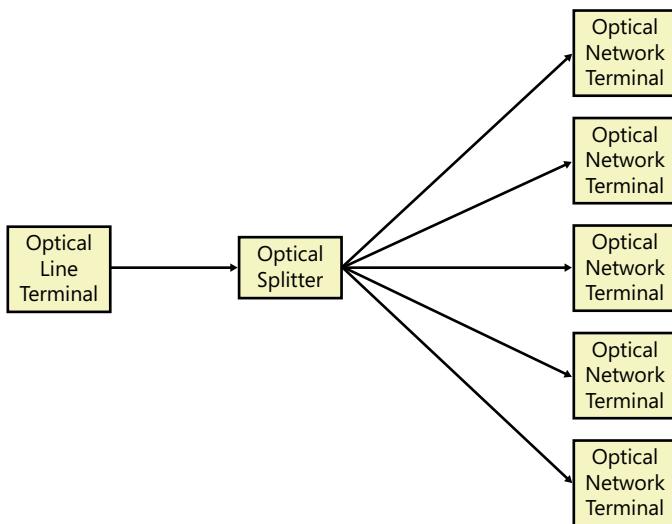


FIGURE 3-12 A passive optical network (PON).

A PON is a lower-cost alternative to an active optical network (AON), because it uses unpowered splitters to duplicate the incoming signals, rather than a complex router that separates the data stream into the packets intended for each subscriber.

Commercial implementations of this technology typically bundle voice, video, and data services into a single package, providing a complete media solution for the subscriber. Internet access is generally available in several tiers, with varying prices. Theoretical speeds can reach as high as 150 Mbps downstream and 25 Mbps upstream.

CELLULAR TECHNOLOGIES

The WAN technologies discussed in the previous sections are all designed to supply Internet access to homes and places of business, but the fastest growing types of wide area networking are those that deliver the Internet to people's pockets and purses. These are the wireless WAN services, most of them based on cellular communications, that enable smartphones, tablets, and other mobile devices to stay connected wherever they go.

EXAM TIP Cellular communications is a highly sophisticated area of the telecommunications field, and candidates for the Network+ exam are not expected to be experts in the field. However, they are expected to be familiar with some of the major technologies, including WiMAX and LTE.

There are literally dozens of standards used in cellular communications, which can be roughly categorized as applying to the first (1G), second (2G), third (3G), or fourth (4G) generation of mobile telecommunications. The industry is currently in the midst of a transition from 3G to 4G.

To be characterized as a 3G technology, a standard must conform to specifications published by the International Telecommunications Union (ITU) in a document called "International Mobile Telecommunications-2000 (IMT-2000)." These standards define all types of mobile communications, including voice telephone, not just Internet access.

One of the most prominent of the 3G standards is *High Speed Packet Access* (HSPA). HSPA supports downstream transfer rates of up to 14 Mbps and often requires only a software upgrade to 3G networks. *Evolved HSPA* (also known as HSPA+), which appeared in 2008 and is now widely implemented, can boost downstream data rates to 84 Mbps, by using MIMO technology.

To be called 4G, standards must conform to an ITU document called "International Mobile Telecommunications – Advanced (IMT-Advanced)." Two of the most prominent 4G standards are as follows:

- **Worldwide Interoperability for Microwave Access (WiMAX)** Based on the 802.16 standard published by the IEEE, WiMAX is a metropolitan area network (MAN) standard offering transfer rates up to 75 Mbps for mobile devices and ranges up to 50 kilometers. Marketed as a wireless last-mile alternative to DSL and CATV networks, WiMAX is now available in some markets in a variety of devices, including smartphones and external modems.
- **Long Term Evolution (LTE)** LTE is the next iteration of the GSM technology that first appeared in the second generation (2G) and which became HSPA in the third (3G). Although not yet compliant with the ITU standard, LTE is generally considered a 4G technology and supports downstream transmission rates of up to 300 Mbps.

True or false: PSTN, ISDN, and DSL are all WAN services that make use of the standard telephone network.

Answer: *True*. Using the Public Switched Telephone Network as a WAN solution requires dial-up modems at both ends. Integrated Services Digital Network and Digital Subscriber Line are both services that use special equipment to send high speed digital traffic over the standard telephone network.

True or false: There is virtually no difference between SONET and SDH.

Answer: *True*. Although there are slight differences between the SONET and SDH signaling methods, the two standards describe technologies that are, for all intents and purposes, the same.

EXAM TIP The Network+ exam requires candidates to have a general knowledge of technologies such as SONET, but the average network administrator is not likely to have any firsthand experience with them. SONET links form the backbone of the Internet and other global telecommunications networks, but they are not the type of WAN technologies that an organization is likely to install for its private use.

True or false: CATV, ADSL, and ISDN are all asymmetrical services.

Answer: *False*. CATV and ADSL are asymmetrical, meaning that they provide more downstream bandwidth than upstream. However, ISDN is a symmetrical service.

True or false: With frame relay, you can use a single connection to a frame relay provider to replace several dedicated leased lines.

Answer: *True*. Using frame relay, you can create a mesh WAN topology connecting the networks at multiple sites by using a single leased line at each location to connect to a common cloud. Because the connections in a frame relay cloud are ephemeral, a single network can simultaneously establish multiple permanent virtual circuits to different destinations.

Categorize WAN technology properties

Selecting a WAN protocol for a particular application requires you to compare their properties. Some of these properties are discussed in the following sections.

Packet-switching and circuit-switching

When multiple computers have to share a single network medium, they transmit their data in small, discrete units called packets. Instead of transmitting an entire file all at once, the protocols running on the computer break it down into packets and transmit them to the destination individually. This way, many computers can gain access to the network and take turns transmitting packets.

Because computers on a network have to break up their transmissions into separate packets, it is conceivable that the packets that compose a single file might take different routes to their destination, and might even arrive at the destination out of order. The destination system must therefore identify the incoming packets and reassemble them in the proper order to recreate the original data transmitted by the sender. This type of arrangement is known as a packet-switching network.

SONET, DSL, frame relay, CATV, satellite, and cellular are all packet-switching WAN technologies.

The opposite of a packet-switching network is a circuit-switching network, in which one system opens a circuit (or path) to another system prior to transmitting any data. The circuit then remains open for the duration of the data transaction. Examples of WAN technologies that use circuit switching are dial-up PSTN connections, ISDN, and leased lines, which are circuits that stay open perpetually. Circuit-switching is not suitable for LANs, because it would monopolize the network medium for long periods.

Speed

For virtually all WAN technologies, speed and cost are directly proportional. The additional bandwidth is there, if you want to pay for it. To the consumer, the main difference between the technologies is how easy or difficult it is to get it.

Leased lines are available at speeds ranging from a fraction of a T-1 to a T-4 at over 274 Mbps. However, changing speeds is a difficult and expensive process, involving the removal of one line and the installation of another.

For most of the Internet access technologies, WAN service is available at various speeds. Most DSL and CATV providers now offer several grades of service, with changes as easy as calling the provider.

Transmission media

For all of the WAN technologies, there are still only four transmission media available, as follows:

- **Coaxial** Rarely used today, coaxial cable has existed since the nineteenth century and was once the dominant medium for long distance connections. Although relatively slow by today's standards, it still can provide longer transmission distances than twisted pair, though not nearly as long as fiber optic.
- **Twisted pair** Faster than coaxial, but slower than fiber optic, twisted pair is the most popular medium for LANs, but is less prevalent for WAN use. Twisted pair distances are relatively short and require frequent repeaters, but its ease of installation and low price make it suitable for smaller installations.
- **Fiber optic** Fiber optic cable can provide higher speeds and longer runs than any other medium, and it also offers protection against EMI and additional security for the data. The drawback is the additional cost for the raw materials and the installation. Nevertheless, most of the new WAN media installed today use this type of cable.
- **Wireless** Slower than any other WAN medium, and spanning shorter distances, the wireless data networking service provided by cellular carriers nevertheless offers a degree of convenience that no other medium can approach.

Distance

Because they use the PSTN for their connections, leased lines are not limited with respect to the distance they can span, but you will pay more for the service as you increase the distance between the connected points. SONET and SDH are designed specifically for long distance fiber optic runs that can be many kilometers long.

ISDN and DSL are not limited in the distance that their signals can travel, but they both have a maximum distance from the nearest telephone company point of presence. If your site is not located within that maximum, you are not eligible to receive the service.

Cellular connections, as with all wireless technologies, are subject to many types of interference, which can weaken their signals. In the same way, as the distance from the nearest cell tower increases, the weaker the signal gets.

True or false: Coaxial cables are capable of spanning longer distances than twisted pair cables.

Answer: *True*. The longest copper link ever allowed on an Ethernet network was a 500 meter coaxial cable called Thick Ethernet. Twisted pair cables are limited to 100 meters on Ethernet.

True or false: The medium on a packet switching network can contain packets from many different systems, mixed in any order.

Answer: *True*. Computers on a network with a shared medium take turns transmitting packets, so the packets on the cable can be in any order.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is the technical name for the standard analog telephone network?
2. Which ISDN service is commonly known as 2B+D?
3. What is the name of the device at both ends of a T-1 connection?
4. What type of cable does a SONET network use at the physical layer?
5. Which DSL type is most commonly used to provide Internet access to residential end users?

Objective 3.5: Describe different network topologies

Selecting the correct cable type is certainly an important factor in the network deployment process, but after you have the cables, it's not as though you can simply start connecting computers together with no thought to a pattern. The physical topology of a network is the pattern you use to connect the computers and other devices together.

When deploying a LAN, the topology you use is directly related to the type of cable you select. Each physical layer specification for data-link layer LAN protocol

has installation requirements that you must observe if the network is to function properly, and the topology is one of those requirements.

A WAN solution for a large network with offices at multiple distant sites can use any one of several topologies, depending on the organization's communication requirements and of course its budget. This exam objective requires Network+ candidates to know the properties of the topologies most commonly used for LANs and WANs.

Exam need to know

- Describe MPLS
For example: How does MPLS ease the burden on routers?
- Describe the point-to-point topology
For example: What types of networks use a point-to-point topology?
- Describe the point-to-multipoint topology
For example: Do wired networks use a point-to-multipoint topology?
- Describe the ring topology
For example: Why do Token Ring networks have hubs if they use a ring topology?
- Describe the star topology
For example: How does a star topology provide fault tolerance?
- Describe the mesh topology
For example: Can a LAN use a mesh topology?
- Describe the bus topology
For example: Why does Ethernet no longer use the bus topology?
- Describe the peer-to-peer topology
For example: Which operating systems use the peer-to-peer model?
- Describe the client-server topology
For example: Are there any operating systems that use the client-server model?
- Describe the hybrid topology
For example: What topologies are combined to form hybrids?

Describe MPLS

Multiprotocol label switching (MPLS) is a packet forwarding technology that is designed to lessen the burden on intermediate routers. On a typical internetwork, packets are forwarded from router to router and at each hop, the router must read the destination address, compare it to its routing table, and decide where to send it next. With MPLS, one router analyzes the packet and inserts a short label into the packet header. The label identifies a complete channel through the internetwork to the destination, so subsequent routers do not have to repeat the analysis.

True or false: Because MPLS is a multiprotocol technology, it can be used with any packet-switching protocol, including IP, frame relay, or ATM.

Answer: *True*. MPLS is defined by some administrators as working at layer 2.5 of the OSI model, between the data-link and network layers. Therefore, the data carried in the packets can be generated by any protocol.

EXAM TIP MPLS is more of a WAN technology than a physical topology, and does not rightly belong in this particular Network+ exam objective.

True or false: Because each router must read the MPLS label in every packet, the efficiency increase the technology provides is marginal.

Answer: *False*. Reading the MPLS label is a far less labor-intensive process than reading the packet's destination address and looking it up in the routing table. Therefore, the increase in efficiency MPLS supplies is substantial.

Describe the point-to-point topology

No topology is simpler than the point-to-point topology, in which one computer is directly connected to another, as shown in Figure 3-13. Obviously, this topology limits the network to two computers.

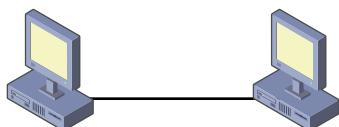


FIGURE 3-13 A point-to-point topology connecting two computers.

It is possible to connect two computers together by using a single Ethernet cable, creating the simplest possible LAN—but this requires a special cable called a crossover cable. In most cases, even two-node LANs connect by using a switch or a hub instead.

Point-to-point connections are also possible with wireless networks. In wireless LAN parlance, two computers connected directly together by using wireless network interface adapters form what is called an ad hoc network.

The point-to-point topology is more commonly found in wide area networking, which consists of all point-to-point links. For example, when you connect your home computer to the Internet, you are establishing a point-to-point WAN connection between your computer and a router on your Internet service provider's network. Corporate networks also use point-to-point WAN links to connect LANs in remote offices together.

True or false: Point-to-point connections require no addresses in their packets.

Answer: *True*. When there are only two devices on the network segment, then all of the packets each one transmits are obviously destined for the other, so there is no need to include an address in each packet.

MORE INFO For more information on the Point-to-Point protocol, the data-link layer protocol used for most point-to-point connections, see “Objective 5.2: Explain the methods of network access security.”

True or false: LANs rarely use point-to-point connections.

Answer: *True*. LANs nearly always use a topology that allows for more than two devices. Two-node LANs using a point-to-point topology are usually just temporary arrangements.

Describe the point-to-multipoint topology

In a point-to-point topology, each node transmits and the other node receives. The alternative to this model is the point-to-multipoint topology, in which a single node transmits and multiple nodes receive the data, as shown in Figure 3-14. All of the other topologies discussed in the other sections are essentially variations on these two models.

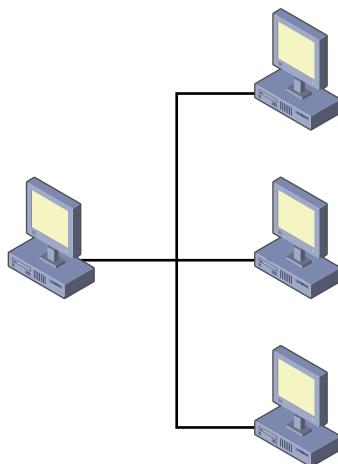


FIGURE 3-14 A point-to-multipoint topology.

The star, mesh, and logical ring topologies, with their variants, are all basically collections of point-to-point links. The only wired LAN topology that uses the point-to-multipoint model is the bus, because the network propagates the signals transmitted by each node to every other node without any intervening switch, hub, MAU, or other device. The other main example of the point-to-multipoint topology on LANs is in wireless networking, where a single system can transmit to multiple destinations simultaneously.

True or false: The only point-to-multipoint topology currently in general use involves wireless networking.

Answer: *True*. Wireless LANs use a point-to-multipoint topology because radio signals, by nature, are transmitted to all receiving stations in range. With a wired

network, switches now enable Ethernet transmissions to use what amounts to a point-to-point topology.

True or false: Ethernet networks using coaxial cable employ a point-to-multipoint topology.

Answer: *True*. Coaxial Ethernet uses a bus topology, which transmits all signals from each node to every other node. Describe the ring topology

A ring topology is essentially a bus with the two ends joined together, as shown in Figure 3-15, so that a signal transmitted by a computer in one direction circulates around the ring, eventually ending up back at its source. The data-link layer protocol traditionally associated with the ring topology is the now-obsolete Token Ring protocol.

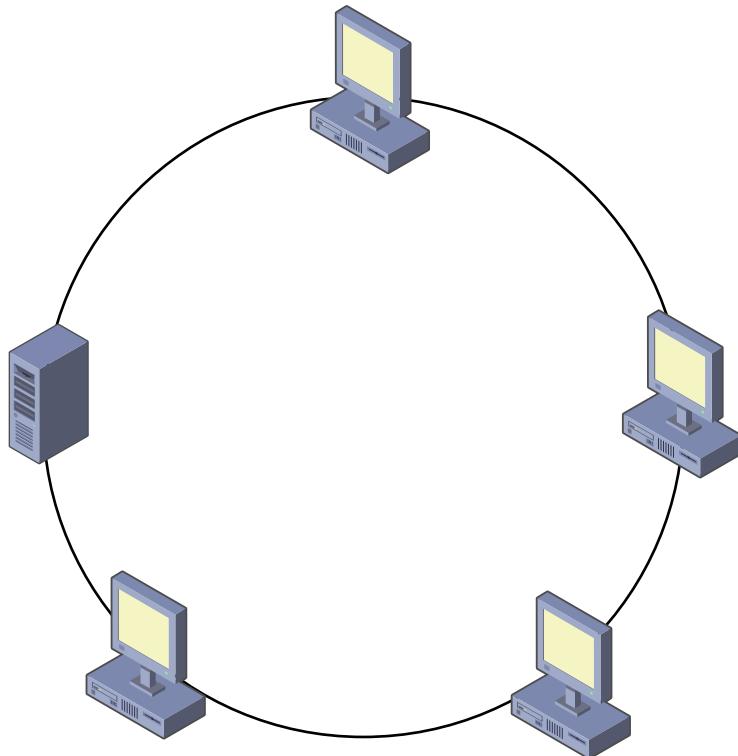


FIGURE 3-15 A LAN using a ring topology.

Token Ring networks use a MAC mechanism called token passing, which requires a special packet called a token to circulate endlessly around the network. Only the computer in possession of the token has permission to transmit data.

Early Token Ring networks were cabled together in an actual physical ring, with each computer connected to the next. However, Token Ring was also the origin

of the hybrid topology, in which a single network contained the attributes of two different topologies.

In the majority of Token Ring networks, the ring is a logical topology implemented in the wiring of the network. Physically, most Token Ring networks take the form of a star topology, with a cable running from each computer to a central cabling nexus called a multistation access unit (MAU). The MAU implements the logical ring by transmitting signals to each node in turn and waiting for the node to send them back before it transmits to the next node. Thus, although the cables are physically connected in a star, the data path takes the form of a ring, as shown in Figure 3-16. This is sometimes referred to as a star-ring topology.

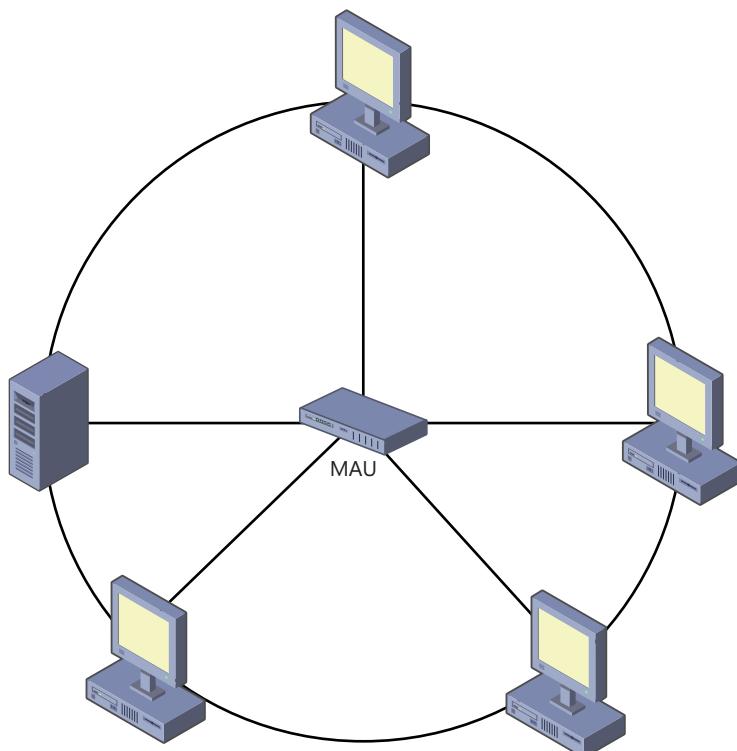


FIGURE 3-16 A LAN using a logical ring topology.

True or false: Implementing the ring topology logically, inside a Token Ring MAU, enables the network to be more fault tolerant.

Answer: *True*. The physical star topology of the logical ring makes it possible for the network to function even when a cable fails. The MAU contains circuitry that can remove a malfunctioning workstation from the ring but still preserve the logical topology. By comparison, a network that is literally cabled as a ring would have no MAU, so a cable break would cause the network to stop functioning completely.

MORE INFO Token Ring was once included in the Network+ exam objectives, as an alternative to Ethernet at the data-link layer. The protocol has now been removed from the objectives, as has the other ring protocol, Fiber Distributed Data Interface (FDDI)

True or false: Ethernet was at one time able to use the ring topology.

Answer: *False*. The original Ethernet standards called for a bus topology, and later specifications required a star topology. There was never an Ethernet ring specification.

Describe the star topology

The most common LAN topology in use today is the star topology, in which each computer or other device is connected by a separate cable run to a central cabling nexus, which can be either a switch or a hub, as shown in Figure 3-17. Compared to the other LAN topologies, the star affords a good deal of fault tolerance. The failure of a cable affects only the connection of one node to the network. The switch or hub does provide a central point of failure, however, which can bring down the entire network, but failure of these devices is comparatively rare.

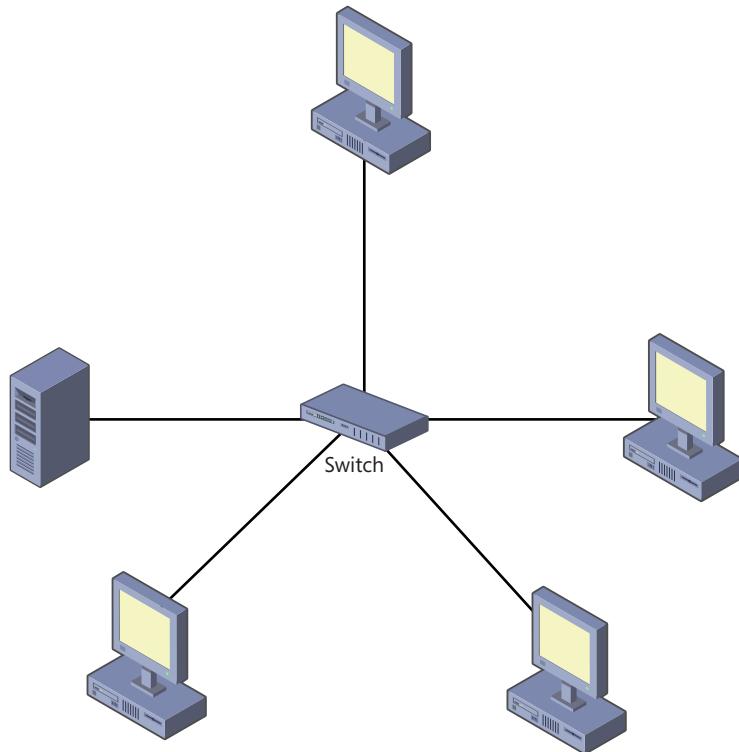


FIGURE 3-17 A LAN using a star topology.

All of the UTP-based Ethernet networks in use today call for the star topology, as do many fiber optic-based LAN protocols. In addition, many of the obsolete LAN

protocols that called for ring topologies in theory actually used a physical star; the ring was implemented logically, inside a specially designed hub.

A basic star network can be as large as the number of ports available in the switch or hub. When a switch or hub is fully populated, it is possible to expand the network further by adding another switch or hub and connecting it to the first. This creates what is known as a hierarchical star topology (or sometimes a branching tree topology), as shown in Figure 3-18. You can continue to expand the network in this way, within certain limits specified in the Ethernet standards.

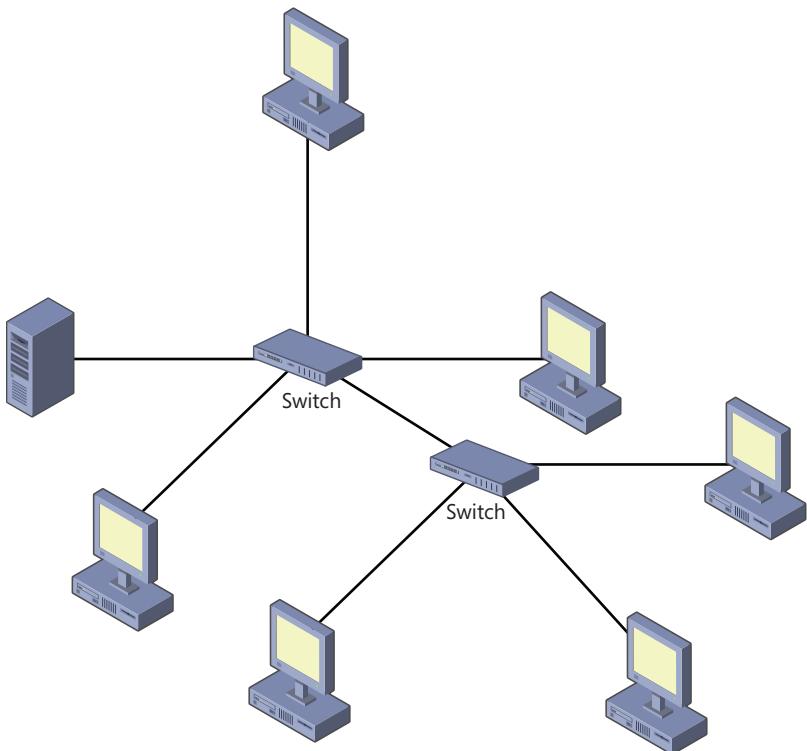


FIGURE 3-18 A LAN using a hierarchical star topology.

True or false: In a star topology, when one cable breaks, the entire network fails.

Answer: *False*. In a star topology, a broken cable only affects the computer connected with that cable.

EXAM TIP By far, the most commonly used cable topology in local area networking is the star topology and its variant, the hierarchical star. The bus and ring topologies are practically obsolete in the LAN world.

True or false: All of the Ethernet UTP and fiber optic specifications call for a star topology.

Answer: *True*. The only Ethernet specifications that did not use the star topology were the 10Base5 and 10Base2 specifications using coaxial cable.

Describe the mesh topology

Data-link layer LAN protocols do not provide much flexibility when it comes to choosing a network topology. The characteristics of specific cable types impose exacting limitations on how you can install them. However, when you are designing an internetwork topology—such as when you are installing multiple LANs at one location, building a large enterprise infrastructure in a data center, or connecting remote networks with WAN links—you have a lot more freedom.

To connect networks together, forming an internetwork, you use routers, and the topology you create when linking networks together is not subject to the same restrictions as the creation of a LAN with a specific cable type.

For example, if your organization has several branch offices located around the country, you can install WAN links to connect them in any way you see fit. By connecting each branch office to the company headquarters, you are using point-to-point connections to build a star WAN topology, as shown in Figure 3-19. In the same way, you could create a bus or a ring topology, although these models are rare in WAN implementations due to their lack of fault tolerance.



FIGURE 3-19 A star topology built by using WAN point-to-point connections.

Another option, in these circumstances, is to create some form of mesh topology, in which each office is connected to two or more other offices. This can be an expensive solution, but its redundancy enables the network to continue functioning even in the case of the failure of one or more links. There are two types of mesh topology: a partial mesh, in which each site has a point-to-point link to at least two

other sites, and a full mesh, in which each site is connected to every other site, as shown in Figure 3-20.



FIGURE 3-20 A full mesh topology built by using WAN point-to-point connections.

In both types of mesh topologies, there are at least two routers on each network. In the event that a router is unable to send data to a specific destination, it can forward the data through another router for transmission by using a different pathway.

True or false: A full mesh topology is impossible on a LAN.

Answer: *False*. On a LAN, a full mesh topology is not impossible, but it is certainly impractical. To create the mesh, each computer would require a separate network interface for every other computer.

EXAM TIP The Network+ objectives include the mesh topology in with the standard LAN cable topologies, but the mesh is possible only in a WAN internetwork configuration.

True or false: A mesh topology is more fault tolerant than a star topology.

Answer: *True*. Because there are multiple connections between each pair of points, a mesh topology can survive a link failure.

Describe the bus topology

A bus topology is one in which each computer is connected to the next one in a line. The first two Ethernet standards called for coaxial cable in a bus topology. The first, Thick Ethernet, used a single length of cable up to 500 meters long, with individual transceiver cables connecting the computers to the main trunk. Thin Ethernet used

the smaller RG-58 cable in lengths, attaching the T connectors on the computers' network interface adapters, forming a bus up to 185 meters long, as shown in Figure 3-21.

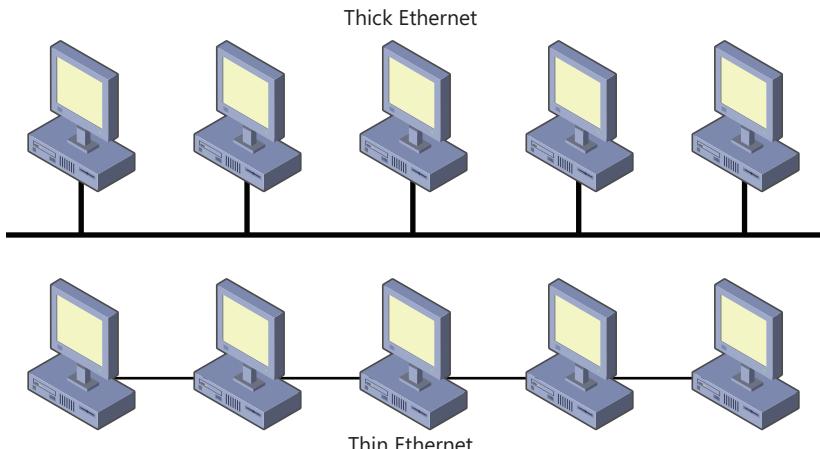


FIGURE 3-21 A LAN using a bus topology.

On a bus network, when any computer transmits data, the signal travels down the cable in both directions. Both ends of the bus must be terminated with resistor packs that negate the signals arriving there. On a bus that is not properly terminated, signals reaching the end of the cable tend to echo back in the other direction and interfere with any newer signals that the computers transmit. This condition is called an impedance mismatch.

The only Ethernet networks that used the bus topology were those wired with coaxial cable. The decline of Ethernet on coaxial cable also meant the decline of the bus topology, which is no longer used today.

True or false: In a bus topology, when one cable breaks, the entire network fails.

Answer: *True*. The inherent weakness of the bus topology is that a single cable failure can disrupt communications for the entire network. A broken cable splits the bus into two halves, preventing the nodes on one side from communicating with those on the other. In addition, both halves of the network are left with one end in an unterminated state, which prevents computers on the same side of the break from communicating effectively.

EXAM TIP The characteristics of the bus topology that Network+ exam candidates must know are that it requires termination at both ends, and that it is much less fault tolerant than the star.

True or false: Ethernet networks using a bus topology do not require hubs or switches.

Answer: *True*. On a bus network, the cable runs from computer to computer. There is no need for an intervening hub or switch.

Describe the client/server topology

The basic functions of a network typically involve one computer or other device providing some kind of service to other computers. This relationship is typically referred to as client/server networking. The server side of the partnership can be a computer that provides storage, access to a printer, email services, webpages, or any number of other services. The client is a computer running a program that accesses the services provided by servers.

In the early days of PC networking, these client and server roles were more clearly defined than they are today. Servers were computers dedicated exclusively to server functions; they could not function as clients.

NOTE One of the few successful network operating system products that operated on a strictly client/server model was Novell NetWare. All Windows, UNIX, and Linux operating systems are capable of using the peer-to-peer model.

True or false: Microsoft was the first major software manufacturer to market a server operating system with a graphical user interface.

Answer: *True*. Prior to Windows NT, servers all used character-based interfaces to conserve system resources. Windows NT was the first server that had a fully functional graphical interface.

EXAM TIP The Network+ exam objectives group client/server and peer-to-peer in with the physical layer topologies, but they are more logical topologies than physical.

True or false: Today, the terms client and server refer to applications more than they do to computers.

Answer: *True*. Few, if any, computers are devoted exclusively to client or server functions. Instead, they can perform a variety of tasks that qualify as client and server applications.

Describe the peer-to-peer topology

Today, virtually all of the computers on a network are capable of functioning as both clients and servers simultaneously, and their roles are more a matter of the administrator's choice than the operating system running on the computer. This relationship is known as peer-to-peer networking. On a peer-to-peer network, for example, a computer can share its drives with the rest of the network and can also access shared drives on other computers, regardless of whether the system is running a server or a client operating system.

Manufacturers of operating systems still tend to market separate server and client versions of their products, but a computer running a server operating system can still function as a client. Many home or small business networks consist solely of

computers running client operating systems, which can also function as servers at the same time.

This might seem confusing, but suffice it to say that in today's computing world, the terms "client" and "server" refer not so much to machines or operating systems as they do to the roles or applications running on those machines or operating systems.

True or false: Client operating systems can run on server computers, and server operating systems can run on client computers.

Answer: *True*. Computer manufacturers design systems specifically for use as servers and clients, but there is nothing about the hardware to prevent someone from installing a server operating system on a client or vice versa.

Describe the hybrid topology

A hybrid topology is one in which two or more of the topologies already described are found on a single network. For example, a star bus topology is one in which multiple star networks are connected by joining their hubs or switches to a bus network. Some FDDI networks use a hybrid star ring topology, with hubs connected to a ring or a double ring.

True or false: In a hybrid topology, computers must have two network interfaces—one to connect to each topology.

Answer: *False*. In a hybrid topology, the computers connect to one type of network and the hubs or switches connect to another. There is no need to connect a computer to two different network topologies.

EXAM TIP The Network+ exam objectives do not define the term hybrid, which can also describe a topology that mixes wired and wireless networks.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What change in network topology must occur to upgrade an existing Thin Ethernet network to Fast Ethernet using UTP cable?
2. An electrician installing a light fixture accidentally severs one of the LAN cables running through the dropped ceiling space. With which topology would the severed cable cause the most amount of disturbance to the network?
3. Which of the physical topologies is used by the majority of new Ethernet networks installed today?
4. Which of the physical network topologies requires the installation of two terminating resistors?

Objective 3.6: Given a scenario, troubleshoot common physical connectivity problems

Installing internal cable runs is a complex and detailed process, and even the most seasoned professionals can make mistakes. It is therefore essential that the installers test each cable run to ensure that it functions properly.

Plugging the two ends of a cable run into a computer and a switch by using patch cables is not an adequate form of testing. There are many types of cabling errors that can cause intermittent performance problems that are not immediately evident. The computer can function normally at first, and only display the symptoms of a cabling problem later.

For this objective, Network+ exam candidates must be familiar with the various problems that can affect cable installations and how to detect and repair them.

Exam need to know

- Troubleshoot cable problems

For example: What equipment can you use to detect bad connectors, bad wiring, open and short circuits, split cables, DB loss, TX-RX reversals, cable placement problems, EMI and interference, excessive distance, and crosstalk?

Troubleshoot cable problems

The first and most essential test that installers must perform on every cable run is a continuity test, which is simply a test for bad wiring that ensures that each wire on both ends of the cable is connected to the correct pin, and only the correct pin. A wire that is not connected to both ends of the cable, either because the connection itself is flawed or because the wire is broken somewhere in the cable, generates a condition called an open circuit. If a pin on one end of a cable run is connected to two or more pins on the other end, you have a short circuit.

It is possible to test cable continuity by using a tone generator and locator or a standard electrical multimeter. However, identifying and testing each cable thoroughly using these tools is a lengthy process; professional installers rarely use them. For a few dollars more, you can purchase a device called a *wiremap tester*, which functions on the same principles but connects to all eight wires at once on both ends and tests them at the same time. A wiremap tester can detect opens and shorts, as well as transposed wires, or TX-RX reversals. However, it cannot detect split pairs. This usually requires a more sophisticated testing device.

A split pair, also called a split wire, is a connection in which two wires are incorrectly mapped in exactly the same way on both ends of the cable. In a properly wired connection, each twisted pair of wires should contain a signal and a ground. In a split pair, you can have two signal wires twisted together as a pair. This can generate excessive amounts of crosstalk, corrupting both of the signals involved. Crosstalk is the bleeding of signals from one wire to another, causing data corruption.

UTP cables have length limitations because signals tend to lose strength (or attenuate) over distance. This condition is sometimes called DB loss. Installers should be conscious of their cable run lengths as they install them, but testing the actual length of the installed cable runs is always recommended. Higher-end cable testing tools have the ability to determine the length of a cable run, locate a break in any one of the wires, and specify the location of the break, in terms of the distance from the cable end.

The ANSI/TIA-568-C standard rates UTP cables by specifying maximum acceptable levels of various types of interference. As you move up to the higher cable categories, there are more tests to conduct, and more stringent limitations to meet. Testing these characteristics in the field therefore requires more sophisticated (and more expensive) equipment.

Cable testing devices at this level of performance are typically called cable certifiers, scanners, or media testers. These are handheld devices that you simply connect to the cable ends, and the certifier runs through a battery of tests for different types of interference, whether caused by the placement of cables near an EMI source or anything else.

Some of the tests that cable certifiers can typically perform include the following:

- **Attenuation** The signal lost along a cable's length as it is transmitted.
- **Near-End Crosstalk (NEXT)** The strength of the signal that bleeds over into the other wires near the end of the cable where the transmitter is located.
- **Far-End Crosstalk (FEXT)** The strength of the signal that bleeds over into the other wires near the far end of the cable.
- **Power Sum NEXT (PS-NEXT)** The crosstalk that is generated when three of the four wire pairs are carrying signals at once.
- **Equal-Level Far-End Crosstalk (EL-FEXT)** The crosstalk at the opposite end of the cable from the transmitter, corrected to account for the amount of attenuation in the connection.
- **Power Sum EL-FEXT (PS-ELFEXT)** The crosstalk that is generated at the far end of the cable by three signal-carrying wire pairs, corrected for attenuation.
- **Alien Crosstalk (AXT)** The measurement of the signal bleedover from a wire pair in one cable to the same pair in an adjacent cable.
- **Propagation delay** The amount of time required for a signal to travel from one end of a cable to the other.
- **Delay skew** The difference between the lowest and the highest propagation delay measurements for the wires in a cable.
- **Return Loss (RL)** Return Loss measures the accumulated signal reflection caused by variations in the cable's impedance along its length. These impedance variations are typically caused by untwisting the wires when making connections.

True or false: Fiber optic cable is less susceptible to attenuation than twisted pair cable.

Answer: *True*. Fiber optic cable, especially singlemode, can span much longer distances because its signals are much less prone to attenuation.

EXAM TIP The Network+ exam objectives split the physical layer troubleshooting process covered here from the coverage of the physical layer troubleshooting tools, found in “Objective 4.2: Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues.”

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is the term that describes the weakening of a signal as it travels through a cable’s length?
2. Which two types of interference are measured on three of the four signal-carrying wire pairs?
3. Which of the following types of wiring faults cannot be detected by a wiremap tester?
4. What is the name of the resulting fault, when one of the eight wires in a UTP cable is broken inside the sheath?
5. Which type of crosstalk involves wire pairs in two separate cables?

Objective 3.7: Compare and contrast different LAN technologies

Any discussion of the history of the local area network (LAN) is also a discussion of the history of Ethernet. Ethernet was the first LAN protocol, originally conceived in 1973, and it has been evolving steadily ever since. Other data-link layer protocols—such as Token Ring and ARCnet—have come and gone, but Ethernet has remained.

The longevity of Ethernet is due primarily to its continual evolution. The earliest commercial Ethernet networks ran at 10 Mbps (megabits per second), and successive iterations of the protocol increased the networks’ transmission speeds to 100, 1,000, and now 10,000 Mbps. For this exam objective, you must be familiar with the basic properties of an Ethernet LAN, and with many of the Ethernet physical layer specifications.

Exam need to know

- Compare and contrast LAN types

For example: What are the defining characteristics of the various Ethernet types, such as Ethernet, 10BaseT, 100BaseT, 1000BaseT, 100BaseTX, 100BaseFX, 1000BaseX, 10GbaseSR, 10GBaseLR, 10GbaseER, 10GbaseSW, 10GbaseLW, 10GbaseEW, and 10GBaseT?

- Compare and contrast LAN properties

For example: How are the following terms significant to local area networking: CSMA/CD, CSMA/CA, broadcast, collision, bonding, speed, and distance?

Compare and contrast LAN types

This first Ethernet standard described a network that used RG-8 coaxial cable in a bus topology up to 500 meters long, with a transmission speed of 10 Mbps. This was commonly known as Thick Ethernet or 10Base5. The consortium published a second version of the standard, called DIX Ethernet II, in 1982. This version added a second physical layer specification, calling for RG-58 coaxial cable. This came to be known as Thin Ethernet or 10Base2. RG-58 cable is thinner and less expensive than RG-8 but more susceptible to interference and attenuation, so the maximum segment length is restricted to 185 meters. Development of the DIX Ethernet standard stopped after the publication of version II.

The original IEEE 802.3 standard retained the coaxial cable specifications from the DIX Ethernet document, but by 1993, the IEEE had published several amendments adding UTP and fiber optic cable specifications, also running at 10 Mbps. The 10 Mbps Ethernet physical layer specifications are listed in Table 3-7.

EXAM TIP 10Base5 and 10Base2, also known as Thick Ethernet and Thin Ethernet, are the only LAN protocols covered on the Network+ exam that use coaxial cable in a bus topology.

TABLE 3-7 10 Mbps Ethernet physical layer specifications.

DESIGNATION	CABLE TYPE	TOPOLOGY	MAXIMUM SEGMENT LENGTH
10Base5	RG-8 coaxial	Bus	500 meters
10Base2	RG-58 coaxial	Bus	185 meters
10Base-T	CAT3 UTP	Star	100 meters
FOIRL	62.5/125 multimode fiber optic	Star	1,000 meters
10Base-FL	62.5/125 multimode fiber optic	Star	2,000 meters

EXAM TIP The Network+ exam objectives contain some, but not all, of the physical layer specifications in the Ethernet standard. The IEEE 802.3 document also includes several other physical layer specifications that were either never implemented or never caught on in the marketplace. These specifications are not included in the tables in this chapter, and knowledge of them is not necessary for the exam.

In addition to cable types and segment lengths, the Ethernet physical layer specifications also limit the number of repeaters that are permitted in a network configuration. A repeater is a physical layer device that enables you to extend the

length of a network segment by amplifying the signals. The 5-4-3 rule states that an Ethernet network can have as many as five cable segments, connected by four repeaters, of which three segments are mixing segments (meaning they contain hosts).

Fast Ethernet

The 100Base-TX specification retains the 100-meter maximum segment length from 10Base-T, as well as the use of two wire pairs. However, to support the higher transmission speeds, the standard calls for a higher grade of cable: CAT5 instead of CAT3.

To provide a direct upgrade path to Fast Ethernet on existing CAT3 cable installations, the IEEE also published a specification called 100Base-T4. This network did not require a cable upgrade, but to compensate, it made use of all four wire pairs instead of just two, and used a different signaling scheme. Together, the two UTP Fast Ethernet specifications are known as 100Base-T. However, for whatever reason, 100Base-T4 never caught on in the marketplace, and CAT5 quickly became the industry standard for UTP cable installations. Table 3-8 provides the Fast Ethernet physical layer specifications.

TABLE 3-8 100 Mbps Fast Ethernet physical layer specifications.

DESIGNATION	CABLE TYPE	MAXIMUM SEGMENT LENGTH
100Base-TX	CAT5 UTP	100 meters
100Base-FX	62.5/125 multimode fiber optic	412 meters (half duplex)/ 2,000 meters (full duplex)

EXAM TIP Although 100Base-T4 never succeeded in the marketplace, it is part of the 100Base-T specification, which is included in the Network+ exam objectives.

On the fiber optic side, 100Base-FX uses the same 4B/5B signaling method as 100Base-TX, and as a result, the two specifications are known collectively as 100Base-X. The maximum segment length on a 100Base-FX network depends both on the cable type and on the use of full duplex communications. Using standard multimode fiber optic cable and full duplex communication, a segment can be as long as 2,000 meters. Half-duplex signaling reduces the length to 412 meters. Using singlemode cable, a full-duplex segment can be 20 kilometers (km) long or more.

The 5-4-3 rule does not apply to Fast Ethernet networks. Fast Ethernet hubs are available in two classes:

- **Class I** Connects different types of Fast Ethernet cable segments together, such as fiber optic to UTP or 100Base-TX to 100Base-T4.
- **Class II** Connects Fast Ethernet cable segments of the same type together.

Each Fast Ethernet hub must be identified by the appropriate Roman numeral in a circle. You can have as many as two Class II hubs on a single LAN, with a total cable length (for all three segments) of 205 meters for UTP cable and 228 meters for

fiber optic cable. Because Class I hubs must perform an additional signal translation, which slows down the transmission process, you can have only one hub on the network, with maximum cable lengths of 200 and 272 meters for UTP and fiber optic, respectively.

The CSMA/CD media access control mechanism is the defining element of the Ethernet protocol, but it is also the source of many of its limitations. The fundamental shortcoming of the Ethernet protocol is that data can travel only in one direction at a time. This is known as half-duplex operation. With special hardware, it is also possible to run Ethernet connections in full-duplex mode, meaning that the device can transmit and receive data simultaneously. In point-to-point communications, this effectively doubles the bandwidth of the network and eliminates the media access control problem.

Gigabit Ethernet

The physical layer specifications for Gigabit Ethernet, running at 1,000 Mbps, appeared in 1998 and 1999. The fiber optic specifications, known collectively as 1000Base-X, were published as IEEE 802.3z, and the 1000Base-T UTP specification as IEEE 802.3ab. These specifications are listed in Table 3-9.

TABLE 3-9 1000 Mbps Gigabit Ethernet physical layer specifications.

DESIGNATION	CABLE TYPE	MAXIMUM SEGMENT LENGTH
1000Base-T	CAT5, CAT5e, or CAT6 UTP	100 meters
1000Base-LX	9/125 singlemode fiber optic	5000 meters
1000Base-LX	50/125 or 62.5/125 multimode fiber optic	550 meters
1000Base-SX	50/125 or 62.5/125 multimode fiber optic	500 meters / 220 meters
1000Base-CX	150-ohm shielded, balanced twinaxial copper cable	25 meters

Gigabit Ethernet once again increased network speeds tenfold, while retaining the same basic UTP configuration and segment length, making upgrades possible in many cases without the need for new cable installations. As with each previous Ethernet speed increase, Gigabit Ethernet increased the requirements for the UTP cables. 1000Base-T uses all four wire pairs, unlike 100Base-TX, and is more susceptible to certain types of crosstalk. To address this issue TIA/EIA created the CAT5e and CAT6 cable grades, which are designed to support Gigabit Ethernet communications.

NOTE The TIA also created an alternative Gigabit Ethernet physical layer specification called 1000Base-TX. Like 100Base-TX, 1000Base-TX uses only two of the four wire pairs in a UTP cable. However, to compensate for that lack of two wire pairs,

the 1000Base-TX specification requires CAT6 UTP cabling. Due perhaps to the cost of the cables, this specification has not been popular in the marketplace.

The 1000Base-X specifications include two fiber optic configurations, essentially long-distance and short-distance options, plus a unique, short-run copper alternative. 1000Base-LX is intended to be the long-distance option, supporting segment lengths up to 5 kilometers with singlemode fiber, and up to 10 kilometers with high-quality optics in a variant called 1000Base-LX10. Some specialized installations also use repeating equipment to create much longer links.

1000Base-LX is designed to be used by large carriers as a long-distance Ethernet backbone, so it is likely that the average network administrator will not ever work with it, nor will you find 1000Base-LX network interface adapters on the shelf at your local computer store.

The 1000Base-LX specification also allows for the use of multimode fiber optic cable at shorter distances, but 1000Base-SX is the shorter-distance specification designed for virtually any type of multimode cable. As with most fiber optic alternatives, 1000Base-SX works well as a link between buildings and on campus networks.

1000Base-CX is a copper specification calling for a twinaxial cable, which is a special shielded 150-ohm cable with two copper cores that uses either 9-pin D-shell or 8-pin Fibre Channel connectors. The maximum segment length is only 25 meters, making the specification good for links within data centers, such as equipment connections within server clusters, and little else. At the time the IEEE published the 802.3z document, 1000Base-CX was the only copper cable specification available, but the 1000Base-T specification appeared a year later, leaving 1000Base-CX as a marginalized technology.

10 Gigabit Ethernet

Once again, the call went out for more Ethernet bandwidth, and the IEEE 802.3 working group responded in 2002 with the first standards defining an Ethernet network running at 10 gigabits per second (Gbps), which is 10,000 Mbps. With 10 Gigabit Ethernet, the developers appear to have reached a turning point, because they have abandoned their previous devotion to backward compatibility. 10 Gigabit Ethernet networks support only four-pair, full-duplex communication on switched networks. Gone is the support for half-duplex communication, hubs, and CSMA/CD. However, the standard Ethernet frame format remains, and as with all of the previous Ethernet standards, there is a copper-based UTP solution that uses RJ-45 connectors and a 100-meter maximum segment length.

As with each of the previous Ethernet speed iterations, the 10 Gigabit Ethernet standards include a variety of physical layer specifications. Some of them have already fallen by the wayside, but the technology is still young enough that the marketplace hasn't yet completed the process of winnowing out the unsuccessful ones. Table 3-10 lists the most prominent of the 10 Gigabit Ethernet physical layer specifications defined by the IEEE standards.

TABLE 3-10 10 Gigabit Ethernet physical layer specifications.

DESIGNATION	CABLE TYPE	WAVE-LENGTH	MAXIMUM SEGMENT LENGTH
10GBase-T	CAT6/CAT6a	N/A	55 meters/100 meters
10GBase-SR	Multimode fiber optic	850 nm	26–400 meters
10GBase-LR	Singlemode fiber optic	1310 nm	10 kilometers
10GBase-ER	Singlemode fiber optic	1550 nm	40 kilometers
10GBase-SW	Multimode fiber optic	850 nm	26–400 meters
10GBase-LW	Singlemode fiber optic	1310 nm	10 kilometers
10GBase-EW	Singlemode fiber optic	1550 nm	40 kilometers

The designers of 10 Gigabit Ethernet intended it to be both a LAN and a WAN solution, and for many administrators, LAN means copper-based UTP cables with a 100-meter maximum segment length. It was not until 2006 that the IEEE published the 802.3an amendment, which defined the 10GBase-T specification, but they knew it had to be done.

Unfortunately, to support transmissions at such high speeds with copper cables, it was necessary to define a new set of cable performance standards. To support 100-meter segments, 10GBase-T requires CAT6a cable, which has an increased resistance to alien crosstalk (interference from signals on other, nearby cables). With standard CAT6 cable, 10GBase-T only supports cable segments up to 55 meters long. The standard does not support UTP cables below CAT6 at all.

The 10 Gigabit Ethernet specifications for fiber optic cable predate the copper and provide a wide variety of options for both LAN and WAN implementations. The possibilities of 10 Gigabit Ethernet as a WAN solution led the developers to create a separate set of specifications that utilize the existing Synchronous Optical Network (SONET) infrastructure to carry Ethernet signals.

As shown in the table, there are three pairs of fiber optic specifications that all begin with the “10Bbase” abbreviation. The first letter of the two-letter code that follows specifies the type and wavelength of the fiber optic cable. The second letter indicates whether the specification is intended for LAN use (“R”) or WAN use (“W”).

The “S” in the 10GBase-SR and 10GBase-SW specifications describes the short wavelength (850 nanometers) of the lasers used to generate the signals on the cable. As with most short-range fiber optic solutions, these specifications call for multimode cable. The maximum segment length depends on the exact cable the network uses. For example, the 62.5-micron multimode fiber commonly used on

FDDI networks (OM1) can only support segments up to 26 meters long. With the newly ratified OM4 cable, 10GBase-SR segments can be as long as 400 meters.

The 10GBase-LR and 10GBase-LW specifications use a long wavelength laser (1310 nm) and singlemode cables to achieve segment lengths of 10 kilometers. The extra-long wavelength of the 10GBase-ER and 10GBase-EW specifications can support segments up to 40 kilometers long.

None of the 10 Gigabit Ethernet physical layer specifications indicate the types of connectors the cables should use. The actual implementation is left up to the equipment manufacturers.

True or false: A Gigabit Ethernet network on UTP requires at least CAT6 cable.

Answer: *False*. A Gigabit Ethernet network can run on CAT6, CAT5e, or CAT5 UTP cable.

EXAM TIP The Network+ exam objectives include technologies that are dated and some that are all but obsolete, particularly in the area of Ethernet. Candidates must be familiar with these older technologies, even if they are rarely seen in the field anymore.

True or false: The 1000Base-X designation includes both fiber optic and copper specifications.

Answer: *True*. The 1000Base-X specifications include two fiber optic configurations, essentially long-distance and short-distance options, plus a seldom-used, short-run copper alternative, called 1000Base-CX.

Compare and contrast LAN properties

In the early implementations of Ethernet, its media access control (MAC) mechanism, called Carrier Sense Multiple Access with Collision Detection (CSMA/CD), was the single most defining characteristic of the network. For multiple computers to share a single network medium, it is critical for there to be an orderly means to arbitrate network access. Each computer must have an equal chance to use the network, or its performance will degrade.

Media access control is the main reason why Ethernet networks have such exacting physical layer specifications. If cable segment lengths are too long, or if there are too many repeaters on the network, the CSMA/CD mechanism does not function properly, causing access control to break down and systems to receive corrupt data.

Carrier sense

The name "Carrier Sense Multiple Access with Collision Detection" describes the successive phases of the media access control process. When a computer on an Ethernet network has data to transmit, it begins by listening to the network to see if it is in use. This is the carrier sense phase of the process. If the network is busy, the system does nothing for a given period and then checks again.

Multiple access

When the network is free, the computer transmits its data packet. This is called the multiple access phase, because all of the stations on the network are contending for access to the same network medium. Even though computers perform an initial check during the carrier sense phase, it is still possible for two systems on the LAN to transmit at the same time, resulting in a signal quality error (SQE) or, as it is more commonly known, a collision. For example, if Computer A performs its carrier sense, and Computer B has already begun transmitting but its signal has not yet reached Computer A, a collision will occur if Computer A transmits. When a collision occurs, both systems must discard their packets and retransmit them. These collisions are a normal and expected part of Ethernet networking; they are not a problem unless there are too many of them or the computers cannot detect them.

Collision detection

The collision detection phase of the transmission process is the most important part of the CSMA/CD process. If the systems cannot detect when their packets collide, corrupted data might reach a packet's destination system and be treated as valid. To avoid this, Ethernet networks are designed so that packets are large enough to fill the entire network cable with signals before the last bit leaves the transmitting computer. Ethernet packets must be at least 64 bytes long; systems pad out short packets to 64 bytes before transmission. The Ethernet physical layer specifications also impose strict limitations on the lengths of cable segments.

The amount of time it takes for a transmission to propagate to the farthest end of the network and back again is called the network's round trip delay time. A collision can occur only during this interval. After the signal arrives back at the transmitting system, that system is said to have "captured the network." No other computer will transmit on the network while it is captured because the source computer will detect the traffic during its carrier sense phase.

Ethernet computers on twisted pair or fiber optic networks assume that a collision has occurred if they detect signals on both their transmit and receive wires at the same time. If the network cable is too long, if the packet is too short (called a runt), or if there are too many hubs, a system might finish transmitting before the collision occurs and be unable to detect it.

NOTE It is possible for a collision to occur after the last bit of data has left the transmitting system. This is called a late collision, and it is an indication of a serious problem, such as a malfunctioning network interface adapter or cable lengths that exceed the physical layer specifications.

When a computer detects a collision, it immediately stops transmitting data and starts sending a jam pattern instead. The jam pattern alerts the other systems on the network that a collision has taken place, that they should discard any partial packets they may have received, and that they should not attempt to transmit any data until the network has been cleared. After transmitting the jam pattern, the system waits a specified period of time before attempting to transmit again. This is called the backoff period. Both of the systems involved in a collision compute the length

of their own backoff periods by using a randomized algorithm called truncated binary exponential backoff. They do this to try to avoid causing another collision by backing off for the same period of time.

Because of the way CSMA/CD works, the more computers you have on a network segment or the more data the systems transmit over the network segment, the more collisions occur. Collisions are a normal part of Ethernet operation, but they cause delays because systems have to retransmit the damaged packets. When the number of collisions is minimal, the delays aren't noticeable, but when network traffic increases, the number of collisions increases and the accumulated delays can begin to have a noticeable effect on network performance. You can reduce the traffic on the LAN by installing a bridge or switch, or by splitting the LAN into two networks and connecting them with a router.

The modern Ethernet

The IEEE 802.3 standard still includes the term “Carrier Sense Multiple Access with Collision Detection” as part of the document name, but the fact is that very few Ethernet networks actually use CSMA/CD anymore. The need for a media access control mechanism hinges on the use of a shared network medium. The early networks that used coaxial cable connected all of the computers to a single cable segment in a bus topology, and the first twisted pair networks used hubs to create a star topology.

A hub is essentially a multiport repeater. When a hub receives a signal through any of its ports, it transmits that signal out through all of the other ports, resulting in a shared network medium. These networks all required a MAC mechanism to arbitrate access to the network.

The big change came when switches began to replace hubs in the marketplace. On a switched network, unicast data arriving at a switch through one of its ports typically leaves through only one of its other ports—the one connected to the destination system. As a result, there is no shared network medium; each pair of computers has a dedicated connection and there is no need for further media access control.

Modern Ethernet variants also use full duplex connections between hosts, which means that computers can transmit and receive data at the same time. This also eliminates the need for media access control. Although the Ethernet standards still include CSMA/CD, for backward compatibility purposes, it is all but obsolete on today's networks.

CSMA/CA

For an IEEE 802.11 wireless network, as with all data-link layer protocols that use a shared network medium, the MAC mechanism is one of the primary defining elements. The 802.11 standard defines the use of a MAC mechanism called Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), which is a variation of the CSMA/CD mechanism used by Ethernet.

One of the characteristics of the wireless networks defined in 802.11 is that stations can repeatedly enter and leave the BSS because of their mobility and the vagaries of the wireless medium. Therefore, the MAC mechanism on a wireless network must be able to accommodate this behavior.

The carrier sense and multiple access parts of the CSMA/CA mechanism are the same as those of an Ethernet network. A computer with data to transmit listens to the network medium and, if the medium is available, begins transmitting its data. If the network is busy, the computer backs off for a specified interval and begins the listening process again.

As with Ethernet, the CSMA part of the process can result in collisions. The difference in CSMA/CA is that systems attempt to avoid collisions in the first place by reserving bandwidth in advance, by specifying a value in the Duration/ID field of the MAC frame or by using specialized control messages called request-to-send (RTS) and clear-to-send (CTS) messages.

The carrier sense part of the transmission process occurs on two levels, the physical and the virtual. The physical carrier sense mechanism is specific to the physical layer medium the network is using and is equivalent to the carrier sense performed by Ethernet systems. The virtual carrier sense mechanism, called a network allocation vector (NAV), involves the transmission of an RTS frame by the system with data to transmit, and a response from the intended recipient in the form of a CTS frame.

Both of these frames have a value in the Duration/ID field that specifies the amount of time needed for the sender to transmit the forthcoming data frame and receive an acknowledgment (ACK) frame in return. This message exchange essentially reserves the network medium for the life of this particular transaction, which is where the collision avoidance part of the mechanism comes in. Because both the RTS and CTS messages contain the Duration/ID value, any other system on the network receiving either one of the two messages observes the reservation and refrains from trying to transmit its own data during that time interval. This way, a station that is capable of receiving transmissions from one computer but not the other can still observe the CSMA/CA process.

In addition, the RTS/CTS exchange also enables a station to determine whether communication with the intended recipient is possible. If the sender of an RTS frame fails to receive a CTS frame from the recipient in return, it retransmits the RTS frame repeatedly until a preestablished timeout is reached. Retransmitting the brief RTS message is much quicker than retransmitting large data frames, which shortens the entire process.

To detect collisions, IEEE 802.11 uses a positive acknowledgment system at the MAC sublayer. Each data frame that a station transmits must be followed by an ACK frame from the recipient, which is generated after a CRC check of the incoming data. If the frame's CRC check fails, the recipient considers the packet to have been corrupted by a collision (or other phenomenon) and silently discards it. The station that transmitted the original data frame then retransmits it as many times as needed to receive an ACK, up to a predetermined limit.

Note that the failure of the sender to receive an ACK frame could be due to the corruption or nondelivery of the original data frame or the nondelivery of an ACK frame that the recipient did send in return. The 802.11 protocol does not distinguish between the two.

Broadcast and collision domains

An Ethernet network that connects devices together with a bus topology or a hub creates a shared network medium, and therefore a single collision domain. A collision domain is a group of network devices connected in such a way that if two devices broadcast at the same time, a collision occurs. Expanding the network with repeaters or additional hubs maintains the single collision domain. If you connect the segments with bridges, switches, or routers instead, you create multiple collision domains.

A bridge splits one LAN into two collision domains, and a switch into many collision domains, but it preserves a single broadcast domain. The broadcast domain of a network is the group of computers that will receive a broadcast message transmitted by any one of its members. Broadcasts are an essential element of Ethernet networking (as are collisions), but bridges and switches always forward all broadcast messages to any part of the network.

This is one of the fundamental differences between a bridge or a switch and a router. A router splits a LAN into two separate collision domains and two separate broadcast domains. A bridge or a switch creates separate collision domains but a single broadcast domain, which enables many of the LAN protocols to continue functioning.

Bonding

Link aggregation, also called bonding, is the ability to install multiple Ethernet network adapters in a computer and combine their bandwidth. If, for example, you have a group of users that routinely accesses large image files on a server, and that traffic is having a negative effect on the performance of the entire network, installing a second network interface adapter in the server and in each of the workstations can, under certain circumstances, be a far less expensive solution than upgrading the entire network to a faster speed.

EXAM TIP The Network+ objectives use the term bonding here to apply to Ethernet LANs, but the term also appears as channel bonding in regard to wireless networking, in "Objective 3.3: Compare and contrast different wireless standards."

Speed

The speed at which a device operates on an Ethernet LAN is based, of course, on the version of the 802.3 standard supported by its hardware. It is also dependent on the capabilities of the device at the other end of the cable. Backwards compatibility has always been a major priority with the designers of the IEEE 802.3 standards. Most of the Ethernet networking hardware on the market today enables a computer

to connect to the network at Gigabit Ethernet speed or negotiate a slower speed connection, if that is all the network supports.

The Fast Ethernet specifications define an optional autonegotiation system that enables a dual-speed device to sense the capabilities of the network to which it is connected and to adjust its speed and duplex status accordingly. The Gigabit Ethernet specifications expand the capabilities of the autonegotiation system, enabling devices to also communicate their port type and clocking parameters. In Gigabit Ethernet that uses copper cable, support for autonegotiation is mandatory.

When two Ethernet devices capable of operating at multiple speeds autonegotiate, they exchange fast link pulse (FLP) packets to determine the best performance level they have in common, and then configure themselves accordingly. The systems use the following list of priorities when comparing their capabilities, with full-duplex 1000Base-T providing the best performance and half-duplex 10Base-T providing the worst:

1. 1000Base-T (full-duplex)
2. 1000Base-T (half-duplex)
3. 100Base-TX (full-duplex)
4. 100Base-T4
5. 100Base-TX (half-duplex)
6. 10Base-T (full-duplex)
7. 10Base-T (half-duplex)

The benefit of autonegotiation is that it permits administrators to upgrade a network gradually with a minimum of reconfiguration. If, for example, you have 10/100/1000 multispeed network adapters in all your workstations, you can run the network at 100 Mbps using 100Base-TX switches. Later, you can simply replace the switches with models supporting Gigabit Ethernet, and the network adapters will automatically reconfigure themselves to operate at the higher speed during the next system reboot. No manual configuration at the workstation is necessary.

Distance

The distances that LANs can traverse are dependent on the media they use. With UTP cable, the matter is simple; 100 meters is the distance limit, whether you are running 10Base-T or 1000Base-T. With fiber optic cable, however, the distances are highly various. Beginning with the Fast Ethernet standard, the fiber optic specifications are split between multimode and singlemode cables, with singlemode providing much longer distances due to its narrower core and its laser light source. In addition, their resistance to EMI makes fiber optic cable suitable for outdoor use, joining buildings together on a campus network, for example.

Gigabit Ethernet began the practice of creating different designations for short haul and long haul segments, with several cable specifications to choose from. 10 Gigabit Ethernet expanded the selection further, with short, long, and extremely long designations, and separate varieties for LAN and WAN use—the intention being to expand the reach of Ethernet beyond the local network with links up to 40 kilometers long.

True or false: 10 Gigabit Ethernet is the first Ethernet version that does not use CSMA/CD.

Answer: *True*. 10 Gigabit Ethernet supports only four-pair, full-duplex communication on switched networks. You cannot use hubs or half-duplex communication, which eliminates the need for media access control and CSMA/CD.

EXAM TIP The description of an Ethernet collision makes it sound like a disastrous occurrence, but a modest number of regular collisions is normal on an Ethernet network and is no cause for concern. It is only massive numbers of collisions that are cause for alarm.

True or false: Wireless networks cannot use the same CSMA/CD mechanism as wired Ethernet networks because they cannot reliably detect collisions as they are happening.

Answer: *True*. To detect collisions, IEEE 802.11 uses a positive acknowledgment system at the MAC sublayer. If the sender of a packet does not receive an acknowledgment message from the recipient, it assumes that the packet was corrupted by a collision and retransmits it.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. On an Ethernet network using CSMA/CD, what prevents two systems that have already experienced a collision from colliding again?
2. To reduce the large amounts of broadcast traffic on your network, you decide to split it into two separate broadcast domains. What hardware changes must you make to do this?
3. What is the maximum cable segment length supported by 10Base-T, 100Base-TX, 1000Base-T, and 10Gbase-T Ethernet networks?
4. Which of the Ethernet physical layer specifications use all four pairs of the wires in the UTP cable?
5. Under what conditions is a UTP-based Ethernet network limited to 55 meter maximum segment lengths?

Objective 3.8: Identify components of wiring distribution

Originally called the ANSI/TIA-568 Commercial Building Telecommunications Cabling Standard, this document was revised in 1995, 2001, and 2009, and is now known as ANSI/TIA-568-C. The 568-C standard defines a structured cabling system for voice and data communications in office environments that has a usable life span of at least 10 years, supporting the products of multiple technology vendors and using twisted pair or fiber optic cable.

For each of the cable types, the standard defines the following elements:

- Cable characteristics and technical criteria determining the cable's performance level
- Installation guidelines, including topology, segment length, pull tension, and bend radius specifications
- Connector specifications and pinouts

The standard also includes specifications for the installation of the cable within a building space, which is called a wiring distribution. This objective requires exam candidates to be familiar with some of the technology used to construct the wiring distribution of a large enterprise network.

Exam need to know

- Identify an IDF in a wiring distribution
For example: What components do you find in an IDF?
- Identify an MDF in a wiring distribution
For example: What components do you find in an MDF?
- Identify a demarc in a wiring distribution
For example: Where is a demarc usually located?
- Identify a demarc extension in a wiring distribution
For example: Who is responsible for the demarc extension?
- Identify a smartjack in a wiring distribution
For example: What services does a smartjack provide, and to whom?
- Identify a CSU/DSU in a wiring distribution
For example: What kind of WAN connection requires a CSU/DSU?

Identify an IDF in a wiring distribution

The 568-C standard divides a building into the following subsystems:

- **Building entrance** The location where the building's internal cabling interfaces with outside cabling.
- **Telecommunications rooms and enclosures** Also known as the intermediate distribution frames (IDF), this is the location of localized telecommunications equipment such as the interface between the horizontal cabling and the backbone.
- **Equipment rooms** The location of equipment providing the same functions as that in a telecommunications room, but which might be more complex.
- **Backbone cabling** The cabling that connects the building's various equipment rooms, telecommunications closets, and the building entrance, as well as connections between buildings in a campus network environment.
- **Horizontal cabling** The cabling and other hardware used to connect the telecommunications rooms to the work areas.
- **Work area** The components used to connect the telecommunications outlet to the workstation.

True or false: In the wiring distribution for a large enterprise network, the intermediate distribution frames (IDFs) contain the demarcs for telephone switches and the WAN interfaces.

Answer: *False*. The IDFs contain the interfaces between the horizontal cabling and the backbone. The demarcs for telephone switches and the WAN interfaces are located in the MDF.

Identify an MDF in a wiring distribution

The components and processes described thus far in this chapter concern the installation of a single LAN. An enterprise network reaches well beyond a single LAN, however. Large corporate networks typically consist of multiple LANs that are joined together by another network, informally referred to as a backbone.

Picture an office building with multiple floors, each of which has its own LAN, consisting of a telecommunications room and horizontal cabling that leads out to the work areas. Picture another network running vertically through the building, connecting all of the telecommunications rooms together, as shown in Figure 3-22. This is the backbone network.

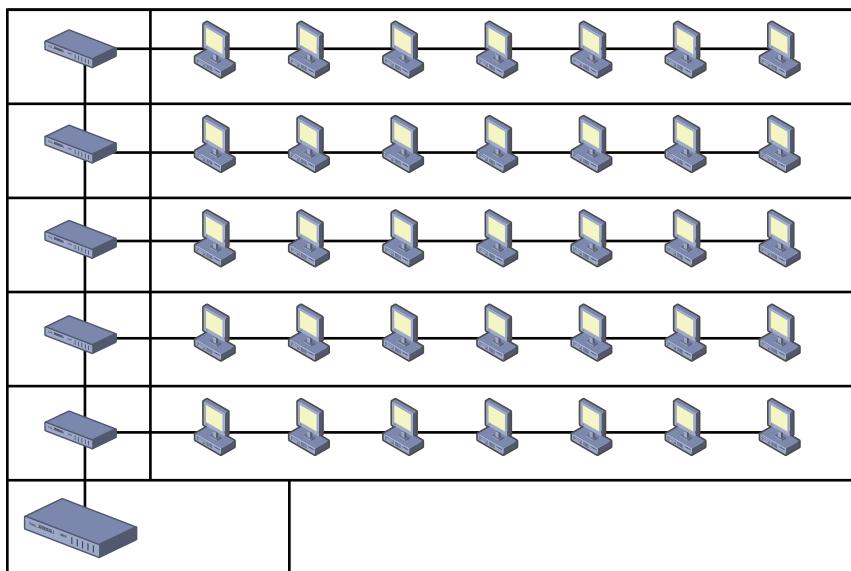


FIGURE 3-22 A backbone network connecting the horizontal cabling on multiple floors.

The ANSI/TIA-568-C standards don't stop at the horizontal cabling and the telecommunications rooms for individual LANs. They also define the networking infrastructure for the entire enterprise, including backbone networks and services originating outside the building, such as telephone and Internet services.

As noted earlier, each telecommunications room contains a switch that actually connects the work area cable runs together. With the cable installation completed,

you use patch cables to connect the ports in the patch panels to the ports in the switch. In the parlance of structured cabling, this switch is the horizontal cross-connect; it joins all of the individual horizontal cable runs into a single LAN.

At this point in the example, what you have is a series of individual LANs, one on each floor. The object, though, is to connect all of the LANs together, so that a computer on any LAN can communicate with a computer on any other LAN. You also want all of the computers on all of the LANs to be able to access a single connection to the Internet.

The backbone is another network, the primary function of which is to connect the LANs together. There are no workstations connected directly to the backbone. In some cases, you might connect servers to it, but many backbones carry nothing but intermediate traffic.

Somewhere in the building, often on the ground floor or in the basement, there is an equipment room containing another patch panel. This patch panel, called the vertical cross connect, contains one end of the vertical cable runs leading to each of the other telecommunications rooms in the building. These vertical cable runs can use the same or a different type of cable as the horizontal runs on each floor.

Also in the equipment room is the backbone switch that connects all of the LANs together into one internetwork. As in the telecommunications rooms, you use patch cables to connect the ports in the vertical cross connect patch panel to the ports in the backbone switch.

A large enterprise network will, at minimum, have a demarc for telephone services and one for a connection to an Internet service provider's network. There might also be connections for WAN links to other offices as well. In an ideal situation, these services will enter the building in the same equipment room that houses the backbone switch. This room is then called the main distribution frame (MDF).

Of course, situations rarely turn out to be ideal, and it is not unusual for the MDF equipment to be spread among several different locations. A single office building might house networks for dozens of different companies, all of whom also require telephone service as well. You might find that various services enter the building at different locations and must be routed to various equipment rooms, each of which is a maze of racks and cables. As with the cabling itself, labeling and organization is an essential part of maintaining a complex installation.

True or false: Every large building has a room that functions as its main distribution frame.

Answer: *False*. Sometimes the components that make up the MDF are scattered among a building's various rooms or locations.

EXAM TIP The Network+ objectives include the acronyms IDF and MDF without explanation. Candidates should be familiar with the concepts of the intermediate distribution frame and the main distribution frame for the exam; candidates are often familiar with the technology, but not with those particular terms.

True or false: A backbone network can use a different data-link layer protocol than the horizontal networks it connects.

Answer: *True*. A backbone network can connect the horizontal networks together using a switch, which maintains a single protocol throughout the installation; or it can use routers, enabling you to use any protocol on the backbone.

Identify a demarc in a wiring distribution

One of the uses of a backbone network is to provide the LANs with access to services arriving from outside the building, such as telephone services and WAN connections. The place where an outside service enters the building is called a demarcation point, or demarc. The demarc is where the responsibility of the network administrator ends. If a problem occurs outside the demarc, it is up to the service provider to fix it. Inside the demarc, it is the network administrator's problem.

A demarc typically takes the form of a hardware device or interface furnished by the service provider, sometimes referred to as a network interface unit (NIU) or a network interface device (NID). The interface could be a simple box with an RJ-45 connector in it, or it could be a more elaborate device, such as a smartjack. On a home network that receives access to the Internet through a cable television company, the little box referred to as the cable modem is the NIU.

True or false: Demarcs are typically found in a building's MDF.

Answer: *True*. The main distribution frame (MDF) contains demarcs, WAN interfaces, and other service connections that originate outside the building.

True or false: If a WAN connection fails because of a cable break on the inside of the demarc, it is the responsibility of the WAN provider to fix it.

Answer: *False*. The demarc is the point at which the WAN provider's responsibility ends. Any failures occurring inside of the demarc are the responsibility of the customer.

Identify a demarc extension in a wiring distribution

To a WAN service provider, the demarc is where the provider's responsibility usually ends. Everything inside the demarc is referred to as customer premises equipment (CPE). In an ideal situation, the demarc is located in the same room as the CPE, so making the connection is simply a matter of plugging a cable into the WAN interface and the backbone switch.

In some cases, however, devices are not in the same place, and a demarc extension is required to connect the WAN interface to the CPE. For example, if the demarc is in a basement wiring closet and the customer's network is on the sixth floor, the demarc extension is the infrastructure that brings the service up from the basement to customer premises. The demarc extension might be solely the customer's responsibility, or it could be supplied by the building management.

True or false: A residence can have a demarc extension, just as a business can.

Answer: *True*. A residential CATV customer typically has a demarc somewhere on the outside of the structure. The wiring leading from the demarc to the jack on the wall inside is the demarc extension. If the customer rents the home, then the demarc extension is the responsibility of the landlord.

Identify a smartjack in a wiring distribution

In the case of a large corporate installation, the NIUs for WAN connections are more than simple wiring devices; they can be much more complex. Most NIUs are equipped with smartjacks, which enable them to perform additional functions, such as signal translation, signal regeneration, and remote diagnostics. A site with multiple WAN connections might have a modular cabinet with a smartjack on a separate expansion card for each service connection.

True or false: A smartjack enables the network administrator to monitor the status of a WAN connection.

Answer: *False*. The capabilities of the smartjack are for the benefit of the WAN provider, not the network administrator.

Identify a CSU/DSU in a wiring distribution

A leased line enters the customer's premises and terminates at the CSU/DSU, which functions as the demarcation point, or demarc. A CSU/DSU is actually two devices that are always combined into a single unit that looks something like an external modem. In fact, CSU/DSUs are sometimes called *digital modems*, a term that is incorrect. The CSU part of the device provides the terminus for the digital link and keeps the link alive when no traffic is passing over it. The CSU also provides diagnostic and testing functions. The DSU part of the device translates the signals generated by the LAN equipment into the bipolar digital signals used by the leased line.

True or false: A CSU/DSU can sometimes take the form of an expansion module for a router.

Answer: *True*. Modular CSU/DSU units are available that plug into an expansion slot in a router. This saves space and eliminates the need for a separate power supply.

EXAM TIP The Network+ exam objectives place the CSU/DSU in with the wiring distribution material, but it could just as easily be a part of the WAN coverage, as found in "Objective 3.4: Categorize WAN technology types and properties."

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Where do the horizontal networks and the backbone network usually meet?
2. What is the term for a switch that connects all of the horizontal networks together to form a single network?

3. What type of WAN connection does a CSU/DSU link to your network?
4. In addition to the components of a wiring distribution, what other topics covered in the Network+ exam are defined in the ANSI/TIA-568-C standard?

Answers

This section contains the answers to the “Can you answer these questions?” sections in this chapter.

Objective 3.1: Categorize standard media types and associated properties

1. CAT6 cables can support 1000Base-T Gigabit Ethernet traffic and, with special installation considerations, 10GBase-T.
2. Fiber optic cable is completely immune to electromagnetic interference and is the best choice for the network.
3. Singlemode fiber optic is better suited for extremely long cable runs, up to several kilometers.
4. By connecting two computers together with a crossover cable, you would create a simple, two-node network.

Objective 3.2: Categorize standard connector types based on network media

1. To attach RJ-45 connectors to patch cables, you need a special tool called a crimper that looks like a large pliers. Installers use the crimper to squeeze the connector closed onto the end of the cable.
2. Three BNC connectors are required. One to connect the network adapter in the computer to the T fitting, and two to connect the cables to the cross members of the T.
3. 8P8C means that the connector has eight positions and eight wire connections, one for each position. An RJ-11 connector is also called a 6P2C, because it has six positions, only two of which are wired.
4. Installers use special types of glue to attach connectors to fiber optic cables. The objective is to connect the cable ends face to face, to permit light to travel through them unobstructed.
5. Wall plates have female RJ-45 connectors, to accept the patch cables that connect them to the computers.

Objective 3.3: Compare and contrast different wireless standards

1. 802.11n designations describing the MIMO capabilities of each device use the format $a:b:c$, where a is the number of transmit antennae in the device, b is the number of receive antennae, and c is the number of data streams that the radio in the device supports. The maximum configuration defined

by the standard is 4×4:4, indicating that a device has four transmit and four receive antennae and can send or receive on four channels at once.

- 2. IEEE 802.11a and 802.11n are the two standards that allow for the use of the 5 GHz band.
- 3. Frame aggregation improves network efficiency by reducing the control overhead incurred by the transmission of several packets into that of one. By combining the payloads from multiple packets into one packet, the system eliminates extra data-link layer frames, acknowledgment messages, spaces between frames, and radio communication transmissions.
- 4. IEEE 802.11g supports the 54 Mbps transmission speed, and is backward compatible with 802.11b equipment.

Objective 3.4: Categorize WAN technology types and properties

- 1. The technical name is the Public Switched Telephone Network (PSTN).
- 2. 2B+D is also known as the Basic Rate Interface (BRI).
- 3. The device at both ends of a T-1 connection is a channel service unit/data service unit (CSU/DSU).
- 4. SONET uses fiber optic cable at the physical layer.
- 5. Asymmetric Digital Subscriber Line (ADSL) is the type most often used for residential users.

Objective 3.5: Describe different network topologies

- 1. All coaxial-based Ethernet networks, including Thin Ethernet, use a bus topology. All UTP-based Fast Ethernet networks use a star topology. Therefore, an upgrade from coaxial to UTP cable must include a reconfiguration from a bus to a star topology.
- 2. With the bus topology. A cable break in a bus network splits the network into halves, preventing the nodes on one side from communicating with those on the other. In addition, both halves of the network are left with one unterminated end, which prevents computers on the same side of the break from communicating effectively.
- 3. Virtually all of the new Ethernet networks installed today use the star or the hierarchical star topology, with one or more switches functioning as a cabling nexus.
- 4. A bus topology requires terminating resistors at each end of the bus to remove signals as they reach the end and prevent them from reflecting back in the other direction and interfering with newly transmitted signals

Objective 3.6: Given a scenario, troubleshoot common physical connectivity problems

1. Attenuation
2. Power Sum NEXT and Power Sum EL-FEXT.
3. A wiremap tester consists of a main unit that connects to all eight wires of a UTP cable at once and a loopback device that you connect to the other end, enabling you to test all of the wires at once. A wiremap tester cannot detect split pairs because in that fault the pins are symmetrically connected.
4. A wire that is not connected to both ends of the cable, either because the connection itself is flawed or because the wire is broken somewhere in the cable, generates a condition called an open circuit.
5. Alien crosstalk is the measurement of the signal bleedover from a wire pair in one cable to the same pair in an adjacent cable.

Objective 3.7: Compare and contrast different LAN technologies

1. The two systems wait for different backoff periods before retransmitting. They calculate their backoff periods using an algorithm called truncated binary exponential backoff.
2. You must insert a router between the two networks, to prevent broadcasts from being forwarded.
3. 100 meters.
4. 1000Base-T and 10Gbase-T.
5. A 10 Gbase-T network using CAT6 cable is limited to 55 meter segment lengths.

Objective 3.8: Identify components of wiring distribution

1. The horizontal networks and the backbone network usually meet in an intermediate distribution frame.
2. The switch that connects all of the horizontal networks together is a vertical cross connect or a backbone switch.
3. A CSU/DSU provides your network with an interface to a leased line, such as a T-1.
4. The ANSI/TIA-568-C also defines the characteristics of the cables used to construct a LAN.

CHAPTER 4

Network Management

The Network Management domain accounts for approximately 20% of the CompTIA Network+ exam, and contains coverage of many of the hardware and software tools that network administrators use to manage, monitor, and troubleshoot their networks.

This chapter covers the following objectives:

- Objective 4.1: Explain the purposes and features of various network appliances
- Objective 4.2: Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues
- Objective 4.3: Given a scenario, use appropriate software tools to troubleshoot connectivity issues
- Objective 4.4: Given a scenario, use the appropriate network monitoring resource to analyze traffic
- Objective 4.5: Describe the purpose of configuration management documentation
- Objective 4.6: Explain different methods and rationales for network performance optimization

Objective 4.1: Explain the purpose and features of various network appliances

For this exam objective, you must be familiar with some of the components that the Network+ objectives refer to as network appliances. These are hardware or software mechanisms that perform high-level administrative services for a network.

Exam need to know

- Explain the purpose and features of a load balancer
For example: How do application servers balance a traffic load amongst themselves?
- Explain the purpose and features of a proxy server
For example: How does a proxy server differ from a NAT router?

- Explain the purpose and features of a content filter
For example: How does a content filter differ from a firewall?
- Explain the purpose and features of a VPN concentrator
For example: How does a VPN concentrator enhance the security of a VPN connection?

Explain the purpose and features of a load balancer

Busy servers, whether on the Internet or not, often receive more traffic than a single computer can comfortably handle. When this happens, the only solution is to add another server. A group of servers, all running the same application to provide a single service, is known as a cluster or a server farm. Load balancing is a technique that distributes incoming traffic equally among the multiple servers in the cluster.

When a client accesses a server, it has only one name, and that name must resolve into one IP address. When you add more servers, they obviously must have their own names and addresses, but to the client, they must appear as one. For example, no one could possibly believe that the Microsoft.com website is running on a single server; the traffic would quickly overwhelm an individual machine. So the work is distributed across many servers, all of which answer to Microsoft.com. This is the trick accomplished by a load-balancing solution.

There are several ways to balance a traffic load among servers. The simplest, and one of the most common, is to use a method called round robin DNS. This method is implemented in the DNS server and relies on the name resolution process performed by each client.

In round robin DNS, the server contains multiple resource records for the same server name, each with a different IP address representing one of the computers running the server application. When a client resolves the server name, the DNS server accesses each of the resource records in turn, so that each address theoretically receives the same number of visitors.

There are problems with this technique, mostly due to the caching of DNS information and of server data all over the Internet. Therefore, the traffic received by each server is not exactly balanced with the others, but the solution works adequately for most situations.

There are many more complicated alternatives to the DNS load-balancing approach. Generally speaking, the load balancing is accomplished by a software or hardware component that receives the incoming traffic and selects one of the servers to receive it. Hardware solutions include special switches, actually devices combining switch and router functions, which continually query the servers in the cluster and choose a server to receive each request based on the responses. These go by various names, including multilayer switches, layer 4 switches, and load-balancing switches.

There are also application layer devices called content switches, layer 7 switches, or web switches, among other names. These devices not only perform load balancing, they also can offload some of the processor load from the web servers by performing Secure Sockets Layer (SSL) transactions.

Software solutions can take the form of standalone programs that monitor incoming traffic and forward requests to the servers in the cluster as needed. Other load-balancing schemes are integrated into comprehensive clustering solutions, such as Failover Clustering in Microsoft Windows Server 2008 R2. In addition to the fault tolerance provided by having multiple servers replicated and synchronized, these solutions balance incoming traffic among the servers in the cluster.

True or false: Load balancing provides both fault tolerance and an increase in performance.

Answer: *True*. Having multiple servers running the same application enables the system to continue functioning if a server should fail. Splitting the client traffic among multiple servers enables each one to function more efficiently.

EXAM TIP The Network+ objectives include “load balancers” in the list of network appliances. But there is no one specific device that can be called a load balancer. There are instead several solutions that make load balancing possible.

True or false: In a round robin DNS load-balancing arrangement, the application servers bearing the load are not aware of each other.

Answer: *True*. A DNS round robin load-balancing solution works by distributing traffic among the IP addresses of the application servers. Each server is only aware of its own incoming traffic, and does not know that there are other servers performing the same function.

MORE INFO For more information on load balancing, see “Objective 4.6: Explain different methods and rationales for network performance optimization.”

Explain the purpose and features of a proxy server

A proxy server is an application layer service that functions as an intermediary for network clients accessing the Internet. Proxy servers protect the clients by preventing direct connections between them and Internet servers, and they also enable administrators to monitor and regulate users’ Internet access.

Like a network address translation (NAT) router, a proxy server receives requests from clients on a private network and forwards those requests to the destination on the Internet, by using its own registered address to identify itself. The primary difference between a proxy server and a NAT router is that the proxy server interposes additional functions into the forwarding process. These functions can include the following:

- **Filtering** Administrators can configure proxy servers to limit user access to the Internet by filtering out requests sent to undesirable sites.
- **Logging** A proxy server can maintain logs of user Internet activity for later evaluation and reporting.

- **Caching** A proxy server can store frequently accessed Internet data in a local cache, which it can then use to satisfy subsequent requests for the same data at higher speeds.
- **Scanning** A proxy server can scan incoming data from the Internet for various types of malware and outgoing data for confidential company information.

Unlike a NAT router, which is invisible to the workstation, a proxy server requires that applications be configured to use it—a process that can be manual or automatic.

True or false: Proxy servers are hardware devices that connect to router or switch ports.

Answer: *False*. Proxy servers are software products that you install and run on a computer with a registered IP address.

EXAM TIP Network+ exam candidates should be able to distinguish between a proxy server and a NAT router, both of which are designed to protect a private network from Internet intruders. For more information on NAT, see "Objective 2.1 Given a scenario, install and configure routers and switches."

True or false: A proxy server must be designed for a specific version of a specific application, which is why they are usually made by the same company as the applications they service.

Answer: *False*. Proxy servers are designed to serve specific applications, based on the port numbers they use for their communications. Therefore, they are not designed for specific application versions, nor are they typically made by the application manufacturers.

Explain the purpose and features of a content filter

A content filter is any mechanism that enables administrators to regulate what data is available to users. The typical application for a content filter is to prevent Internet users from accessing objectionable websites. Administrators can usually specify the users or groups of users whose access is to be restricted, as well as what content is permitted or denied, based on a variety of criteria such as domain names, content types, or media types.

Content filters can take the form of client-side software products or be integrated into search engines, but in the context of the Network+ exam, they are typically implemented as part of a proxy server package.

True or false: A content filter performs the same function as a firewall.

Answer: *False*. Firewalls are designed to prevent intruders outside of a network from accessing the resources inside. A content filter is designed to prevent users inside of a network from accessing content on the outside.

EXAM TIP The Network+ exam objectives refer to a content filter as a network appliance, but they are more often implemented as part of a larger package, such as a proxy server, rather than a standalone product.

True or false: Internet Explorer has a built-in content filter.

Answer: *True*. It has only limited capabilities, but the Content Advisor feature in Internet Explorer enables users to block specific types of content from display in the browser.

Explain the purpose and features of a VPN concentrator

A virtual private network (VPN) is a remote access method in which a client establishes a connection to a remote server using the Internet as a network medium. To secure the connection, the systems use a method called tunneling, in which the data exchanged by the client and the server are encapsulated and encrypted for security purposes.

A VPN concentrator is a hardware device located between the VPN client and server that performs several tasks that enhance the security of the connection, including the following:

- Function as a tunnel endpoint
- Construct the tunnel
- Authenticate VPN users
- Encrypt data for tunnel transmission
- Monitor tunnel data transfers

True or false: A VPN concentrator is required to establish a client/server VPN connection.

Answer: *False*. A VPN client can establish a connection to a VPN server without a concentrator. A concentrator just provides additional functionality.

EXAM TIP The Network+ objectives also include virtual private networking in the discussion of security in “Objective 5.2: Explain the methods of network access security.”

True or false: A VPN concentrator is typically located on the server side of a remote access connection.

Answer: *True*. A VPN concentrator performs tasks that would otherwise be the responsibility of a remote access server.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. How does a proxy server protect clients on private network?
2. How does DNS provide load-balancing services for an application?
3. At what layer of the OSI reference model do proxy servers operate?

Objective 4.2: Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues

For this exam objective, you must be familiar with the various hardware tools that network administrators use when building, testing, and troubleshooting a network. Most of these tools are used for troubleshooting physical layer problems, such as cable faults.

Exam need to know

- Troubleshoot using a cable tester
For example: What kinds of faults can a cable tester detect?
- Troubleshoot using a cable certifier
For example: What tests can a cable certifier perform?
- Troubleshoot using a crimper
For example: How do you use a crimper to create patch cables?
- Troubleshoot using a butt set
For example: Of what use is a butt set to a LAN administrator?
- Troubleshoot using a toner probe
For example: How do you use a toner probe to perform continuity testing?
- Troubleshoot using a punch down tool
For example: How do you use a punch down tool to install internal cable runs?
- Troubleshoot using a protocol analyzer
For example: What security hazard does a protocol analyzer present?
- Troubleshoot using a loopback plug
For example: Why is a loopback plug necessary for cable continuity testing?
- Troubleshoot using a TDR
For example: What kinds of cable faults can a TDR detect?
- Troubleshoot using an OTDR
For example: How does an OTDR differ from a TDR?
- Troubleshoot using a multimeter
For example: How can you use a multimeter to test cable continuity?
- Troubleshoot using an environmental monitor
For example: What conditions do environmental monitors read?

Troubleshoot using a cable tester

Identifying and testing each cable thoroughly by using a tone generator and locator or multimeter is a lengthy and arduous process; professional installers rarely use them. For a few dollars more, you can purchase a device better suited to the job.

A cable tester, also called a wiremap tester, connects to all eight wires at once on both ends and tests them at the same time. Wiremap testers, like most cable testers, consist of a main handheld unit that you connect to one end of the cable, and a loopback device that you connect to the other end. A wiremap tester can detect

opens and shorts, as well as transposed wires. However, it cannot detect split pairs. This usually requires a more sophisticated testing device.

A split pair is a connection in which two wires are incorrectly mapped in exactly the same way on both ends of the cable. In a properly wired connection, each twisted pair of wires should contain a signal and a ground. In a split pair, you can have two signal wires twisted together as a pair. This can generate excessive amounts of crosstalk, corrupting both of the signals involved.

True or false: A cable tester cannot detect a split pair because while the wires in the cable are connected incorrectly, the signals are sent and received on the appropriate connector pins.

Answer: *True*. In a properly wired cable, each wire carrying a signal is twisted with its corresponding ground wire. If you mistakenly have two signal-carrying wires twisted together, the signals can interfere with each other, even though the signals are ending up in the correct places.

EXAM TIP The Network+ exam objectives use the term cable tester, which can of course refer to any type of cable testing tool, but the wiremap tester described here is the most likely device to which this objective refers.

True or false: Unlike a tone generator, a wiremap tester is capable of testing all eight of a cable's connections at once.

Answer: *True*. The wiremap tester generates signals on each wire and a loopback device at the other end of the cable sends them back again.

Troubleshoot using a cable certifier

The ANSI/TIA-568-C standard rates unshielded twisted pair (UTP) cables by specifying maximum acceptable levels of various types of interference. As you move up to the higher categories, there are more tests to conduct, and more stringent limitations to meet. Testing these characteristics in the field therefore requires sophisticated (and more expensive) equipment.

Cable testing devices at this level of performance are typically called cable certifiers, scanners, or media testers. These are handheld devices of great sophistication, which fortunately make the testing process utterly foolproof. After connecting the main unit to one cable end, as shown in Figure 4-1, and the remote unit to the other end, you select a battery of tests and the certifier runs through them all in a matter of seconds.

Most cable certifiers are already configured with the appropriate tests and performance levels specified by the ANSI/TIA-568-C standards for each UTP cable category. Prices of cable certifiers can be very high, depending on the other features they provide. Higher-end units have faster testing speeds, can interface with a computer to maintain records of your tests, and might have built-in printers or support for add-on modules that enable you to test other cable types.



FIGURE 4-1 A cable certifier.

True or false: A cable certifier eliminates the need for cable testers and tone generators.

Answer: *True*. A cable certifier can detect all of the faults that these other devices can detect, plus a great deal more.

EXAM TIP The Network+ exam does not cover the features or operation of any particular make or model of cable certifier. The capabilities of the products offered by different manufacturers vary widely, as do their prices. The exam only covers their basic capabilities, such as testing the levels of various types of crosstalk as required by the ANSI/TIA-568-C standards.

True or false: Cable certifiers must be reconfigured whenever a new cable specification is standardized.

Answer: *True*. Most cable certifiers enable you to create your own battery of tests, which allows the device to retain its usefulness even when cable standards change.

Troubleshoot using a crimper

Attaching male RJ-45 connectors to make a patch cable requires a tool similar in appearance to pliers, called a crimper. After stripping off the cable sheath you line the conductors up in the proper order, insert them into the connector, and use the crimper, shown in Figure 4-2, to squeeze the connector closed around them and to pierce the insulation with the eight pins of the connector, thereby making the electrical connections.

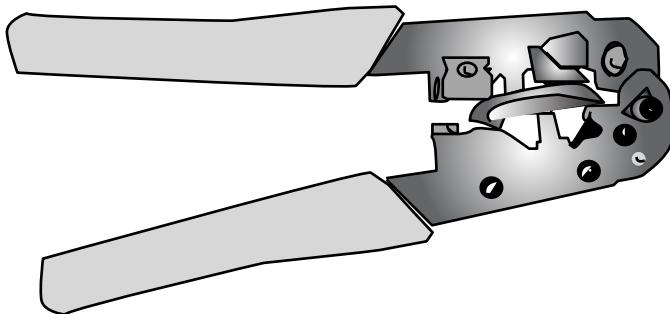


FIGURE 4-2 A crimping tool for attaching RJ-45 connectors.

In many cases, faulty patch cables are the result of improperly applied connectors. Although many administrators find it easier to simply replace a faulty patch cable, you can also troubleshoot it by cutting off the existing connectors and attaching new ones.

True or false: You need to purchase a separate crimping tool for each type of cable to which you want to attach connectors.

Answer: *False*. While some crimpers are designed for a single cable/connector combination, there are many that have multiple or replaceable bits, enabling them to work on a variety of cables and connectors.

EXAM TIP While this Network+ exam objective associates devices such as crimpers and punch down tools as troubleshooting tools, they are more commonly associated with the cable installation process. The average network administrator might never have need for them.

True or false: Making your own patch cables by applying connectors is always more economical than buying prefabricated patch cables.

Answer: *False*. Buying bulk cable and connectors and making patch cables yourself can conceivably be cheaper than purchasing prefabricated cables. However, when you factor in the time needed to attach the connectors, the learning curve required to attach the connectors correctly, and the failure rate requiring the re-application of connectors, it is generally more economical to purchase prefabricated patch cables in quantity instead.

Troubleshoot using a butt set

A lineman's handset, commonly known as a butt set, is a diagnostic and troubleshooting tool for analog networks, and specifically for telephony. The device looks like an ordinary telephone handset, complete with a dial keypad and several types of cable connectors. A butt set typically has a standard RJ-11 connector, and possibly an RJ-45, enabling the user to plug into any modular phone jack. However, the real versatility of the device comes from a set of alligator clips that enable the device to connect to any punch down block or cable. The teeth of the clips can

penetrate the insulation on the cable, so that the device can listen to (or butt in on) the connection, hence the name butt set.

True or false: A butt set has no application on a digital network.

Answer: *True*. A butt set is only capable of working with analog signals. As a result, it has little value to a LAN troubleshooter, except as a means of distinguishing telephone cables from network cables.

True or false: Using a butt set, you can listen in on Voice over IP calls.

Answer: *False*. VoIP is a digital service, and butt sets can only listen in on analog calls.

Troubleshoot using a toner probe

One of the simplest and least expensive cable testing tools is called a tone generator and locator, as shown in Figure 4-3. The same device can also be called a toner probe. The tone generator is a device that you connect to one end of a cable run and which generates a signal over the cable. At the other end of the run, when you touch the locator to the right wire, it produces an audible sound as it detects the signal. You can use this type of device to locate and identify the one cable you need in a bundle, or the one wire you need in a cable.

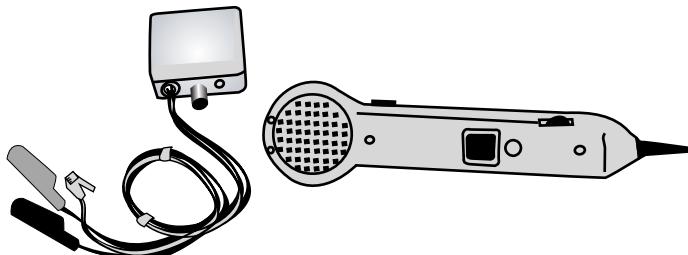


FIGURE 4-3 A tone generator and locator.

The first and most essential test that installers must perform on every cable run is a continuity test, which ensures that each wire on both ends of the cable is connected to the correct pin—and only the correct pin. A wire that is not connected to both ends of the cable, either because the connection itself is flawed or because the wire is broken somewhere in the cable, generates a condition called an open circuit. If a pin on one end of a cable run is connected to two or more pins on the other end, you have a short circuit.

True or false: Continuity testing using a multimeter is a slow and arduous process.

Answer: *True*. To perform a continuity test using a multimeter, you must individually test each wire in each cable, which is a time-consuming procedure.

EXAM TIP The Network+ exam objectives use the term toner probe to refer to the

tool that is also known as a tone generator and locator or a fox and hound tester.

When preparing for the exam, you might encounter references using various names for the same tool.

True or false: When you apply the tone generator to a wire at one end of the cable, but you fail to detect a tone on the same wire at the other end, this indicates the existence of an open fault in the cable.

Answer: *True*. The failure to detect a tone indicates either that there is a break in the wire somewhere inside the cable or a bad contact with the connector at one or both ends.

Troubleshoot using a punch down tool

The process of attaching a jack to a cable end is called punching down, and the process is essentially the same whether you are connecting the cable to a wall plate or a patch panel. In each case, you must strip off the outer sheath to expose the individual wires inside the cable and punch down each wire into the correct pin on the jack.

Punching down the wires to the correct pins is an essential part of the process. UTP cable for LANs is wired straight through, meaning that each pin at one end of the cable must be connected to the corresponding pin at the other end. A connection between two devices on a network must contain a signal crossover somewhere. The signal transmitted over the send pin at one end must somehow end up at the receive pin on the other end, but data networks rely on network devices such as switches, routers, and hubs to supply the crossover circuit.

In a typical UTP installation, there are two types of female RJ-45 connectors that installers have to punch down. The keystone connector is the snap-in module that fits into a standard keystone wall plate. Keystone modules are available in a variety of configurations, enabling you to mix different types of cables in a single plate. Most of the connectors on the market today are color coded. Some are available in T568A and T678B versions; most have color codes for both pinouts. The other connector is a 110 block, which you find on the back of a patch panel.

MORE INFO For more information on 110 blocks, see “Objective 3.2: Categorize standard connector types based on network media.”

The tools required for punching down are minimal. You need a pair of scissors (or specialized cable snips) for cutting off the cable sheath and for cutting the ends of the wires, and you need a specialized punch down tool, like that shown in Figure 4-4. Punch down tools are available with a variety of bits to support different cable types. The elaborateness and expense of the tool you need depends on the types of cables you will be installing.

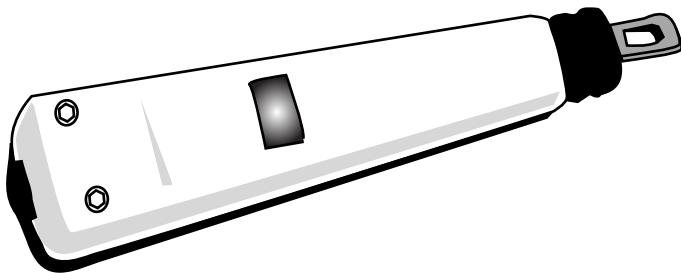


FIGURE 4-4 A punch down tool.

The process of punching down a keystone connector or a patch panel connection consists of the following steps:

1. Strip a few inches of the insulating sheath off the cable end to expose the wires, being careful not to cut into the wires themselves.
2. Separate each of the four twisted wire pairs at the end.
3. Place each of the eight wires into the appropriately colored contact in the connector.
4. Holding your punch down tool vertically, punch each of the eight wires down into its receptacle.

The punch down tool is critical to this operation. In one motion, the tool presses the wire down into place, allowing the connector to pierce its insulation—making both a mechanical and an electrical connection—and cuts off the excess at the end. Without this tool, the process of punching down would require you to strip the insulation off of each individual wire, press it into place, and trim off the end, which would be laborious and time-consuming.

True or false: You cannot use a punch down tool to create patch cables.

Answer: *True*. To create patch cables, you must line the wires up against the connector and squeeze the connector closed with a crimper

True or false: When punching down, you must use the T568A pinout at one end and the T678B pinout at the other end.

Answer: *False*. You can use either the T568A or the T678B pinout, but you must use the same pinout at both ends of the cable.

Troubleshoot using a protocol analyzer

A protocol analyzer, sometimes called a packet sniffer, is one of the most powerful tools for learning about, understanding, and monitoring network communications. A protocol analyzer captures a sample of the traffic passing over the network, decodes the packets into the language of the individual protocols they contain, and lets administrators examine them in minute detail. Some protocol analyzers can also compile network traffic statistics, such as the number of packets that are using each protocol and the number of collisions that are occurring on the network.

Using a protocol analyzer to capture and display network traffic is relatively easy, but interpreting the information that the analyzer presents and using it to troubleshoot the network requires a detailed understanding of the protocols running on the network. However, there is no better way to acquire this type of knowledge than to examine the actual data transmitted over a live network.

A protocol analyzer is typically a software product that runs on a computer connected to a network. On an Ethernet network that uses hubs, protocol analyzers work by switching the network interface adapter that they use to access the network into promiscuous mode. When a network interface adapter is in promiscuous mode, it reads and processes all the traffic that is transmitted over the network, not just the packets that are addressed to it. This means that the system can examine all of the traffic transmitted on the network from one computer.

On today's networks, however, switches are more common than hubs, and as a result, capturing traffic for the entire network is more difficult. Because switches forward incoming unicast traffic only to its intended recipient, a protocol analyzer connected to a standard switch port only has access to one computer's incoming and outgoing traffic, plus any broadcasts transmitted over the local network segment.

To capture all of the traffic transmitted on the network, you must plug the computer running the protocol analyzer into a switch that supports port mirroring. Switches that support port mirroring have a special port to which they send all incoming traffic.

After you have captured a network traffic sample, a list of packets appears in the Frame Summary pane, as shown in Figure 4-5.

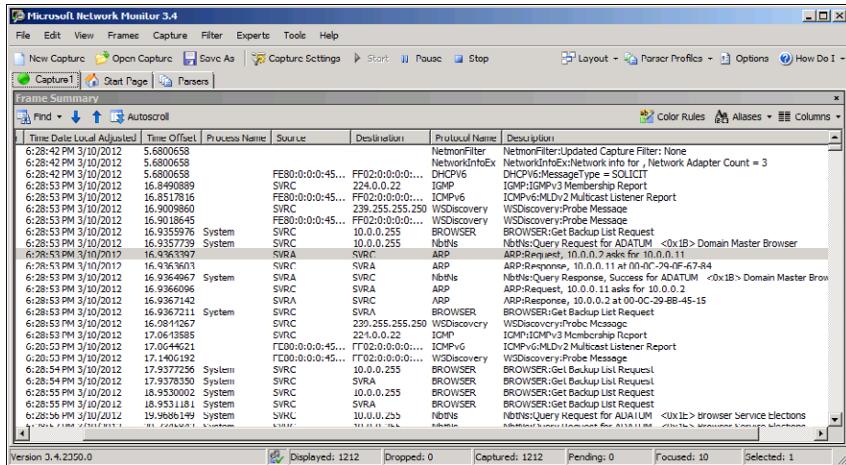


FIGURE 4-5 The Network Monitor Frame Summary pane.

This pane displays a chronological list of the packets in your sample, including the following information:

- **Frame Number** Shows the number of the frame (or packet) in the sample.
- **Time Date** Indicates the time and date that the packet was captured.
- **Time Offset** Indicates the time (in seconds) that the packet was captured, measured from the beginning of the sample.
- **Process Name** Identifies the process that generated the packet.
- **Source** Specifies the name or IP address of the network interface in the computer that transmitted the packet.
- **Destination** Specifies the name or IP address of the network interface in the computer that received the packet.
- **Protocol Name** Identifies the dominant protocol in the packet. Each packet contains information generated by protocols running at several different layers of the Open Systems Interconnection (OSI) reference model. The protocol specified here indicates the primary function of the packet.
- **Description** Specifies the function of the packet, using information specific to the protocol referenced in the Protocol field. For an HTTP packet, for example, this field indicates whether the packet contains an HTTP GET Request or a Response message.

From this main display, you can track the progress of transactions between specific pairs of computers on your network. For example, you can see that an exchange of ARP messages between two systems consists of a request and a response.

When you select one of the packets listed in the Frame Summary pane, the frame details and Hex Details panes display the contents of the selected packet, as shown in Figure 4-6.

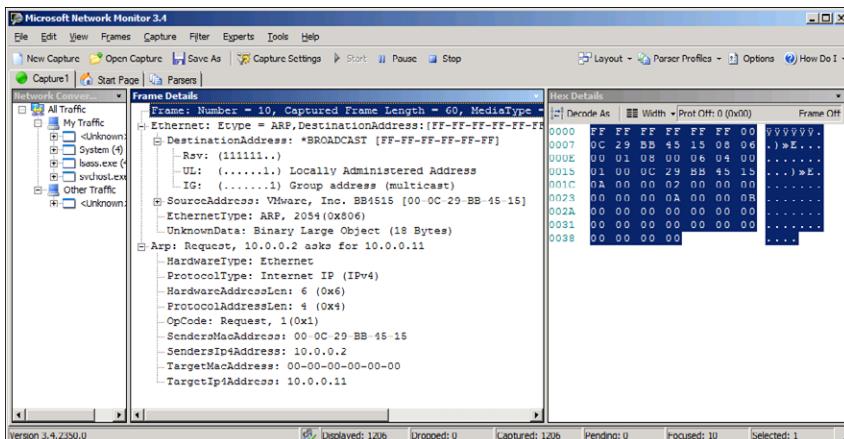


FIGURE 4-6 The Network Monitor Frame Details and Hex Details panes.

The Frame Details pane contains the contents of the selected packet in a fully interpreted, expandable display. The Hex Details pane contains the raw, uninterpreted contents of the packet in hexadecimal and alphanumeric form.

The Frame Details pane is where you can learn the most about the contents of each packet. The analyzer interprets the data in the packet and separates it into the headers for the protocols operating at the OSI model layers. Clicking the plus sign next to a protocol expands it to display the contents of the various header fields.

The Hex Details pane is used primarily to view the application layer data carried as the payload inside a packet. For example, when you look at an HTTP Response packet transmitted by a web server to a browser, you see the HTML code of the webpage the server is sending to the browser.

True or false: To troubleshoot using a protocol analyzer, you must be familiar with the layers of the OSI reference model and the protocols that operate at each layer.

Answer: *True*. A protocol analyzer captures frames and displays their contents, including the individual header fields created by the protocols at the data-link, network, and transport layers. However, to interpret the exchanges between the computers on the network, you must be familiar with the protocols and how they operate.

EXAM TIP The Network+ exam objectives classify protocol analyzers as hardware tools, and there are some dedicated hardware products that are essentially special purpose, portable computers running a protocol analyzer program. However, the objectives also include protocol analyzers as software tools, and there are many such products that you can run on any computer. The most commonly used protocol analyzer is the Microsoft Network Monitor application, mostly because it is available as a free download from the Microsoft website.

True or false: Protocol analyzers can be a network security risk.

Answer: *True*. Protocol analyzers are useful tools in the hands of experienced network administrators, but they can also be used for malicious purposes. In addition to displaying the information in the captured packets' protocol headers, the analyzer can also display the data carried inside the packets. This can sometimes include confidential information, such as unencrypted passwords and personal correspondence.

Troubleshoot using a loopback plug

A loopback plug is a device that connects to a networking device or a cable and reflects all outgoing signals back in to the source. Many of the devices used to test and troubleshoot cables and other network hardware require a loopback plug to function.

For example, cable or wiremap testers and cable certifiers consist of two components: one that generates signals, which you connect to one end of a cable; and a loopback plug, which sends the signals back over the appropriate wires.

True or false: To test the actual connection hardware of a network interface adapter, you must use a loopback plug.

Answer: *True*. Many network adapters have integrated diagnostic testing capability, and many claim to perform a loopback test, but to perform an actual hardware test on the adapter's RJ-45 socket you must connect a loopback plug to it.

True or false: The loopback plug for a UTP cable tester takes the arriving signals and sends them back over the same wires.

Answer: *False*. UTP cables have wires dedicated to carrying signals in each direction, so a loopback plug must take the signals arriving over the receive wires and send them back over the transmit wires.

Troubleshoot using a TDR

UTP cables have length limitations because signals tend to lose strength (or attenuate) over distance. Installers should be conscious of their cable run lengths as they install them, but testing the actual length of the installed cable runs is always recommended.

As you move upward in the cable tester market, you find units that have the ability to determine the length of a cable run, locate a break in any one of the wires, and specify the location of the break, in terms of the distance from the cable end. The technique that provides this capability is called time domain reflectometry (TDR). The tester transmits a signal over the cable and measures how long it takes for a reflection of the signal to return from the other end. Using this information and the cable's nominal velocity of propagation (NVP), a specification supplied by the cable manufacturer, the device can calculate the length of a cable run. If the result is not the same for all eight wires, then the device knows that there is a break somewhere.

True or false: TDR testing requires a loopback plug.

Answer: *True*. To reflect the signal back from the far end of the cable to the source requires a hardware loopback plug at the other end.

EXAM TIP The Network+ exam objectives refer to TDR only by its initials.

Candidates for the exam must know that TDR stands for time domain reflectometry and be familiar with the function of the device and the cable faults it can diagnose.

True or false: Cable certifiers typically include TDR capabilities.

Answer: *True*. TDR functionality is often incorporated into cable certifier devices that can perform a full battery of tests on a cable run at one time.

Troubleshoot using an OTDR

Cable testing for fiber optic cables is, not surprisingly, a completely different world from copper cable testing, and a far more expensive one as well. Fiber optic cables are subject to different types of interference and have largely different performance level requirements. The ANSI/TIA-568-C standards include specifications for fiber optic as well as copper cables, however, and there are cable certifiers designed for testing fiber optic too.

Fiber optic cables do not have multiple wire pairs, so wire mapping is not a concern, and neither is any type of crosstalk. However, cable length and cable breaks are as much of a concern in fiber optic as in copper cable installations. The fiber optic equivalent of a TDR is called an optical time domain reflectometer (OTDR), and it can determine cable lengths and locate cable breaks, just like the copper testing device.

Some of the other primary characteristics that fiber optic cable certifiers usually test include the following:

- **Attenuation** Weakening of the light signal as it travels over long distances
- **Dispersion** Spreading of the light signal as it travels over long distances
- **Light leakage** Signal damage caused by exceeding the cable's bend radius during installation
- **Modal distortion** Signal damage unique to multimode fiber optic cable caused by the various modes having different propagation velocities

True or false: Locating a cable break precisely on a fiber optic network can be more critical than on a copper cable network.

Answer: *True*. Fiber optic cable runs can be several kilometers long, and can be buried in conduits, under pavement, or in other places that are not readily accessible. Precisely locating a break in a fiber optic cable can be a critical part of the troubleshooting process.

True or false: Fiber optic cable testing equipment is far more expensive than copper cable testing tools.

Answer: *True*. Prices of OTDRs can reach well into the five figures in USD, greatly exceeding the average prices of copper cable certifiers.

Troubleshoot using a multimeter

A multimeter, sometimes called a Volt-Ohm Meter (VOM), is an electronic device that can measure voltage, current, and resistance. As a troubleshooting device, a multimeter is invaluable for testing power supplies and batteries. It is also possible to test cable continuity by using a standard multimeter.

True or false: Testing cable continuity with a multimeter requires precise accuracy of measurement.

Answer: *False*. To test cable continuity using a multimeter, you touch the meter's probes to opposite ends of a cable and set the readout to Ohms. You should see a result close to 0 Ohms if you have a proper connection, and infinite Ohms if you do not. You do not need a precise measurement; the device only has to read the presence or absence of resistance.

True or false: A multimeter can be used to troubleshoot fiber optic cabling.

Answer: *False*. Multimeters are devices that measure electricity. Fiber optic cables do not use electric signals, so there is nothing for a multimeter to measure.

Troubleshoot using an environmental monitor

The MDF and IDFs for a large network, like all of the telecommunications rooms in a building, contain a lot of critically important, and extremely expensive, equipment. To keep the equipment functional and safe, you must see to it that these rooms remain cool, dry, and secure. You must also ensure that all of the equipment has access to a high-quality source of power. Adequate locks, alarms, and HVAC equipment are essential. You might also want to install monitoring equipment to ensure that the conditions are maintained, even when the room is vacant. A voltage event recorder can monitor the power for interruptions, and a temperature monitor can ensure that climate conditions remain constant.

True or false: Some computers have built-in temperature monitoring equipment.

Answer: *True*. Computers can have temperature sensors that cause system fans to speed up or slow down as needed.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is the term for a device that determines the length of a cable by transmitting a signal at one end and measuring how long it takes for a reflection of the signal to return from the other end?
2. A continuity test of a newly installed UTP cable run with a wiremap tester indicates an open circuit on one of the wires. What would be the most likely cause for the problem?
3. You are working on a CAT5 UTP network that is several years old. Over time, some of the paper labels the original cable installers used to identify wall plates and patch panel connectors have worn away or fallen off. As a result, there are some cable runs that you are unable to identify. After checking with your supervisor, you discover that the company has no cable testing equipment and is unwilling to hire a consultant just to identify cable runs. What is the most inexpensive tool you can buy to associate unlabeled wall plates with the correct patch panel ports?
4. Which of the tools covered in this objective is only usable on analog connections?
5. Which tool might you use to connect internal UTP cable runs to the keystone connectors that snap into wall plates?

Objective 4.3: Given a scenario, use appropriate software tools to troubleshoot connectivity issues

Virtually every operating system with networking capabilities includes support for the TCP/IP protocols. In most cases, the TCP/IP stack includes utilities you can use to gather information about the various protocols and the network. Traditionally, these utilities run from the command line, although there are some graphical versions. In many cases, TCP/IP utilities use the same syntax, even on different operating

systems. This objective requires you to be familiar with some of the most common TCP/IP utilities and other software troubleshooting tools.

Exam need to know

- Troubleshoot using a protocol analyzer
For example: How can a protocol analyzer be used for troubleshooting?
- Troubleshoot using a throughput tester
For example: How can a system calculate the throughput of a network?
- Troubleshoot using connectivity software
For example: How is connectivity software useful for troubleshooting?
- Troubleshoot using Ping
For example: Why has Ping become a less reliable tool in recent years?
- Troubleshoot using Traceroute
For example: How can Traceroute help to locate a faulty router?
- Troubleshoot using dig
For example: How is dig an improvement over nslookup?
- Troubleshoot using Ipconfig/ifconfig
For example: What can ifconfig do that Ipconfig.exe cannot?
- Troubleshoot using Nslookup.exe
For example: What DNS information can Nslookup.exe display?
- Troubleshoot using ARP
For example: How can you use ARP to speed up a system's performance?
- Troubleshoot using nbtstat
For example: What information can Nbtstat.exe display?
- Troubleshoot using netstat
For example: How can you use netstat to identify unauthorized programs running on a computer?
- Troubleshoot using Route.exe
For example: How do you create a routing table entry with Route.exe?

Troubleshoot using a protocol analyzer

A protocol analyzer is a device or application that captures samples of network traffic and displays the information inside the frames. In addition to the application layer data carried inside the frame, the analyzer displays the contents of the protocol headers at each layer of the OSI model.

MORE INFO The Network+ exam objectives list protocol analyzers under hardware, as well as software, tools. For more information on protocol analyzers, see Objective 4.2 Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues.

Troubleshoot using a throughput tester

A throughput tester is a software program that measures the actual speed at which data is transferred over a network connection. By transmitting a specific amount of data to a destination and measuring the exact time it takes for the data to travel there and back, the software computes the throughput—measured in kilobits per second (Kbps), megabits per second (Mbps), or gigabits per second (Gbps).

True or false: Throughput and bandwidth are the same things.

Answer: *False*. Bandwidth is a measurement of a network's potential speed, while throughput is a measurement of the network's actual speed. Installing Gigabit Ethernet equipment provides 1 Gbps of bandwidth, but there are many factors that can affect the actual speed of the network. A throughput tester measures the actual speed realized by the hardware, with all mitigating factors considered.

EXAM TIP Candidates for the Network+ exam should be aware of the distinction between throughput and bandwidth.

Troubleshoot using connectivity software

Many network administrators are responsible for troubleshooting computers at distant locations. To save travel time and expense, there are software products that enable an administrator to access remote computers and interact with them as though they were sitting at the console. For example, all Windows computers include a feature called Remote Desktop, which is a client/server application that provides full access to any computer accessible through a local network or the Internet. For troubleshooting purposes, administrators can examine the event logs on the remote computer, configure virtually any setting, and install software, as though they were sitting at the computer.

True or false: Remote Desktop access requires the permission of the user.

Answer: *True*. The remote computer must be configured to accept remote access requests before a client can connect to it. It is also possible to control access by requiring remote users to authenticate before access is granted.

EXAM TIP The Network+ exam objectives do not refer to any connectivity software products by name. Candidates for the exam need only know that such products exist and that troubleshooters use them to access remote systems.

Troubleshoot using Ping

Ping is a utility included with all TCP/IP implementations that can tell you if the TCP/IP stack of another system on the network is functioning normally and if you have connectivity to it. Ping can also verify name resolution by making calls to DNS. The Ping program works by generating a series of Echo Request messages using the Internet Control Message Protocol (ICMP), and transmitting them to the computer

whose host name or IP address you specify on the command line. The basic syntax of the Ping program is as follows:

```
ping target
```

The target variable contains the IP address or host name of a computer on the network. You can use IP addresses, Fully-Qualified Domain Names (FQNDs), or Network Basic Input/Output System (NetBIOS) names in Ping commands. When provided with a name, Ping resolves the name into an IP address before sending the Echo Request messages, and it then displays the address in its readout. Most Ping implementations also have command-line switches with which you can modify the operational parameters of the program, such as the number of Echo Request messages it generates and the amount of data in each message.

TCP/IP computers respond to any Echo Request messages they receive that are addressed to them by generating Echo Reply messages and transmitting them back to the sender. When the pinging computer receives the Echo Reply messages, it produces a display like the following:

```
Pinging cz1 [192.168.2.10] with 32 bytes of data:  
Reply from 192.168.2.10: bytes=32 time<1ms TTL=128  
Ping statistics for 192.168.2.10:  
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),  
Approximate round trip times in milli-seconds:  
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

In the case of this Ping implementation (from Windows 7), the display shows the IP address of the computer receiving the Echo Requests, the number of bytes of data included with each request, the elapsed time between the transmission of each request and the receipt of each reply, and the value of the Time to Live (TTL) field in the IP header. In this particular example, the target computer was on the same LAN, so the time measurement is very short—less than 1 millisecond. When you are pinging a computer on the Internet, the interval is likely to be longer.

A successful Ping test such as this one indicates that the target computer's networking hardware is functioning properly, as are the protocols, at least as high as the network layer of the OSI model. If the Ping test fails, there is a problem in one or both of the computers, in the connection between them, or in the name resolution process.

True or false: A Ping test failure always indicates the existence of a malfunction in one of the computers or the network.

Answer: *False*. The usefulness of Ping has been hampered in recent years by the tendency for firewalls to block ICMP Echo Request messages by default. By using command-line parameters, attackers can easily make Ping bombard a target system with an endless stream of large ICMP packets, creating what is known as a denial of service attack. Therefore, many administrators use firewalls to prevent those

messages from reaching their servers. As a result, a failed Ping test does not always mean that there is a problem.

EXAM TIP Candidates for the Network+ exam should be familiar with the output from the Ping program. Different operating systems might have slight variations in the appearance of the output, but the information should be fundamentally the same.

True or false: When a Ping test using a computer's host name fails, but the test succeeds using the same computer's IP address, this indicates that there is a name resolution problem.

Answer: *True*. Any time a host name connection fails but an IP address connection succeeds, the computer is failing to resolve the name into an IP address. This can be due to an incorrect DNS server address or a malfunction in the DNS server itself.

Troubleshoot using Traceroute

Traceroute is a variant of the Ping program that displays the path that packets take to their destination. Because of the nature of IP routing, paths through an internet-work can change from minute to minute, and Traceroute displays a list of the routers that are currently forwarding packets to a particular destination. The program is called traceroute on UNIX and Linux systems and Tracert.exe on Windows-based systems.

Traceroute uses ICMP Echo Request and Echo Reply messages just like Ping, but it modifies the messages by changing the value of the Time to Live (TTL) field in the IP header. The values in the TTL field prevent packets from getting caught in router loops that keep them circulating endlessly around the network. The computer generating the packet normally sets a relatively high value for the TTL field; on computers running Windows, the default value is 128. Each router that processes the packet reduces the TTL value by one. If the value reaches zero, the last router discards the packet and transmits an ICMP error message back to the original sender.

When you run Traceroute with the name or IP address of a target computer, Traceroute generates its first set of Echo Request messages with TTL values of 1. When the messages arrive at the first router on their path, the router decrements their TTL values to 0, discards the packets, and reports the errors to the sender. The error messages contain the router's address, which Traceroute displays as the first hop in the path to the destination. Traceroute's second set of Echo Request messages use a TTL value of 2, causing the second router on the path to discard the packets and generate error messages. The Echo Request messages in the third set have a TTL value of 3, and so on. Each set of packets travels one hop farther than the previous set before causing a router to return error messages to the source. The list of routers displayed by Traceroute as the path to the destination is the result of these error messages. The following is an example of a Traceroute display:

```
Tracing route to www.fineartschool.co.uk [173.146.1.1] over a maximum of 30
hops:
1  <10 ms    1 ms  <10 ms  192.168.2.99
2  105 ms   92 ms   98 ms  qrv1-67terminal01.cpandl.com [131.107.24.67.3]
```

```

3 101 ms 110 ms 98 ms qrvl.cpndl.com [131.107.67.1]
4 123 ms 109 ms 118 ms svcr03-7b.cpndl.com [131.107.103.125]
5 123 ms 112 ms 114 ms clsm02-2.cpndl.com [131.107.88.26]
6 136 ms 130 ms 133 ms sl-gw19-pen-6-1-0-T3.fabrikam.com [157.54.116.5]
7 143 ms 126 ms 138 ms sl-bb10-pen-4-3.fabrikam.com [157.54.5.117]
8 146 ms 129 ms 133 ms sl-bb20-pen-12-0.fabrikam.com [157.54.5.1]
9 131 ms 128 ms 139 ms sl-bb20-nyc-13-0.fabrikam.com [157.54.18.38]
10 130 ms 134 ms 134 ms sl-gw9-nyc-8-0.fabrikam.com [157.54.7.94]
11 147 ms 149 ms 152 ms sl-demon-1-0.fabrikam.com [157.54.173.10]
12 154 ms 146 ms 145 ms ny2-back-1-ge021.router.fabrikam.com
   [157.54.173.121]
13 230 ms 225 ms 226 ms tele-back-1-ge023.router.adatum.co.uk
   [157.60.173.12]
14 233 ms 220 ms 226 ms tele-core-3-fxp1.router.adatum.co.uk
   [157.60.252.56]
15 223 ms 224 ms 224 ms tele-access-1-14.router.adatum.co.uk
   [157.60.254.245]
16 236 ms 221 ms 226 ms tele-service-2-165.router.adatum.co.uk
   [157.60.36.149]
17 220 ms 224 ms 210 ms www.fineartschool.co.uk [206.73.118.65]

```

Trace complete.

In this example, Traceroute displays the path between a computer in Pennsylvania and one in the United Kingdom. Each of the hops contains the elapsed times between the transmission and reception of three sets of Echo Request and Echo Reply packets.

True or false: Ping can tell you that a problem exists, but Traceroute can tell you approximately where the problem is located.

Answer: *True*. You can use Traceroute to isolate the location of a network communications problem. A failure to contact a remote computer with Ping could be due to a problem in your workstation, in the remote computer, or in any of the routers in between. Traceroute can tell you how far your packets are going before they run into the problem.

EXAM TIP Network+ exam candidates should be aware that many routers block or deprioritize ICMP processes in favor of packet forwarding and other critical router tasks. When a router is busy, it might delay the processing of a Ping or Traceroute packet, and the resulting latency numbers will be higher than the delay experienced by actual data packets crossing the network.

True or false: In the sample trace, you can see the point at which the packets begin traveling across the Atlantic Ocean.

Answer: *True*. At hop 13, the elapsed times increase from approximately 150 to 230 milliseconds (ms) and stay in that range for the subsequent hops. This additional delay of only 80 ms is the time it takes the packets to travel the thousands of miles across the Atlantic Ocean.

Troubleshoot using dig

Dig is a name resolution utility that has largely replaced nslookup in most UNIX and Linux distributions, due in part to its extensive list of options and its more detailed output. The basic syntax for dig is as follows:

```
dig @server name type
```

This yields an output like that shown in Figure 4-7.

```
; <>> DiG 9.7.3 <>> @192.168.2.1 www.microsoft.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 28002
;; flags: qr rd ra; QUERY: 1, ANSWER: 4, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.microsoft.com.           IN      A

;; ANSWER SECTION:
www.microsoft.          1066    IN      CNAME   toggle.www.ms.akadns.net.
toggle.www.ms.akadns.net. 60      IN      CNAME   g.www.ms.akadns.net.
g.www.ms.akadns.net.     72      IN      CNAME   lb1.www.ms.akadns.net.
lb1.www.ms.akadns.net.   293     IN      A       207.46.19.254

;; Query time: 25 msec
;; SERVER: 192.168.2.1#53(192.168.2.1)
;; WHEN: Sat Mar 17 15:31:38 2012
;; MSG SIZE  rcvd: 123
```

FIGURE 4-7 Output from the dig program.

Some of the many command-line options for dig are as follows:

```
dig [@server] [-b address] [-f filename] [-k filename] [-m] [-p port#]
[-t type] [-x address] [-4] [-6] [name] [type]
```

- **@server** Specifies the name or address of the DNS server to be queried.
- **-b address** Specifies the source address to be used in the query.
- **-f filename** Runs the program in batch mode, using the commands in the file identified by the filename variable.
- **-k filename** Signs the queries by using transaction signatures found in a key file identified by the filename variable.
- **-m** Enables memory usage debugging.
- **-p port#** Specifies the port number to which the program should send queries. The default value is 53.
- **-t type** Specifies the type of query to be performed.
- **-x address** Performs a reverse lookup of the IP address specified by the address variable.
- **-4** Performs IPv4 queries only.
- **-6** Performs IPv6 queries only.

- **name** Specifies the DNS name to be resolved.
- **type** Specifies the type of resource records the program should display. The default value is A, for host records.

Troubleshoot using Ipconfig/ifconfig

UNIX and Linux systems have a program called ifconfig (the name is derived from the words *interface configuration*) that you use to configure the properties of network interface adapters and assign TCP/IP configuration parameters to them. Running ifconfig with just the name of an interface displays the current configuration of that interface.

Windows has a version of this program, Ipconfig.exe, which omits most of the configuration capabilities and retains the configuration display. When you run Ipconfig.exe with the /all parameter at the Windows Server 2008 R2 command line, you see a display like the following:

```
Windows IP Configuration
  Host Name . . . . . : wkstn12
  Primary Dns Suffix . . . . . : adatum.local
  Node Type . . . . . : Broadcast
  IP Routing Enabled. . . . . : No
  WINS Proxy Enabled. . . . . : No
  DNS Suffix Search List. . . . . : adatum.local

Ethernet adapter Local Area Connection:
  Connection-specific DNS Suffix . : adatum.local
  Description . . . . . : Realtek PCIe GBE Family Controller
  Physical Address. . . . . : 60-EB-69-93-5E-E5
  DHCP Enabled. . . . . : Yes
  Autoconfiguration Enabled . . . . . : Yes
  Link-local IPv6 Address . . . . . : fe80::7441:4473:f204:ec1d%10
                                         (Preferred)
  IPv4 Address. . . . . : 192.168.2.9(Preferred)
  Subnet Mask . . . . . : 255.255.255.0
  Lease Obtained. . . . . : Wednesday, March 14, 2012 8:05:32 AM
  Lease Expires . . . . . : Thursday, March 22, 2012 8:05:29 AM
  Default Gateway . . . . . : 192.168.2.99
  DHCP Server . . . . . : 192.168.2.1
  DHCPv6 IAID . . . . . : 241232745
  DHCPv6 Client DUID. . . . . : 00-01-00-01-14-81-CC-39-60-EB-69-93-
                               5E-E5
  DNS Servers . . . . . : 192.168.2.1
  Primary WINS Server . . . . . : 192.168.2.1
  NetBIOS over Tcpip. . . . . : Enabled
```

Ipconfig.exe also has another function associated with the Dynamic Host Configuration Protocol (DHCP): it is the easiest way in Windows to see what IP address and other parameters the DHCP server has assigned to a computer.

True or false: You can also use Ipconfig.exe to manually release IP addresses obtained through DHCP and renew existing leases.

Answer: *True*. By running `Ipconfig.exe` with the `/release` or `/renew` command-line parameters, you can release or renew the IP address assignment of one of the network interfaces in the computer or of all of the interfaces at once.

EXAM TIP Candidates for the Network+ exam should be familiar with the output from the `Ipconfig.exe` command. The sample shown here contains only the information for a single network interface, but each interface in the computer has its own section in the display, as do the IPv6/IPv4 compatibility mechanisms, such as Teredo.

True or false: You can use `Ipconfig.exe` to assign a static IP address to a Windows computer from the command line.

Answer: *False*. You cannot configure an IP address on Windows using `Ipconfig.exe`. However, you can configure a UNIX or Linux TCP/IP client using the `ifconfig` utility.

Troubleshoot using `nslookup`

The `nslookup` (in UNIX/Linux) and `Nslookup.exe` (in Windows) command-line utilities enable you to generate DNS request messages and transmit them to specific DNS servers on the network. The advantage of `nslookup` is that you can test the functionality and the quality of the information on a specific DNS server by specifying it on the command line.

The basic command-line syntax of `Nslookup` is as follows:

```
nslookup DNSname DNSserver
```

- ***DNSname*** Specifies the DNS name that you want to resolve.
- ***DNSserver*** Specifies the DNS name or IP address of the DNS server that you want to query for the name specified in the *DNSname* variable.

There are also many additional parameters that you can include on the command line to control the server query process. The output generated by `Nslookup.exe` in Windows XP looks like the following:

```
C:\Users\craigz.ZACKER>nslookup www.microsoft.com 192.168.2.1
Server:  cz1.zacker.local
Address: 192.168.2.1
```

Non-authoritative answer:

```
Name:      lb1.www.ms.akadns.net
Address:   207.46.131.43
Aliases:   www.microsoft.com
           toggle.www.ms.akadns.net
           g.www.ms.akadns.net
```

True or false: The `nslookup` utility has two operational modes: command-line and interactive.

Answer: *True*. When you run `nslookup` with no command-line parameters, the program displays its own prompt from which you can issue commands to specify the default DNS server to query, resolve multiple names, and configure many other aspects of the program's functionality.

EXAM TIP The Network+ exam typically limits its coverage of the nslookup utility to the command line interface. The interactive method has a great many commands and parameters that exam candidates do not need to know.

True or false: When you specify a DNS server name on the nslookup command line, the program displays only resource records from that server.

Answer: *True*. Unlike the normal DNS name resolution process, nslookup by default does not let the specified server issue any queries, whether iterative or recursive, to other servers.

Troubleshoot using ARP

The Address Resolution Protocol (ARP) enables a TCP/IP computer to convert IP addresses to the hardware addresses that data-link layer protocols need to transmit frames across local subnets. IP uses ARP to discover the hardware address to which each of its datagrams will be transmitted. To minimize the amount of network traffic ARP generates, the computer stores the resolved hardware addresses in a cache in system memory. The information remains in the cache for a short time (usually between two and ten minutes) in case the computer has additional packets to send to the same address.

UNIX, Linux, and Windows all include a command-line utility that you can use to manipulate the contents of the ARP cache. In UNIX and Linux, this utility is called arp; in Windows, it is called Arp.exe. You can use arp or Arp.exe to add the hardware addresses of computers you contact frequently to the cache, saving time and reducing network traffic during the connection process.

The arp and Arp.exe utilities use a similar syntax, with many identical command-line arguments. This syntax and some of the most important command-line arguments are as follows:

```
arp [[-a {ipaddress}] [-n ifaddress]] [-s ipaddress hwaddress {ifaddress}]  
[-d ipaddress {ifaddress}]
```

- **-a {ipaddress}** Displays the contents of the ARP cache. The optional ipaddress variable specifies the address of a particular cache entry to be displayed.
- **-n ifaddress** Displays only the contents of the ARP cache associated with the network interface specified by the ifaddress variable.
- **-s ipaddress hwaddress {ifaddress}** Adds a static entry to the ARP cache. The ipaddress variable contains the IP address of the host. The hwaddress variable contains the hardware address of the same host. The ifaddress variable contains the IP address of the network interface in the local system for which you want to modify the cache.
- **-d ipaddress {ifaddress}** Deletes the entry in the ARP cache that is associated with the host represented by the ipaddress variable. The optional ifaddress variable specifies the cache from which the entry should be deleted.

The ARP table of a computer running Windows 7, as displayed by Arp.exe, appears as follows:

Interface: 192.168.2.9 --- 0xa	Internet Address	Physical Address	Type
192.168.2.1	00-0b-cd-cf-e3-b6	dynamic	
192.168.2.3	00-22-64-34-12-4b	dynamic	
192.168.2.4	00-11-2f-46-c1-46	dynamic	
192.168.2.99	00-1c-10-08-f6-1a	dynamic	
192.168.2.255	ff-ff-ff-ff-ff-ff	static	
224.0.0.22	01-00-5e-00-00-16	static	
224.0.0.251	01-00-5e-00-00-fb	static	
224.0.0.252	01-00-5e-00-00-fc	static	
239.255.255.250	01-00-5e-7f-ff-fa	static	
255.255.255.255	ff-ff-ff-ff-ff-ff	static	

True or false: Addresses that you add to the ARP cache manually are static, not dynamic.

Answer: *True*. The system does not delete manually-added addresses from the ARP cache after the usual expiration period. They remain in memory until you manually remove them or reboot the computer.

MORE INFO For more information on ARP, see “Objective 1.6: Explain the function of common networking protocols.”

True or false: It is possible to preload the ARP cache with frequently-used addresses whenever you start the computer.

Answer: *True*. If you want to preload the ARP cache whenever you boot your system, you can create a script containing arp or Arp.exe commands and execute it by using an rc file in UNIX/ Linux or by placing it in the Startup program group in Windows.

Troubleshoot using Nbtstat.exe

Nbtstat.exe is a Windows command-line program that displays information about the NetBIOS Over TCP/IP (NetBT) connections that Windows uses when communicating with other computers running Windows on a TCP/IP network. The syntax for Nbtstat.exe is as follows:

```
nbtstat [-a name] [-A ipaddress] [-c] [-n] [-r] [-R] [-s] [-S] [-RR]
```

- **-a name** Displays the NetBIOS names registered on the computer identified by the name variable.
- **-A ipaddress** Displays the NetBIOS names registered on the computer identified by the ipaddress variable.
- **-c** Displays the contents of the local computer’s NetBIOS name cache.
- **-n** Displays the NetBIOS names registered on the local computer.

- **-r** Displays the number of NetBIOS names registered and resolved by the local computer, using both broadcasts and Windows Internet Name Service (WINS).
- **-R** Purges the local computer's NetBIOS name cache of all entries and reloads the Lmhosts file.
- **-s** Displays a list of the computer's currently active NetBIOS settings (identifying remote computers by name), their current status, and the amount of data transmitted to and received from each system.
- **-S** Displays a list of the computer's currently active NetBIOS settings (identifying remote computers by IP address), their current status, and the amount of data transmitted to and received from each system.
- **-RR** Sends name release requests to WINS, then starts refresh.

The NetBIOS cache listing as displayed by Nbtstat.exe -c on a computer running Windows XP appears as follows:

Local Area Connection:

Node IpAddress: [192.168.2.11] Scope Id: []

NetBIOS Remote Cache Name Table				
Name	Type	Host Address	Life [sec]	
CZ4	<20>	UNIQUE	192.168.2.2	602
CZ9	<20>	UNIQUE	192.168.2.18	602
CZ11	<20>	UNIQUE	192.168.2.14	602
CZ5	<20>	UNIQUE	192.168.2.12	455
CZ1.ZACKER.LOCA<4C>	UNIQUE	192.168.2.1		172
CZ10	<20>	UNIQUE	192.168.2.21	602
CZ1	<20>	UNIQUE	192.168.2.1	602
COMPAQ-XP	<20>	UNIQUE	192.168.2.27	602

The list of NetBIOS names registered by a computer appears as follows:

Local Area Connection:

Node IpAddress: [192.168.2.11] Scope Id: []

NetBIOS Local Name Table		
Name	Type	Status
CZ8	<00>	UNIQUE
CZ8	<20>	UNIQUE
ZACKER	<00>	GROUP
ZACKER	<1E>	GROUP

True or false: Nbtstat.exe is typically used to troubleshoot DNS name resolution problems.

Answer: *False*. Nbtstat.exe only displays information about NetBIOS over TCP/IP connections, so its primary troubleshooting application is for NetBIOS name resolution problems.

EXAM TIP Unlike most Windows utilities, the command-line parameters for Nbtstat.exe are case-sensitive. When taking the Network+ exam, be alert for questions that differentiate between parameters using uppercase and lowercase.

True or false: Nbtstat.exe is available on Windows, UNIX, and Linux operating systems.

Answer: *False*. UNIX and Linux do not use NetBIOS names, and therefore do not include an implementation of Nbtstat.exe.

Troubleshoot using netstat

Netstat is a command-line program that displays status information about the current network connections of a computer running TCP/IP and about the traffic generated by the TCP/IP protocols. In UNIX and Linux, the program is called netstat, and in Windows, it is called Netstat.exe. The command-line parameters differ for the various implementations of Netstat, but the information they display is roughly the same.

The syntax for Netstat.exe on Windows is as follows:

```
netstat [interval] [-a] [-p protocol] [-n] [-e] [-r] [-s]
```

- **interval** Refreshes the display every interval second until the user aborts the command.
- **-a** Displays the current network connections and the ports that are currently listening for incoming network connections.
- **-p protocol** Displays the currently active connections for the protocol specified by the protocol variable.
- **-n** Causes the program to identify computers using IP addresses instead of names, when combined with other parameters.
- **-e** Displays incoming and outgoing traffic statistics for the network interface. The statistics are broken down into bytes, unicast packets, non-unicast packets, discards, errors, and unknown protocols.
- **-r** Displays the routing table plus the current active connections.
- **-s** Displays detailed network traffic statistics for the IP, ICMP, TCP, and UDP protocols.

Part of the default network connection listing displayed by Netstat.exe –a –n on a computer running Windows 7 appears as follows:

Active Connections

Proto	Local Address	Foreign Address	State
TCP	0.0.0.0:135	cz12:0	LISTENING
TCP	0.0.0.0:445	cz12:0	LISTENING
TCP	0.0.0.0:554	cz12:0	LISTENING
TCP	0.0.0.0:1025	cz12:0	LISTENING
TCP	127.0.0.1:1039	cz12:5354	ESTABLISHED
TCP	127.0.0.1:1088	cz12:0	LISTENING
TCP	127.0.0.1:1138	cz12:27015	ESTABLISHED
TCP	127.0.0.1:1140	cz12:19872	ESTABLISHED

```

TCP 127.0.0.1:5354          cz12:0          LISTENING
TCP 192.168.2.9:1141        v-client-4b:https  CLOSE_WAIT
TCP [::]:135                cz12:0          LISTENING
TCP [::]:445                cz12:0          LISTENING
TCP [::]:554                cz12:0          LISTENING
UDP 0.0.0.0:123             *:*             LISTENING
UDP 0.0.0.0:500             *:*             LISTENING
UDP [fe80::7441:4473:f204:ec1d%10]:1900  *:*
UDP [fe80::7441:4473:f204:ec1d%10]:55010  *:*

```

The interface statistics produced by Netstat.exe –e on a computer running Windows 7 look like this:

Interface Statistics

	Received	Sent
Bytes	2127854975	751579877
Unicast packets	3151802	2833248
Non-unicast packets	64418	995
Discards	0	0
Errors	0	0
Unknown protocols	84	65

The routing table display produced by Netstat.exe –r appears as follows:

```
=====
Interface List
15...00 26 c7 7e 00 e1 .....Microsoft Virtual WiFi Miniport Adapter
12...00 26 c7 7e 00 e0 .....Intel(R) WiFi Link 1000 BGN
10...60 eb 69 93 5e e5 .....Realtek PCIe GBE Family Controller
1..... ....Software Loopback Interface 1
=====
```

IPv4 Route Table

Active Routes:

Network	Destination	Netmask	Gateway	Interface	Metric
0.0.0.0	0.0.0.0	192.168.2.99	192.168.2.9	192.168.2.9	10
127.0.0.0	255.0.0.0	On-link	127.0.0.1	127.0.0.1	306
127.0.0.1	255.255.255.255	On-link	127.0.0.1	127.0.0.1	306
127.255.255.255	255.255.255.255	On-link	127.0.0.1	127.0.0.1	306
192.168.2.0	255.255.255.0	On-link	192.168.2.9	192.168.2.9	266
192.168.2.9	255.255.255.255	On-link	192.168.2.9	192.168.2.9	266
192.168.2.255	255.255.255.255	On-link	192.168.2.9	192.168.2.9	266
224.0.0.0	240.0.0.0	On-link	127.0.0.1	127.0.0.1	306
224.0.0.0	240.0.0.0	On-link	192.168.2.9	192.168.2.9	266
255.255.255.255	255.255.255.255	On-link	127.0.0.1	127.0.0.1	306
255.255.255.255	255.255.255.255	On-link	192.168.2.9	192.168.2.9	266

IPv6 Route Table

Active Routes:

If	Metric	Network	Destination	Gateway
1	306	::1/128		On-link
10	266	fe80::/64		On-link

```
10      266 fe80::7441:4473:f204:ec1d/128  On-link
  1      306 ff00::/8                      On-link
10      266 ff00::/8                      On-link
```

The command-line parameters for the UNIX/Linux netstat tool are similar to those used in Netstat.exe, but not identical. The UNIX/Linux version has additional parameters. The default netstat display on a UNIX or Linux system is shown in Figure 4-8.

```
[root@localhost ~root]# netstat
Active Internet connections (w/o servers)
Proto Recv-Q Send-Q Local Address          Foreign Address        State
Active UNIX domain sockets (w/o servers)
Proto RefCnt Flags       Type      State          I-Node Path
unix  10  [ ]      DGRAM    980  /dev/log
unix  2  [ ]      DGRAM    1340
unix  2  [ ]      DGRAM    1298
unix  2  [ ]      DGRAM    1264
unix  2  [ ]      DGRAM    1236
unix  2  [ ]      DGRAM    1162
unix  2  [ ]      DGRAM    1128
unix  2  [ ]      DGRAM    1023
unix  2  [ ]      DGRAM    992
unix  2  [ ]      STREAM   CONNECTED  490
[root@localhost ~root]#
```

FIGURE 4-8 The default netstat display.

To display the statistics for the system, you run netstat with the **-s** parameter. The statistics display is shown in Figure 4-9.

```
[root@localhost ~root]# netstat -s | more
Ip:
  6129 total packets received
    0 forwarded
    0 incoming packets discarded
  76 incoming packets delivered
  38 requests sent out
Icmp:
  4 ICMP messages received
  0 input ICMP message failed.
  ICMP input histogram:
    echo requests: 4
  10 ICMP messages sent
  0 ICMP messages failed
  ICMP output histogram:
    destination unreachable: 6
    echo replies: 4
Tcp:
  3 active connections openings
  0 passive connection openings
  0 failed connection attempts
  0 connection resets received
  0 connections established
  17 segments received
--More--
```

FIGURE 4-9 The netstat statistics display.

To display the system's current connections, you run netstat with the **-l** parameter. The connections display is shown in Figure 4-10.

True or false: The netstat utility can help administrators to detect Trojan Horse attacks.

Answer: *True*. By displaying information about the ports in use on a computer, netstat can sometimes locate an unauthorized program on a zombie computer listening for incoming activity from a controller on the Internet.

EXAM TIP The routing table display produced by the netstat -r command is identical to that generated by the route print command. Network+ exam candidates should be aware that the two commands are interchangeable.

```
[root@localhost ~root]# netstat -l | more
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address          Foreign Address        State
tcp     0      0 *:32768                 *:*
tcp     0      0 *:sunrpc                *:*
tcp     0      0 localhost.localdom:smtp  *:*
udp     0      0 *:32768                 *:*
udp     0      0 *:847                  *:*
udp     0      0 *:sunrpc                *:*
Active UNIX domain sockets (only servers)
Proto RefCnt Flags     Type      State         I-Node Path
unix  2      [ ACC ]   STREAM   LISTENING    1295   /tmp/.font-unix/fs7100
unix  2      [ ACC ]   STREAM   LISTENING    1261   /dev/gpmctl
[root@localhost ~root]#
```

FIGURE 4-10 The netstat connections display.

True or false: All TCP/IP systems must have an implementation of netstat.

Answer: *False*. While the Windows and UNIX/Linus operating systems include netstat, there is no directive stating that they must include it.

Troubleshoot using Route.exe

All Windows-based operating systems include a command-line program called Route.exe, which you can use to modify the contents of the system's routing table. The syntax for Route.exe is as follows:

```
ROUTE [-f] [-p] [command [destination] [MASK netmask]
          [gateway] [METRIC metric] [IF interface]]
```

- **-f** Deletes all entries from the routing table. When used with the ADD command, deletes the entire table before adding the new entry.
- **-p** Creates a persistent entry in the routing table, when used with the ADD command. A *persistent route* is one that remains in the table permanently, even after the system is restarted. When -p is used with the PRINT keyword, the system displays only the persistent routes in the table.
- **command** Contains one of the following keywords that specifies the function of the command:
 - **PRINT** Displays the contents of the routing table. When used with the -p parameter, displays only the persistent routes in the routing table.

- **ADD** Creates a new entry in the routing table.
- **DELETE** Deletes an existing entry from the routing table.
- **CHANGE** Modifies the parameters of an existing entry in the routing table.
- **destination** Specifies the network or host address of the table entry being managed.
- **MASK *netmask*** Specifies the subnet mask to be applied to the address indicated by the *destination* variable.
- **gateway** Specifies the IP address of the router that the system should use to reach the host or network indicated by the *destination* variable.
- **METRIC** Specifies a value that indicates the relative efficiency of the route in the table entry.
- **IF *interface*** Specifies the number of the network interface that the system should use to reach the router specified by the *gateway* variable.

For example, if you were using the network configuration shown in Figure 4-11 to create an entry that informs Router A of the existence of network 192.168.3.0—which is accessible through Router B—you would execute a Route.exe command like the following at Router A system's command line:

```
ROUTE ADD 192.168.3.0 MASK 255.255.255.0 192.168.2.7 IF 1 METRIC 1
```

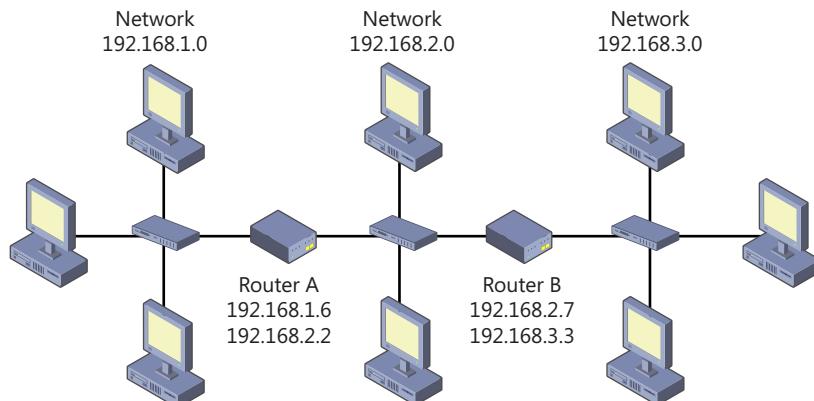


FIGURE 4-11 When the administrator adds a static route to the routing table in the Router A system, Router A can forward packets to the 192.168.3.0 network through Router B.

The functions of the Route.exe parameters in this particular command are as follows:

- **ADD** Indicates that the program should create a new entry in the existing routing table.
- **192.168.3.0** The address of the network to be added to the Router A routing table.
- **MASK 255.255.255.0** The subnet mask to be applied to the destination address.

- **192.168.2.7** The local address of the gateway (Router B) that provides access to the destination network.
- **IF 1** The number of the interface in Router A that provides access to the network containing the specified gateway (Router B).
- **METRIC 1** Indicates that the destination network is one hop away.

This new routing table entry essentially tells Router A that when it has traffic to send to any address on the 192.168.3.0 network, it should send the traffic to the router with the address 192.168.2.7, using the Router A network interface adapter designated by the system as interface 1.

True or false: The following route command causes all traffic destined for the 192.168.2.7 host to be sent to the default gateway on the 192.168.3.0 network.

```
ROUTE ADD 192.168.3.0 MASK 255.255.255.0 192.168.2.7 IF 1 METRIC 1
```

Answer: *False*. The command causes all traffic destined for the 192.168.3.0 network to be sent to the 192.168.2.7 gateway.

EXAM TIP UNIX and Linux operating systems also have a route utility, which performs many of the same functions, but which has a different syntax and parameters. Network+ exam candidates should know that the utility exists on various platforms, but need not memorize the switches for both.

True or false: All systems with routing tables have a mechanism for creating routing table entries from the command line.

Answer: *False*. There is no requirement stating that there must be a means of creating routing table entries from the command line.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which TCP/IP tool uses ICMP messages and manipulates IPv4 Time to Live values to disclose the route packets take through an internetwork?
2. Which utility can you use on a Windows computer to view resource record information on a particular DNS server?
3. What Windows command line utility can you use to release an existing DHCP address lease and request a new one?
4. Which two utilities can display a Windows computer's routing table from the command line?

Objective 4.4: Given a scenario, use the appropriate network monitoring resource to analyze traffic

The activities of a network are, by nature, hidden from view unless you use special tools to observe them. There are a variety of tools that administrators can use to gather information about what is happening on a network and display it for analysis. Some

of these tools require an extensive investment in compliant hardware, whereas others require a considerable level of expertise from the operator. This objective requires you to be familiar with some of the most common network monitoring tools.

Exam need to know

- Monitor a network using SNMP
For example: What components are required for an SNMP network management product?
- Monitor a network using SNMPv2
For example: Why was SNMPv2 not acceptable to many network administrators?
- Monitor a network using SNMPv3
For example: How did SNMPv3 regain the approval of network administrators?
- Monitor a network using syslog
For example: For what application was syslog originally invented?
- Monitor a network using system logs
For example: What events do system logs record?
- Monitor a network using history logs
For example: What events do history logs record?
- Monitor a network using general logs
For example: What events do general logs record?
- Monitor a network using traffic analysis
For example: What conditions can traffic analysis disclose?
- Monitor a network using a network sniffer
For example: How does a sniffer differ from an analyzer?

Monitor a network using SNMP

When an application or an operating system experiences a problem, it usually generates an error message. You can easily monitor these error messages by reviewing logs, but receiving error messages from other network components, such as routers and switches, can be more difficult.

A standalone router doesn't have a screen on which it can display error messages, but it does usually have an administrative interface accessible through a remote connection. However, even with this capability, it is difficult for an administrator who is responsible for dozens or hundreds of devices to monitor them all. In this case, it is possible to arrange for many networking devices to supply administrators with information about their status.

Network management products are designed to provide administrators with a comprehensive view of network systems and processes, by using a distributed architecture based on a specialized management protocol, such as the Simple Network Management Protocol (SNMP).

The Simple Network Management Protocol (SNMP) is a TCP/IP application layer protocol and query language that specially equipped networking devices use to communicate with a central console. Many of the networking hardware and

software products on the market, including routers, switches, network adapters, operating systems, and applications, are equipped with SNMP agents.

An SNMP agent is a software module that is responsible for gathering information about a device and delivering it to a computer that has been designated as the network management console. The agents gather specific information about the network devices and store them as managed objects in a management information base (MIB). At regular intervals, the agents transmit their MIBs to the console by using SNMP messages, which are carried inside User Datagram Protocol (UDP) datagrams. The agents use UDP port 161, and the management console uses port 162.

The network management console processes the information that it receives from the agents in SNMP messages and provides the administrator with a composite picture of the network and its processes. The console software can usually create a map of the interconnections between network devices, as well as display detailed log information for each device. When there is a serious problem, an agent can generate a special message called a trap, which it transmits immediately to the console, causing it to alert the administrator of a potentially dangerous condition. In many cases, the console software can be configured to send alerts to administrators in a variety of ways, including by email and text messaging.

In addition to their network reporting capabilities, network management products can provide other functions, including the following:

- Software distribution and metering
- Network diagnostics
- Network traffic monitoring
- Report generation

Network management products are available with a wide range of abilities, ranging from relatively modest open-source packages to extremely complex and expensive commercial products.

The first version of the SNMP standard, which the Internet Engineering Task Force (IETF) published in 1988 as RFC 1065, RFC 1066, and RFC 1067, provides the protocol's basic functionality but is hampered by shortcomings in security. SNMPv1 messages contain no protection other than a community string, which functions as a password, and which the systems transmit in clear text.

True or false: To effectively monitor a network using SNMP, you must be sure that all of the equipment you purchase when designing and building your network supports the protocol.

Answer: *True*. When you see a network interface adapter, switch, router, access point, or other device that purports to be managed or that claims to have network management capabilities, this usually means that the device includes an SNMP agent.

EXAM TIP The Network+ exam objectives list three versions of the SNMP protocols.

The basic framework of the protocol is the same for all three; the main difference

between the versions is the security mechanism built into each one. Candidates for the exam should be familiar with the basic SNMP structure and the differences between the versions.

True or false: SNMP is not only the name of a protocol; it is also the name of a network management product.

Answer: *False*. SNMP is not the name of a network management product; it is just the name of the protocol that provides a framework for the interaction of the various components in a network management product.

Monitor a network using SNMPv2

SNMPv2, released in 1993, added some improvements in functionality to the original SNMP protocol. One such improvement was a new protocol data unit (PDU) called GetBulkRequest that enables systems to send large amounts of management data in a single message, rather than using multiple GetNextRequest PDUs.

Version 2 also included a new security system that many people criticized as being overly complex. The resistance to this system was so strong that an interim version appeared, called SNMPv2c, which consisted of SNMPv2 without the new security system and with the old version 1 community string instead.

True or false: SNMPv1 and SNMPv2 rely on a community string as their only means of security.

Answer: *False*. SNMPv1 uses a community string, but SNMPv2 does not. The interim version SNMPv2c retains the community string from version 1 in place of the new version 2 security system.

EXAM TIP Candidates for the Network+ exam should be sure to differentiate between SNMPv2 and SNMPv2c.

Monitor a network using SNMPv3

In 2002, the IETF published an SNMP standard with a workable security solution, which became version 3 and was ratified as an Internet standard. SNMPv3 includes all of the standard security services that administrators have come to expect, including authentication, message integrity, and encryption.

True or false: Most of the network management products on the market today support SNMPv3.

Answer: *True*. In addition, many network management products that implement SNMPv3 also include support for the earlier, unprotected versions, such as SNMPv1 and SNMPv2c.

Monitor a network using syslog

One of the oldest tools for the generation of logs is called syslog. Syslog is a protocol designed to send log entries generated by a device or process called a facility across an IP network to a message collector, called a syslog server.

Syslog messages begin with two numerical codes. The first code identifies the facility that generated the message, some examples of which are as follows:

- **0** Kernel messages
- **1** User-level messages
- **2** Mail system
- **3** System daemons
- **4** Security/authorization messages
- **5** Messages generated internally by syslog
- **6** Line printer subsystem
- **7** Network news subsystem
- **8** UUCP subsystem
- **9** Clock daemon
- **10** Security/authorization messages
- **11** FTP daemon
- **12** NTP subsystem

The second code uses the following values to specify the severity of the message:

- **0** Emergency: system is unusable
- **1** Alert: action must be taken immediately
- **2** Critical: critical conditions
- **3** Error: error conditions
- **4** Warning: warning conditions
- **5** Notice: normal but significant condition
- **6** Informational: informational messages
- **7** Debug: debug-level messages

Subsequent parts of the syslog message format include a structured data section, which contains information in a format easily parsable by other software entities, and a free-form message section intended to carry more specific information about the event.

True or false: Virtually all UNIX and Linux distributions include a version of syslog.

Answer: *True*. Syslog has become ubiquitous on the UNIX and Linux platforms and is common on other operating systems as well, including Windows.

EXAM TIP Candidates for the Network+ exam must know what syslog does and that it is primarily a UNIX/Linux utility, but they do not have to know how applications use it.

True or false: Syslog is the means by which an application can write information to a log file.

Answer: *True*. Syslog was developed in the 1980s for use with sendmail, the de facto standard in Simple Mail Transfer Protocol (SMTP) mail servers. It enables sendmail to record its activities in a text-based log file. Since then, it has come to be used for the same purpose with a variety of applications and processes.

Monitor a network using system logs

In UNIX and Linux operating systems, logs are traditionally text files, but the Windows operating systems use a graphical application called Event Viewer to display the log information gathered by the operating system and certain applications running on it.

All computers running Windows maintain the same basic logs. These are essentially system logs that display information about operating system events, application activities, and serious system errors. Servers performing certain roles have additional logs, such as those tracking directory service and replication activities.

The primary function of the Windows Eventing engine is to record information about system activities as they occur and package that information in individual units called events. When you launch the Event Viewer console, you see the Overview and Summary display shown in Figure 4-12.

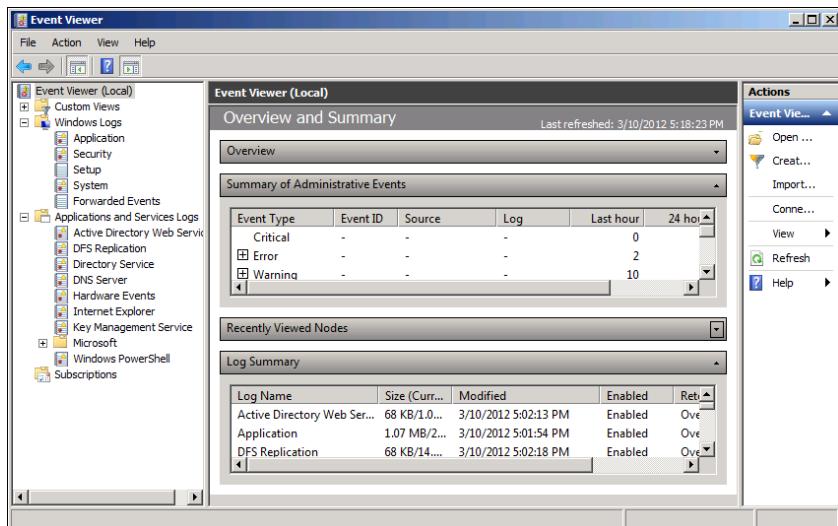


FIGURE 4-12 The Overview and Summary display in the Event Viewer console.

True or false: System logs are the first place a network administrator should look for information about server activities.

Answer: *True*. System logs track a server's startup activities and the ongoing status of device drivers and services. When a fault or a change in status occurs, system logs can tell you exactly what happened and when.

EXAM TIP Network+ exam candidates should be aware that Windows system logs are accessible through the Event Viewer console, and that on many UNIX or Linux systems, the logs are located in the `/var/log` directory.

True or false: The Windows Event Viewer console classifies entries into Information, Error, Warning, and Critical events

Answer: *True*. Windows system logs do not contain only reports of problems; they also include informational events that describe the status of ongoing processes.

Monitor a network using history logs

History logs are records of users' activities on a system or on a network. Web browsers typically maintain a history of all the sites users access, but network administrators rarely have need to consult these. Of more interest to the administrator is the ability of a Windows server to maintain a history of the network resources users access.

In Windows, to create this type of history log, an administrator must activate the operating system's auditing capability and select the type of activities to record. When auditing is active, the system records the selected activities to the security log, accessible through the Event Viewer console.

True or false: Windows systems can audit both successful and failed activities.

Answer: *True*. An administrator can choose to audit logon successes and failures, in which case the log will contain a record of each user's logon times. If an unauthorized individual attempts to access an account by trying different passwords, the failed logon attempts are in the log also.

EXAM TIP Network+ exam candidates should be aware of the distinction between system logs and history logs.

True or false: Web servers maintain history logs that identify clients who connect to the server and the pages they access.

Answer: *True*. All web servers maintain logs of client activities; administrators use these logs to study traffic patterns and document the server's visitors and hits.

Monitor a network using general logs

General logs typically contain information about application updates and similar events, which can sometimes be useful to the network troubleshooter looking for what changes to the system configuration might be the cause of a new problem. In Windows, the Event Viewer console contains individual logs for many of the system's features, services, and applications.

True or false: General logs record the application of patches and updates.

Answer: *True*. General logs are primarily informational, and enable administrators to track changes to application installations and see a history of status updates.

Monitor a network using traffic analysis

Traffic analysis is a technique for deriving information based on the pattern and frequency of messages transmitted over a network, rather than by their contents. For example, a marked increase in the number of incoming packets addressed to a web server is significant in itself, even if you do not know the contents of the packets.

Generally speaking, you perform a traffic analysis using a tool that either monitors or captures network traffic for itself, or works with logs you already have to create a graphical representation of your network traffic.

True or false: To analyze website traffic you use an application that interprets the web server's log files.

Answer: *True*. Most web servers maintain text logs of all hits and visits. The logs contain a wealth of information, but they provide little insight to the administrator in their raw form. A traffic analysis application reads the logs files and presents the information in a more coherent manner, usually in the form of tables and graphs.

EXAM TIP Network+ exam candidates should be aware that traffic analysis is a technique, not the product of any one particular tool. Administrators can use nearly any kind of networking monitoring utility to analyze traffic patterns; some, however, are specialized tools designed to detect specific phenomena.

True or false: Traffic analysis can help to disclose attacks against your systems.

Answer: *True*. For example, a traffic sample that contains an inordinate number of ICMP packets can be a sign that a denial of service attack is in progress.

Monitor a network using a network sniffer

There is often confusion amongst administrators as to what differentiates a network or protocol analyzer and a network sniffer. Analyzers and sniffers both capture network traffic samples. The difference between the two lies in what the application does with the data once it is captured.

Analyzers read the internal contents of the packets they capture, parse the individual data units, and display information about each of the protocols involved in the creation of the packet. Sniffers, on the other hand, look for trends and patterns in the network traffic without examining the contents of each packet.

True or false: Some network monitoring products are both analyzers and sniffers.

Answer: *True*. There are network monitoring products that can display network traffic trends as well as analyze individual packets.

EXAM TIP Technically speaking, analyzers and sniffers perform different functions. But in practice, as well as on the Network+ exam, the terms network sniffer, network analyzer, packet sniffer, and protocol analyzer are often used interchangeably.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which versions of the Simple Network Management Protocol do not include any security protection other than a clear text community string?
2. What program was created to provide logging services for the UNIX sendmail program?
3. What type of log does a web server maintain, recording all of the client visits and hits?
4. What is the primary difference between a protocol analyzer and a network sniffer?

Objective 4.5: Describe the purpose of configuration management documentation

Documentation is a critical part of all network administration and troubleshooting tasks. Knowing what you have done previously, and what others have done, prevents people from needlessly repeating the same actions and pursuing incorrect theories. This objective covers some of the most important types of network documentation that all administrators should work with on a frequent basis.

Exam need to know

- Describe the purpose of wire schemes
For example: What information should be found on a wire scheme?
- Describe the purpose of network maps
For example: What applications can create network maps?
- Describe the purpose of documentation
For example: Why is it necessary to document network administration activities?
- Describe the purpose of cable management
For example: Why are cable management policies necessary in a data center?
- Describe the purpose of asset management
For example: What types of tasks are stipulated by asset management policies?

- Describe the purpose of baselines
For example: What do you do with a baseline performance level after you have captured it?
- Describe the purpose of change management
For example: What types of tasks are stipulated by change management policies?

Describe the purpose of wire schemes

Documentation of your network's cable installation is particularly important, both because much of it is probably hidden from view and because your organization probably had an outside contractor install it. The purpose of having this documentation is so that if something goes wrong with a cable run, or if you want to expand the network, you know where the existing hardware is and don't have to go poking holes in walls and lifting ceilings for nothing.

In many cases, the best way to ensure that you have all of the documentation you need is to begin with the five classic questions posed by journalists: who, what, where, when, and how. In the case of a cable installation, ask yourself the following:

- **Who installed it?** If the cables were installed by an outside contractor, you must have contact information for them and copies of the original contract. Every aspect of the arrangement should be documented; no oral agreements.
- **What was installed?** Your documentation should include a complete list of all the hardware used in the cable installation, including the bulk cable itself and all connectors, wall plates, patch panels, and other components. Save receipts and invoices attesting to the rating of the cable components. If a contractor agrees to use CAT6 hardware throughout the installation, that contractor should provide documentation to prove that CAT6 was indeed used.
- **Where was it installed?** A wiring schematic or cabling diagram is essential to your document collection. Cable installers must document the exact path of every cable run through walls, floors, and ceilings. The best way to accomplish this is to obtain a copy of the original plan or blueprint for the site and add the cable runs to it. Documents should also record the numbers assigned to each cable end and connector (because that sticky note attached to the patch panel might someday fall off).
- **When was it installed?** For warranty purposes and to track conformance to ever-changing standards, you should record when your cables were installed, especially if you installed different parts of the network at different times.
- **How was it installed?** It is essential to record the decisions made during the cable installation process, such as whether the pinouts conform to the

T568A or T568B standard. This enables you to ensure that future cabling work in the network conforms to the same standards.

True or false: A wire scheme should indicate the complete route of every internal cable run from wall plate to patch panel.

Answer: *True*. The main purpose of a wire scheme is to indicate where cables are located in walls and ceilings.

EXAM TIP The Network+ exam objectives have used differing terminology for this information throughout the years. The most recent objectives refer to "wire schemes," whereas the N10-004 objectives called them "wiring schematics." Be prepared to see either term on the exam.

True or false: Wire schemes are a troubleshooting tool.

Answer: *True*. When you experience problems on a cabled network, studying the wiring scheme can provide clues as to possible causes. For example, cable runs that are overlong, run too near sources of interference, or use incorrect components can cause transmission problems.

EXAM TIP The Network+ exam usually contains at least one question requiring the candidate to study and evaluate a wiring scheme.

Describe the purpose of network maps

The terminology is not always consistent, but a network diagram differs from a cable diagram in that its intention is to illustrate the relationships between network components; it is not usually drawn to scale and does not necessarily include architectural elements of the site, such as walls, ducts, and fixtures.

What a network diagram does have is a representation of every device and component on the network and all the connections between them. This means that the diagram includes not only computers, but all of the switches, routers, access points, wide area network (WAN) devices, and other hardware components that make up the network infrastructure, as shown in Figure 4-13.

There are numerous software tools you can use to create network diagrams, the most popular of which is probably Microsoft Visio. In most cases, these products use generic icons to represent network hardware components, but there are packages that provide genuine depictions of specific products, enabling you to create a realistic diagram of the racks in your data center, for example.

Network diagrams typically specify only the names assigned to components and depict the connections between them. A network map provides more detail, such as the IP addresses and/or hardware addresses of each component and the speeds at which links operate.

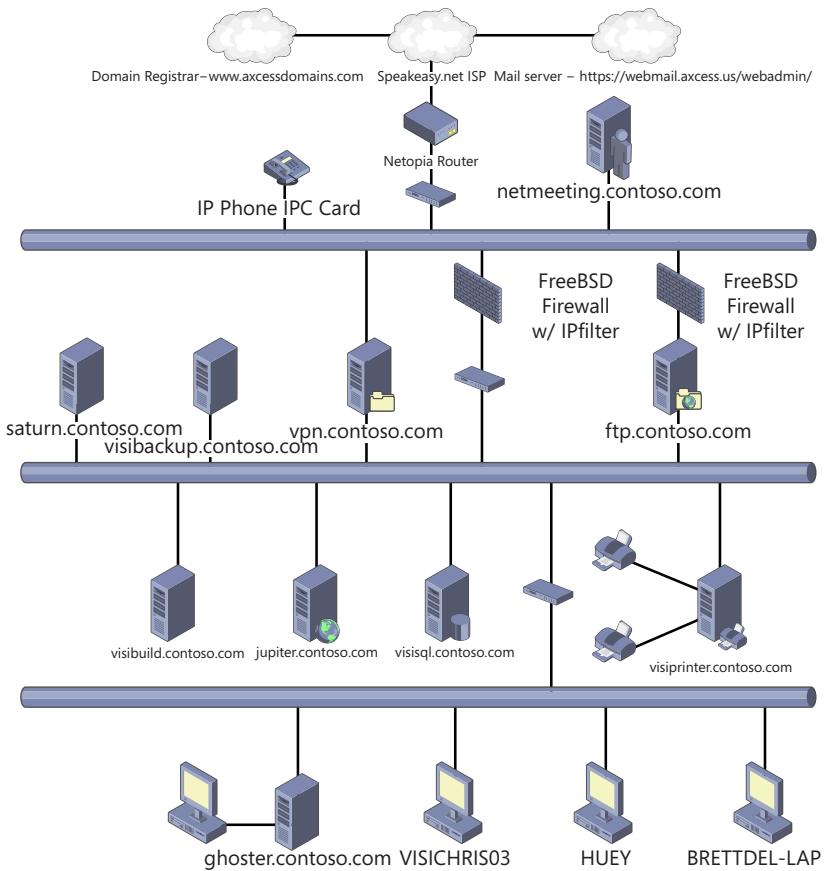


FIGURE 4-13 A network diagram created with Microsoft Visio.

There are several utilities that can automatically create a network map by scanning a network and reading its properties. A map created by the graphical interface for the Nmap utility is shown in Figure 4-14.

True or false: Network maps are typically drawn to scale.

Answer: *False*. Network maps illustrate the relationships between devices and provide information about the links connecting them, but they usually do not indicate the exact location of each device.

True or false: Network maps typically contain more information than network diagrams.

Answer: *True*. Network diagrams specify device types and connections, but network maps can also include IP addresses, link speeds, and other information.

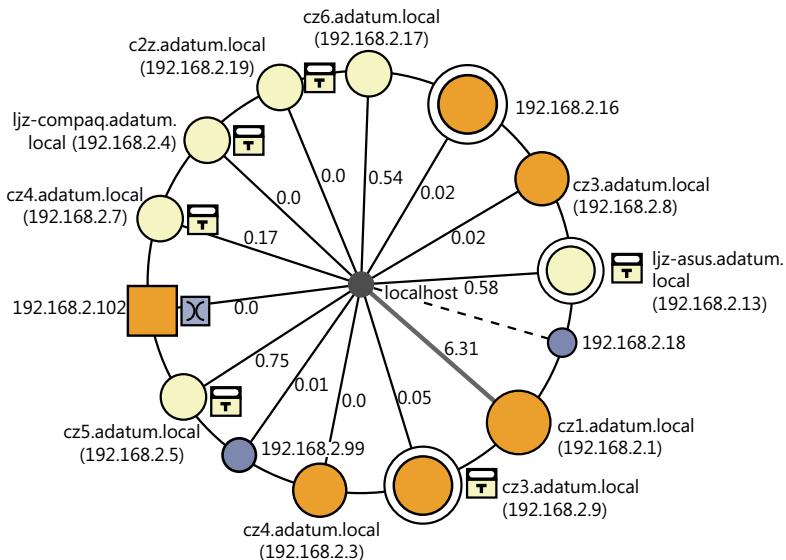


FIGURE 4-14 A network map created with Nmap.

Describe the purpose of documentation

When maintaining a home or small office network, it is sometimes possible for an administrator to work “on the fly,” dealing with issues as they arise, solving problems as they happen, and keeping all of the details about the network in his or her head. When you get beyond a four- or five-node network, however, this method becomes increasingly unmanageable. You begin to forget some of the details, things slip by that you should have remembered, and you find yourself repeating tasks, unnecessarily.

In truth, this network management philosophy is impractical and unprofessional for even the smallest network. Documentation is a critical part of any network management plan, and the time to start thinking about it is well before you install your first network hardware. The planning phase of the network must also be documented, so that the people who have to work on it later know what has been done.

There are many types of documentation that network administrators use and maintain. How and in what form you choose to create these documents is a matter of personal preference and company policy, but the important factor is that everyone involved in the network management and administration processes knows where the documentation is and can access it.

True or false: Network documentation must be read/write accessible to all personnel working on the network.

Answer: *False*. While universal accessibility is definitely desirable, there are some individuals who should not be able to access documentation, including temporary workers and non-IT computer users.

Describe the purpose of cable management

To the unaccustomed visitor, even a well-organized data center can seem like a confused clutter of multicolored cables running between racks, overhead, and under the floor. To keep a complex data center or wiring closet organized, there must be cable management policies in place that specify how cables and jacks are labeled.

True or false: On a network with proper cable management policies in place, an administrator should be able to look at the identifying number on a wall plate anywhere in the building and know where the corresponding patch panel port is located.

Answer: *True*. A proper numbering system should identify the IDF where each cable run leads, as well as providing more specific information, such as rack and patch panel numbers.

EXAM TIP Candidates for the Network+ exam should be conscious of the distinction between cable management policies and the wire schemes and cable diagrams discussed earlier in this chapter.

Describe the purpose of asset management

When a new software product is released, and it looks as though you might have to upgrade the hardware in some of your computers to run it, how do you know for sure which computers need the upgrade and which can support it without? There are various tools that can inventory the hardware in your computers, but how many of them can give you all the information you need?

The best way to keep track of your computers and their configurations is to document them yourself. Large enterprise networks typically assign their own identification numbers to their computers and other hardware purchases as part of an asset management process that controls the entire life cycle of each device, from recognition of a need to retirement and disposal.

The record for each device should contain all available information about it, including the original documentation for the computer, an inventory of its internal components, and detailed information about its software configuration. This way, anyone seeking to upgrade or troubleshoot the computer can find out what's inside without having to travel to the site and open the case.

The record for each computer should also document any changes that administrators make to it, whether in the hardware or software, so that the information is continually updated.

True or false: Most hardware inventory scanners can provide all of the information you need to determine whether a computer can support a hardware upgrade.

Answer: *False*. Hardware scanners can inventory the components that are inside a computer, but they can't necessarily tell what expansion capabilities are available, such as drive bays or expansion slots.

EXAM TIP The Network+ exam objectives include the term asset management as a type of configuration management documentation, but they do not specify a method or a product for managing assets. As a result, exam questions are likely to be general ones about recognizing the need, rather than about implementing a solution.

True or false: Asset management concerns hardware tracking only.

Answer: *False*. IT departments have assets other than hardware. An asset management solution should also specify the software installed on a computer, as well as configuration settings.

Describe the purpose of baselines

One of the basic principles of network management is to observe and address any changes that might occur in the performance of a system, whether it is a computer or a network. You do this by comparing the system's performance levels at various times. A baseline is the starting point for these comparisons.

You can use any tools or any criteria that you want to measure performance. Windows includes a Performance Monitor tool that can display information about hundreds of different system and network performance characteristics, called counters, as shown in Figure 4-15. There are also many third-party products.

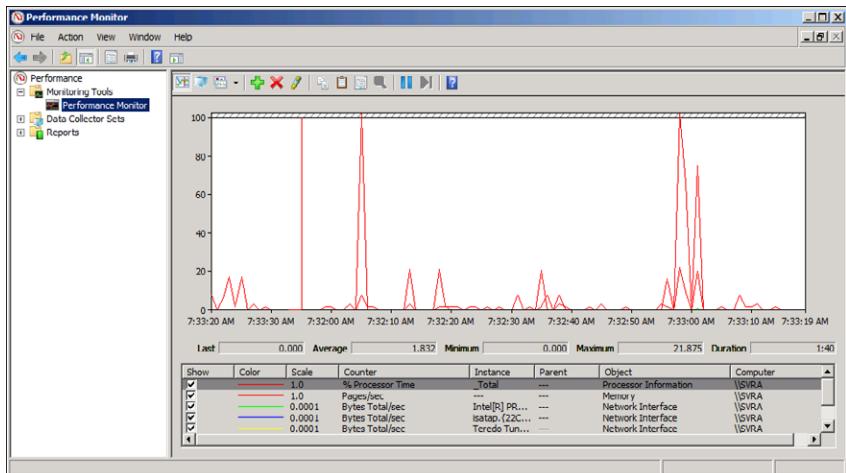


FIGURE 4-15 A Windows Performance Monitor graph.

In addition to displaying performance data in real time, as shown in the figure, Performance Monitor can also capture data to log files—called data collector sets, in some versions—over extended periods of time. To capture an effective baseline,

you might want to capture data over the course of several hours, days, or even longer.

What is most important is that you document both your exact testing procedure and the initial results of your tests, which will function as your baseline for future comparisons. These documents should become part of the permanent record for the system.

At regular intervals, you should then repeat your tests, using the same tools and the same procedures, and compare the results to your previous ones. If you find major discrepancies between your new results and your earlier ones, you should make an effort to determine why. This basic technique can help you to identify trends in performance that enable you to address problems before they become severe.

True or false: A baseline consists of the performance levels that you expect networking hardware to be capable of achieving, based on published standards and the manufacturers specifications.

Answer: *False*. A baseline should consist of actual readings of system performance under real world conditions that are easily reproducible.

EXAM TIP The Network+ exam does not require candidates to know precisely how to operate Performance Monitor or any other specific data capturing tool. Candidates should, however, understand the concept of the baseline and how it can benefit them as administrators.

Describe the purpose of change management

A properly documented network also has written policies regarding how things are supposed to be done. For example, when an administrator troubleshoots a computer and in doing so replaces a hard drive, there should be more to it than taking the drive out of the box and installing it in the computer. There should be a change management policy that leads the administrator through all the ancillary tasks related to the hard drive replacement.

For example, the administrator might have to update the parts inventory to show one less drive in stock; check the warranty status of the failed drive and, if necessary, file a claim; update the computer's record with the serial number and characteristics of the new drive; rebuild the user's local data from backups; and any number of other related tasks.

The same sort of documents should be on file for network-related tasks, including expansions and upgrades, so that the policies used to build the network in the first place are maintained throughout its life cycle.

True or false: Change management policies ensure that all network documentation remains current.

Answer: *True*. All network documentation is a work in progress. Change management provides instructions that outline the modification that must be made to network

records in response to the regular activities of network administrators and troubleshooters.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which Windows application would you most likely use to create a baseline of system or network performance?
2. What is the term usually applied to a graphical representation of network devices, automatically compiled, and containing information such as IP addresses and connection speeds?
3. What is the reason for having a cable management policy?

Objective 4.6: Explain different methods and rationales for network performance optimization

Network performance is a constant concern for the administrator. When the network slows down, the administrator's phone starts ringing, even when the problem is beyond his or her control. To try to avoid these fluctuations, there are several technologies that administrators can use to optimize network performance. Basically, the objective is to keep the data that users need readily available to them, which administrators can do by concentrating on the following tasks:

- Anticipating needs
- Prioritizing traffic
- Providing redundancy

There are a variety of technologies that address these tasks, some of which are covered in this objective.

Exam need to know

- Explain different methods of network performance optimization
For example: How do the following performance optimization methods work: QoS, traffic shaping, load balancing, high availability, caching engines, fault tolerance, and CARP?
- Explain different reasons for network performance optimization
For example: Why are the following elements reasons for optimizing network performance: latency sensitivity; high bandwidth applications, such as VoIP, video applications, and unified communications; and uptime?

Explain different methods of network performance optimization

There are several ways that administrators can optimize network performance, as covered in the following sections.

Quality of Service

Administrators can regulate traffic on a network so that certain types of data receive priority access to the available bandwidth. The easiest way to ensure that data types requiring large amounts of bandwidth receive it is to overprovision the network to accommodate the largest conceivable requirements, so that there is never any traffic congestion. This can be an expensive proposition, however.

Another way is to use Quality of Service (QoS), a method for assigning priorities to various types of traffic. Administrators control the types of traffic that receive higher priorities by specifying protocols, ports, IP addresses, users, bit loss rates, or other criteria. The prioritization is achieved in one of the following ways:

- **Integrated services (IntServ)** A method in which applications communicate with routers to reserve a specific amount of bandwidth. The devices communicate by using a specialized protocol, such as the Resource Reservation Protocol (RSVP).
- **Differentiated services (DiffServ)** A method that uses bits in the IP header's Type of Service (TOS) field to identify the priority to be assigned to the data in the packet.

At the present time, the DiffServ method is proving to be more popular—especially with Internet traffic—because it requires less overhead and less participation from routers.

Traffic shaping

Traffic shaping is a means of prioritizing packets without prior negotiation between applications and routers and without tagging packets. Implemented in an end system, a router, or a switch, traffic shaping is essentially a system that delays packets by storing them in a temporary buffer, so that others conforming to specified criteria receive priority access to the network.

True or false: Traffic shaping is a technology that prioritizes packets by adding tags to the packet headers.

Answer: *False*. Traffic shaping prioritizes traffic by creating queues on end systems, routers, and switches, where the devices temporarily store lower-priority packets until they have finished transmitting the higher-priority ones.

EXAM TIP The Network+ objectives have, at various times, used the terms “traffic shaping” and “bandwidth shaping” when referring to the same technologies.

EXAM TIP In addition to being a performance optimization tool, the Network+ objectives also classify it as a network appliance. For more information on load balancing, see “Objective 4.1 Explain the purpose and features of various network appliances.”

High availability

High availability is a design principle that calls for a system to achieve a previously determined level of performance and reliability. For example, when a service provider enters into a contract with a client, the provider might agree to include a high availability clause stating that the service must be up and running 99.9 percent of the time over the course of a year. This allows for approximately eight hours of downtime per year; otherwise, the contractor must pay a penalty.

Technologies such as Redundant Array of Independent Disks (RAID) are as much high availability solutions as they are fault tolerance solutions. A properly configured RAID array ensures that even when a hard disk fails, all of the data stored on the array continues to be available. To maintain that availability, many RAID solutions enable administrators to replace a drive while the unit is running. This is called *hot swapping*.

Fault tolerance

Fault tolerance is the ability of a system to continue functioning after a failure of some kind. The concept of high availability is similar to that of fault tolerance. In fact, some technologies are capable of providing both.

Servers often have redundant hardware as a hedge against a device failure. Dual power supplies and RAID arrays are common; some organizations even maintain clusters of redundant servers solely for the purpose of fault tolerance.

A cluster is a group of two or more servers bound into a relationship that synchronizes their applications and data. If one of the servers fails, then another immediately takes its place so that the services the cluster provides are always available.

The most common type of fault tolerance, one that every administrator should have, is a recently made backup of all important data. Enterprise networks traditionally use magnetic tape as a backup medium, and many administrators still rely on this technology. With this type of linear medium, the typical practice is to perform periodic full backups of all the data needing protection and, in between full backups, partial backups of any files that have changed.

There are two basic types of partial backups:

- **Incremental** A backup of all of the files that have changed since the last full or incremental backup. To perform a full restore, you must restore the most recent full backup and then each of the subsequent incremental backups, in order.
- **Differential** A backup of all of the files that have changed since the last full backup. To perform a full restore, you must restore the most recent full backup and then just the single most recent differential backup.

Network backup software products typically implement a job schedule and a media rotation scheme that enables the administrator to set up a reliable backup plan that only requires someone to insert the appropriate tape into the drive.

Magnetic tape is an effective backup medium and has its own advantages, but the ever-dropping prices of hard disk storage and the ever-increasing bandwidth of Internet connections has resulted in some new backup strategies.

Hard disk backups are now a practical alternative, due both to the low prices of hard drives and the advent of high-speed external interfaces, enabling administrators to take drives offsite for storage. Another advantage of using hard disks for backups is that they are random access devices. You can create a job that copies only the files that have changed since the last job, and have them simply replace the old versions of the files on the backup disk. You are then left with a full, up-to-date image of the protected data after every job, and no additional incremental or differential backups to restore.

Broadband and other inexpensive, high-speed Internet access services have made the Internet into a viable backup medium. There are now service providers who rent space on well-protected servers, enabling administrators to upload their data instead of backing it up locally. The advantage of this arrangement is that, because the data is stored offsite, it is protected against local fire, theft, and disaster.

Caching engines

Caching is a technique for storing frequently used data in a more available medium, in the anticipation that users will need to access it again. Computers can store data in many different places and at many different speeds. Accessing data from a server on the Internet is relatively slow. Data stored on a local network server can be accessed faster. A hard drive in the local computer is faster still, and data stored in local memory is fastest of all. Caching is an attempt to anticipate users' needs by storing the most needed data in the fastest practical medium.

Computers and networks have many caching systems. All contemporary PCs have a level-2 memory cache that is faster than the main system memory array, and all hard disk drives have an on-board cache for frequently used data. In TCP/IP networking, the ARP maintains a cache of MAC addresses on every computer, and intermediary DNS servers routinely cache names and IP addresses for repeated use.

There are other systems, called caching engines, which can provide a performance boost for your network's internal users by storing frequently used Internet data on a local network server. Some products implement just the web caching capability from the proxy server concept, solely for the purpose of conserving an organization's Internet bandwidth. Some of these products combine a variety of performance optimization capabilities in one package, including load balancing among multiple caches and selective caching of certain traffic types.

CARP

For networks with multiple proxy servers, there is a way to increase the benefit provided by the caching feature. Protocols such as Cache Array Routing Protocol (CARP) enable clients to route requests for specific URLs to the one proxy server containing that cached data. This way, the proxy servers do not maintain duplicate data in their caches, effectively increasing the amount of cached data available to all of the clients on the network.

True or false: Many of the same technologies that provide fault tolerance also provide high availability.

Answer: *True*. Technologies such as RAID enable systems to survive a component failure, but they also ensure that the system continues to function reliably.

EXAM TIP CompTIA, in the Network+ objectives, refers to the caching server products as caching engines. Be familiar with both terms for the exam.

Explain different reasons for network performance optimization

Networks today often carry more than just programs and document files; they also carry voice and video of various types, as well as other kinds of high-bandwidth traffic. These traffic types can have different priorities to the organization running the network, and some are more latency-sensitive than others, purely by their nature. Latency is the delay incurred as traffic travels across a network.

For example, a brief delay while downloading a large image file might be moderately annoying, but constant stutters during an important video conference can be positively infuriating and can also make a company look bad to its clients, partners, or stockholders. This is what is meant by "latency-sensitive;" streaming video can tolerate a few lost bits, but after a certain point, they become evident to the user.

Sensitivity is not solely determined by the data type, either. Interruptions during a streaming video are less of a concern when you're watching a music video than when you are making a sales presentation to potential customers.

Network administrators strive for a maximum amount of network uptime, because the alternative means ringing phones, lost productivity, and perhaps an unexpected career change. Optimizing network performance can help to eliminate these delays for critical applications.

True or false: The easiest way to optimize network performance is to increase the network's available bandwidth.

Answer: *True*. It is rarely the most inexpensive way, but upgrading a network to a faster speed provides more bandwidth, enabling it to carry more data.

True or false: The increasing use of Voice over IP is a common reason for wanting to optimize network performance.

Answer: *True*. People are accustomed to a high level of telephone service, and it doesn't take many problems to make a VoIP service unacceptable for business use. Therefore, administrators are concerned with optimizing the performance of their networks.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is the term used to describe the DNS-based method of load balancing?
2. Which of the two main QoS methods uses tags in the IP header to specify the priority of the data in the packet?
3. What is the name of the protocol that enables proxy servers to exchange information among one another?
4. Which type of QoS requires the use of a specialized protocol, such as the Resource Reservation Protocol (RSVP)?

Answers

This section contains the answers to the “Can you answer these questions?” sections in this chapter.

Objective 4.1: Explain the purpose and features of various network appliances

1. The proxy server functions as an intermediary between the clients on the private network and the Internet by substituting its own registered IP address for the private IP addresses in the clients’ Internet service requests. The Internet servers then reply to the proxy server, and the proxy forwards the replies to the clients.
2. When a DNS server is configured with resource records providing multiple IP addresses for a single server name, the server can respond to name resolution requests for that server with a different IP address each time. This enables each IP address to receive an approximately equal portion of the traffic intended for that server.
3. Proxy servers operate at the application layer of the OSI model.

Objective 4.2: Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues

1. A time domain reflectometer (TDR).
2. The most likely cause is an incorrect punch down in one of the connectors. You would repair it by removing the cable from the connector and punching it down again.
3. A tone generator and locator is the most inexpensive cable testing tool that can accomplish the desired task.
4. A linesman’s handset or butt set is only usable on analog networks.
5. You use a punch down tool to connect keystone connectors to cables.

Objective 4.3: Given a scenario, use appropriate software tools to troubleshoot connectivity issues

1. Traceroute.
2. nslookup.
3. Ipconfig.exe.
4. Route.exe and Netstat.exe.

Objective 4.4: Given a scenario, use the appropriate network monitoring resource to analyze traffic

1. Version 1 and version 2c.
2. syslog.
3. A history log.
4. Analyzers display information about each of the protocols involved in the creation of each packet, while sniffers look for trends and patterns in the network traffic.

Objective 4.5: Describe the purpose of configuration management documentation

1. Performance Monitor.
2. A network map.
3. A cable management policy specifies how cable runs, wall plates, and patch panels should be labeled, so that administrators can easily locate specific connections.

Objective 4.6: Explain different methods and rationales for network performance optimization

1. Round robin DNS.
2. Differentiated services.
3. The Cache Array Routing Protocol (CARP).
4. Integrated services (IntServ).

CHAPTER 5

Network Security

The Network Security domain accounts for approximately 19% of the CompTIA Network+ exam, and contains coverage of protocols, applications, and appliances that enable administrators to protect their networks from unauthorized access—whether intentional or not, whether malicious or not.

This chapter covers the following objectives:

- Objective 5.1: Given a scenario, implement appropriate wireless security measures
- Objective 5.2: Explain the methods of network access security
- Objective 5.3: Explain methods of user authentication
- Objective 5.4: Explain common threats, vulnerabilities, and mitigation techniques
- Objective 5.5: Given a scenario, install and configure a basic firewall
- Objective 5.6: Categorize different types of network security appliances and methods

Objective 5.1: Given a scenario, implement appropriate wireless security measures

For this exam objective, you must be familiar with the most important security mechanisms used by IEEE 802.11 networks. On a wireless local area network (LAN), security measures are not optional. Running a wireless network without them leaves it wide open to anyone happening by with a laptop or other mobile device.

Exam need to know

- Implement encryption protocols
For example: What are the differences between WEP, WPA, WPA2, and WPA Enterprise?
- Implement MAC address filtering
For example: How can you protect a wireless network using MAC address filtering?
- Implement appropriate device placement
For example: How is device placement a security precaution?
- Implement appropriate signal strength
For example: Can you have wireless network signals that are too strong?

Implement Encryption Protocols

Wireless networks using the Institute of Electrical and Electronics Engineers (IEEE) 802.11 protocols at the data-link layer are highly susceptible to intrusion. Any person with a wireless-equipped computer can conceivably access the network, simply by wandering into the transmission range of an access point (AP) or ad hoc partner. Encryption of transmissions is therefore imperative on wireless networks, and there are several different protocols available to protect them.

WEP

Wired Equivalent Privacy (WEP) is a wireless security protocol that protects transmitted data by using a shared secret—a text string possessed by the AP and the clients—as an encryption key. Encrypting the data before it is transmitted helps to prevent unauthorized users from accessing the information in the packets.

To use WEP, administrators must configure all of the devices on the wireless network with the same shared secret. The devices use that one key to encrypt and decrypt all of their transmissions. Anyone who gains knowledge of that key can, at the very least, read the contents of the transmitted packets, and at worst, participate on the network.

Unfortunately, some crafty attackers with experience in cryptography developed methods to discover the shared secret of a WEP network by analyzing captured traffic, even though the packets were encrypted. These attackers quickly exploited the weaknesses they discovered in WEP and spread their key penetration techniques in the form of software tools that they released on the Internet. The combination of free WEP-cracking tools, the ease with which attackers can capture wireless traffic, and the increased popularity of wireless networks, led to WEP becoming one of the most frequently cracked network encryption protocols.

Apart from its weak cryptography, the other limitation resulting in WEP's vulnerability is that the protocol standard does not define any mechanism for changing the shared secret. On a wireless network with hundreds of hosts, manually changing the shared secret on each device is a huge task, and one that is highly prone to error. As a result, WEP networks tend to use the same shared secret for long periods, if not indefinitely. In any cryptographic system, the longer the ciphers use the same key, the more likely it is that an attacker will be able to penetrate it. A WEP installation with an unchanging shared secret gives attackers the cryptographic opportunity they need, and all the time they could want to work on infiltrating the network.

If administrators were able to change the shared secret on a regular basis, however, they would prevent an attacker from gathering enough data to crack the key. This would significantly improve WEP's security. There are techniques for dynamically and automatically changing the shared secret to dramatically reduce WEP's weaknesses.

NOTE A standard WEP implementation that uses an unchanging shared key is sometimes referred to as *static WEP*. When a WEP implementation has a mechanism for automatically changing the shared secret, it is called *dynamic WEP*.

If a wireless network relies on a static shared secret for security, you cannot trust WEP to protect the data transmitted over the network or to prevent unauthorized users from accessing the wireless network. However, if you are forced to use static WEP because you have devices on your network that can support nothing else, there are a few things you can do to improve its security, such as the following:

- Use the highest level of encryption possible: 128-bit or 256-bit. Short keys might be sufficient in some encryption scenarios, but WEP's 40-bit encryption is very vulnerable.
- Locate your access points on a perimeter network to restrict access to internal resources. If users need access to the internal network from a wireless network, they can use a VPN connection.
- Position your APs so that wireless connectivity is limited to locations that you can physically secure, such as the interior of your building.

WPA

Even in the best deployment scenarios, WEP still has security weaknesses. WEP uses a separate static key for broadcast packets, which an attacker can analyze to build a map of the network's private IP addresses and computer names. In addition, the frequent renewal of WEP keys places an additional burden on the RADIUS service.

To address these weaknesses, the Wi-Fi Alliance developed a new encryption protocol called Wi-Fi Protected Access (WPA). WPA can use the same authentication mechanisms and encryption algorithms as WEP, which enables hardware manufacturers to support WPA within their existing product designs with only a simple firmware modification. However, WPA virtually eliminates WEP's most exploited vulnerability by using a unique encryption key for each packet.

When you enable WPA, you establish a passphrase that is automatically associated with the dynamically generated security settings. This passphrase is stored on the AP and on each of the networked computers. Only wireless devices with the WPA passphrase can join the network and decrypt network transmissions.

There are two encryption options for WPA, as follows:

- Temporal Key Integrity Protocol (TKIP)
- Advanced Encryption System (AES)

TKIP is the encryption algorithm that WEP uses, and many WPA implementations continue to support it. WPA improves upon WEP's implementation of TKIP, however. WPA with TKIP reuses initialization vectors (IVs) less frequently than WEP with TKIP, and as a result reduces the likelihood that an attacker will collect enough traffic to compromise the encryption. Additionally, WPA with TKIP creates a unique encryption key for every frame, whereas WEP can use the same key for weeks or months.

In WPA version 2, known as WPA2, the protocol includes support for AES, an encryption algorithm that is more secure than TKIP. The use of WPA2 in combination with IEEE 802.1X using RADIUS is known as WPA2-Enterprise.

Although it is possible to upgrade the firmware in existing WEP wireless equipment to support WPA, older equipment cannot be upgraded to support AES.

As a result, wireless networks with older hardware will probably not be able to use WPA2 encryption unless the organization chooses to upgrade the hardware.

True or false: The cryptographic weakness of the original WEP standard was due in part to US export restrictions.

Answer: *True*. The original WEP implementations used a key that was only 40 bits long, because US export restrictions prevented them from being any longer. The shorter the key, the easier it is to crack.

EXAM TIP Candidates for the Network+ exam must be familiar with the differences between WEP and WPA, and with the differences between WPA, WPA2, and WPA2-Enterprise.

True or false: An administrator has to choose one encryption standard for a wireless access point, and all of the clients using that access point must support the same standard.

Answer: *True*. The biggest problem with implementing encryption on a wireless network is that a device can only use one encryption protocol configuration. Wireless encryption is not like PPP or other protocols, in which systems can negotiate the strongest possible encryption for each connection. You can have a network with 100 wireless clients, all of which support WPA2, but if there are 10 older computers that can only use WEP, then the administrator must configure all the devices to use WEP, in order to support those 10 computers. The only alternatives are to upgrade the 10 older computers or set up a separate WEP access point for those 10 (which would still be the weak point of the network).

Implement MAC address filtering

One common technique that many wireless implementations use to make it more difficult for an unauthorized user to connect to the network is to configure APs to allow network access only to a predefined set of media access control (MAC) addresses. Just as with wired Ethernet, manufacturers assign a unique MAC address (or hardware address) to every wireless network interface adapter.

When an administrator configures an AP to use MAC address filtering, the device ignores any messages from wireless adapters with MAC addresses not on the approved list. Although this does improve security, it has significant drawbacks as far as manageability is concerned. First, administrators must manually maintain the list of MAC addresses on the AP, which can be a difficult task when there are a lot of computers or multiple APs involved. Second, APs typically have limited memory and might not be able to store a large organization's complete list of MAC addresses. Third, if an attacker is knowledgeable and determined enough to circumvent your WEP or WPA encryption, the attacker might also be able to identify and spoof—that is, impersonate—an approved MAC address.

True or false: MAC address filtering on wireless networks can be additive or subtractive.

Answer: *True*. Many access points enable administrators to specify a list of MAC addresses that should be granted or denied access to the network. Administrators can therefore deny access to all but the specified systems, or conversely, grant access to all but the specified systems.

MORE INFO For more information on MAC addressing, see “Objective 1.3: Explain the purpose and properties of IP addressing.”

True or false: To create MAC address filters on a wireless network, an administrator must have each computer’s OUI.

Answer: *True*. The organizationally unique identifier (OUI) forms the first half of a computer’s MAC address. So, if you have a computer’s MAC address, then you have its OUI as well.

Implement appropriate device placement

Physical security is difficult to implement for a wireless network, because you cannot protect radio signals with a locked door the way you can with cables. Many wireless access points provide sufficient range to allow users located in adjacent offices, or even outside the building, to connect. This kind of accessibility can furnish attackers with the time they need to penetrate the security protocols protecting the network. However, there are steps that an administrator can take to minimize the chances of intruders accessing the network from outside.

In addition to distance, the effective range of an access point is based on several unquantifiable factors, including climate conditions, proximity to sources of interference, and composition of walls and other structures. Selecting central locations for your APs within your space is generally recommended, but you also want to make sure that you provide coverage for all the spaces in which your users might need to connect.

The perfect configuration, in which all of your interior space is covered and all of the spaces outside your walls are not, is difficult, if not impossible, to achieve. However, by combining strategic placement of APs with careful selection of antennas and regulation of signal strength, a situation close to this ideal is sometimes possible.

True or false: By locating your wireless access points properly, you can minimize the chances of unauthorized users accessing your network.

Answer: *True*. By keeping your access points nearer to the center of your space, you can limit the signals reaching outside of your office or building.

EXAM TIP Questions about wireless device placement on the Network+ exam can concern security, or they can concern constructing a network that provides access to all the clients that need it. For more information on selecting access point locations, see “Objective 2.2: Given a scenario, install and configure a wireless network.”

True or false: Special types of antennae can help to limit wireless coverage to the interior of a building.

Answer: *True*. Directional antennae pointed inward can prevent wireless access points from projecting signals outside of the building.

Implement appropriate signal strength

For wireless network administrators, achieving sufficient signal strength to service all of their clients is often a major concern. However, from a security standpoint, it is possible to have too much signal strength, enabling individuals outside of the organization to access the network. As noted earlier, time is the main enemy of any security mechanism; if you give attackers enough time, they can pick any lock, crack any safe, or penetrate any encryption algorithm.

True or false: To provide the best possible signal strength for clients, wireless network administrators should locate access points as near as possible and deploy the highest gain antennas they can.

Answer: *False*. It is possible to have signals that are too strong on a wireless network, enabling unauthorized users to access the network. The object should be to generate a strong enough signal to service the clients, while containing it within a defensible structure or office space.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. How does MAC address filtering increase the security of a wireless LAN?
2. Which encryption algorithm did the second version of the WPA protocol add to the standard?
3. What differentiates WPA2 from WPA2-Enterprise?
4. What are the three factors discussed in this chapter that weakened WEP?

Objective 5.2: Explain the methods of network access security

Securing the network has always been one of the administrator's many jobs, but in recent years it has become a constantly escalating arms race, with attackers and defenders both equipping themselves with ever-more-powerful technology. Security is now a concern in virtually every task an administrator performs. Every new piece of hardware or software offers a potential threat, and keeping up with the latest defensive strategies and tactics is a matter of necessity. This objective examines some of the protocols and techniques that network administrators use in the course of securing a network.

Exam need to know

- Explain the use of access control lists

For example: How do the following technologies use access control lists: MAC filtering, IP filtering, port filtering?

- Explain the use of tunneling and encryption

For example: How do the following protocols use tunneling and/or encryption: SSL VPN, VPN, L2TP, PPTP, IPsec, ISAKMP, TLS, TLS2.0, site-to-site, and client-to-site?

- Explain the use of remote access

For example: How does remote access use the following protocols: RAS, RDP, PPPoE, PPP, ICA, and SSH?

Explain the use of access control lists

To authorize a user, a computer system typically checks an access control list (ACL) that is stored with the resource being protected. An ACL is a list of users, or groups of users, who are permitted to access a resource, as well as the degree of access each user or group is permitted. Network devices use ACLs for a variety of purposes. Servers have permission systems that use ACLs to restrict access to file systems, registry settings, and directory service objects. Routers use ACLs to restrict the types of traffic that are permitted to enter a network.

Some of the ACLs used by routers are as follows:

- **IP address filtering** IP address filtering lets you limit network access to specific computers. For example, if you have an Internet web server on a LAN with other computers, and you want Internet clients to be able to access only the web server, you can create a filter permitting only those packets addressed to the web server to enter the network from the Internet. You can also use IP address filtering to protect part of a private network. You can create filters that give only certain computers access to the protected LAN, while preventing all others from accessing it.
- **MAC address filtering** MAC address filtering (or filtering hardware addresses) provides the same basic functionality as IP address filtering. However, it is more difficult to spoof a hardware address than it is an IP address, so MAC address filters are inherently more secure than IP address filters. MAC address filtering is rarely used on Internet routers or firewalls, but for internal filtering and especially wireless LANs, MAC address filtering may be a useful means of restricting access to specific resources.
- **Port number filtering** Port number filtering, also known as *service-dependent filtering*, is the most common and flexible type of packet filtering. Because port numbers represent specific applications, you can use them to prevent traffic generated by these applications from reaching a network. For example, to protect a perimeter network containing your company's web servers, you can create filters that allow only traffic using port 80 to enter from the Internet, blocking all other application ports.

MORE INFO For more information on the well-known port numbers assigned to specific applications and services, see “Objective 1.5: Identify common TCP and UDP default ports.”

True or false: IP address filters are easier for attackers to penetrate than MAC address filters.

Answer: *True*. It is simple to change a system’s IP address to bypass a filter, because IP addresses are readily changeable. However, because MAC addresses are read directly from the system hardware, they are more difficult to change. This makes MAC addresses more difficult to penetrate than IP addresses.

EXAM TIP Candidates for the Network+ exam should be familiar with the various applications that use MAC address, IP address, and port filtering.

True or false: Routers connecting peripheral networks (or DMZs) to the Internet are more likely to use ACLs with port filters than with IP address or MAC address filters.

Answer: *True*. Implementing and maintaining IP and MAC address filters on an Internet access router can be painstaking and time consuming. It is much more common for administrators to restrict access to the peripheral networks by allowing access only to traffic destined for specific ports.

Explain the use of tunneling and encryption

Applications and operating systems use security protocols to protect data as it is transmitted over a network. These protocols generally use specific types of data encryption and a mechanism called tunneling to exchange information securely. Tunneling is a method for transferring a payload across an insecure network by encapsulating the frame with an additional header generated by a tunneling protocol.

Some of these encryption and tunneling protocols are discussed in the following sections.

Virtual private networking

A virtual private network (VPN) is a connection between two computers across a network or the Internet that enables them to communicate in a manner that mimics the properties of a dedicated private network. In most cases, a VPN is functionally similar to a WAN technology, except that the Internet functions as the network medium.

Virtual private networking was developed as a means to provide users with a relatively inexpensive method for connecting to a network at a remote location, as shown in Figure 5-1. VPNs enable users working at home or on the road to connect securely to a remote access server by using the routing infrastructure provided by a public network such as the Internet. From the user’s perspective, the VPN is a point-to-point connection between the user’s computer and a remote access server. The nature of the intermediate network (also called the transit network) is irrelevant, because it appears as if the computers are transmitting data over a dedicated private link.

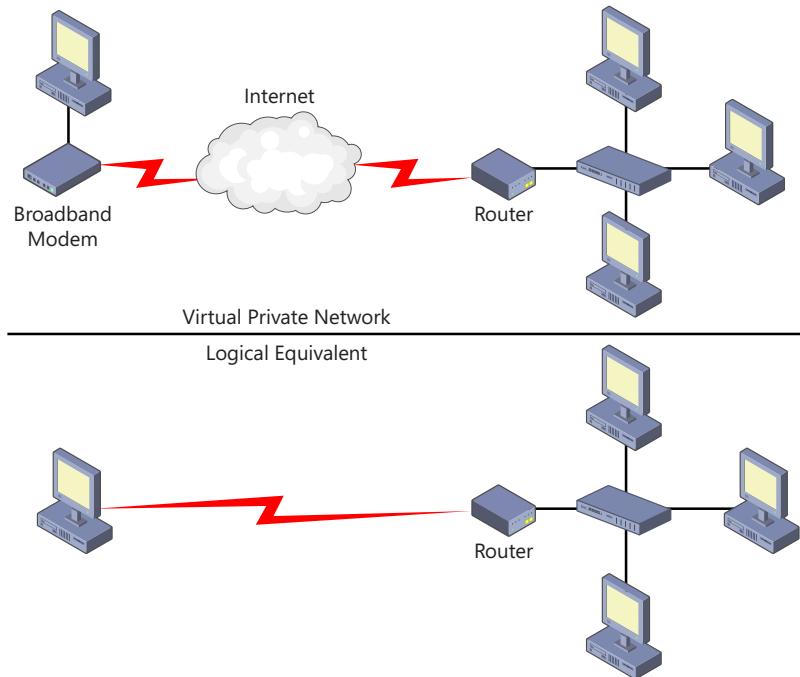


FIGURE 5-1 A VPN connection.

A VPN enables both the remote access client and server computers to connect to the Internet by using a local ISP, which keeps telecommunication charges to a minimum. The client then establishes a connection to the server across the Internet, and the server provides routed access to the corporate network. The connection across the transit internetwork appears to the user as a virtual WAN, providing private network communication over a public internetwork, hence the term *virtual private network*.

The problem with using the Internet for private network communications in this manner is, of course, security. To allow the client and server to exchange confidential information over a public internetwork such as the Internet, the client and server must have a mechanism for securing the data. The method used by VPNs is called tunneling.

Tunneling, also known as *encapsulation*, is a method for transferring a payload across an insecure internetwork infrastructure. The payload might be the frames generated by another protocol, such as Point-to-Point Protocol (PPP), or even a LAN protocol, such as Ethernet. Instead of transmitting the frame as produced by the originating node, the system encapsulates the frame with an additional header generated by a tunneling protocol. The tunneling protocol may also encrypt the original frame. This way, even if someone intercepts the packets as they pass over the Internet, the information inside them remains secure.

The additional header provides routing information so that the encapsulated payload can traverse the intermediate network. The encapsulated packets are then routed between the tunnel endpoints over the transit network, as shown in Figure 5-2. When the encapsulated frames reach the far endpoint, the system de-encapsulates the frames and forwards them to their final destinations.

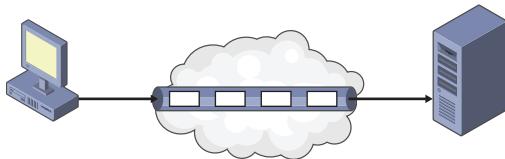


FIGURE 5-2 A VPN tunnel.

This entire process (the encapsulation and transmission of packets) is what is known as tunneling. The tunnel is the logical path through which the encapsulated packets travel across the transit internetwork. The name is derived from the way that the tunneling protocol creates a conduit between two points on the transit network. The original frame produced by the sending computer passes through the tunnel without being accessed or modified in any way, so that the information inside remains intact.

This original type of VPN connection is called a client-to-site connection, because it enables a single remote user to access a home or office network. Some organizations also use VPN connections as a WAN solution to connect two networks together. This is called a site-to-site connection, and it enables any user at either site to access the network at the other site. This is a relatively low-cost branch office WAN solution that some organizations use as a backup to a leased line or other WAN technology.

EXAM TIP The Network+ exam objectives list site-to-site and client-to-site without clarifying their context. Candidates for the exam must know that these are types of VPNs.

The most common tunneling protocols used to create VPNs are the Point-to-Point Tunneling Protocol (PPTP) and the Layer 2 Tunneling Protocol (L2TP), as described in the following sections.

POINT-TO-POINT TUNNELING PROTOCOL (PPTP)

PPTP is an extension of PPP that encapsulates PPP frames into IP datagrams for transmission over an IP network such as the Internet. PPTP can also be used in private LAN-to-LAN networking. PPTP uses a TCP connection for tunnel maintenance and uses modified GRE-encapsulated PPP frames for tunneled data. (GRE stands for Generic Routing Encapsulation.) The payloads of the encapsulated PPP frames can be encrypted and compressed.

PPTP tunnels use the same authentication mechanisms as PPP connections, and also inherit payload encryption and compression capabilities from PPP. PPP encryption provides confidentiality between the endpoints of the tunnel only. If stronger security

or end-to-end security is needed, Internet Protocol Security (IPsec) is the preferred tunneling protocol.

LAYER 2 TUNNELING PROTOCOL (L2TP)

L2TP is a combination of PPTP and Layer 2 Forwarding (L2F), a technology created by Cisco Systems. L2TP is a hybrid of the best features in PPTP and L2F.

L2TP is a network protocol that encapsulates PPP frames to be sent over IP, frame relay, or ATM networks. When utilizing IP as its datagram transport, L2TP can function as a tunneling protocol over the Internet, or it can be used in private LAN-to-LAN networking.

L2TP uses User Datagram Protocol (UDP) and a series of L2TP messages for tunnel maintenance. L2TP also uses UDP to send L2TP-encapsulated PPP frames as the tunneled data. The payloads of encapsulated PPP frames can be encrypted and compressed. Windows systems use IPsec to encrypt the data inside L2TP packets instead of PPP encryption. However, it is possible for other implementations of L2TP to use PPP encryption.

L2TP is similar to PPTP in the way it functions. An L2TP tunnel is created between the L2TP client and an L2TP server. The client might already be attached to an IP network (such as a LAN) that can reach the tunnel server, or the client might have to connect to an ISP to establish IP connectivity.

L2TP tunnel authentication uses the same mechanisms as PPP connections. L2TP inherits PPP compression but not encryption. L2TP does not use PPP encryption because it does not meet the security requirements of L2TP. PPP encryption can provide confidentiality but not per-packet authentication, integrity, or replay protection. Instead, L2TP uses data encryption provided by IPsec.

Both PPTP and L2TP use PPP for point-to-point WAN connections (to provide an initial envelope for the data), and then they append additional headers for transport through the transit network. However, there are some differences between PPTP and L2TP:

- PPTP requires that the transit network use IP. L2TP requires only that the tunnel medium provide packet-oriented point-to-point connectivity. L2TP can run over IP (by using UDP), frame relay permanent virtual circuits, or ATM virtual circuits.
- L2TP provides header compression capability. When header compression is enabled, L2TP operates with 4 bytes of overhead, compared to 6 bytes for PPTP.
- L2TP provides tunnel authentication, whereas PPTP does not. However, when either PPTP or L2TP runs over IPsec, it provides tunnel authentication—making Layer 2 tunnel authentication unnecessary.
- PPTP uses PPP encryption; L2TP does not. The Windows L2TP implementation requires IPsec for encryption.

SSL VPN

Another type of VPN, which does not require any special software on the client side, is called an SSL VPN. SSL VPNs enable a user to connect to a remote site using a standard web browser, with the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocol providing encryption services. To establish the connection, a user connects to an SSL VPN gateway at the remote site, completes an authentication process, and is then able to access one or more resources on the remote network.

IPsec

IPsec is the term used to describe a series of draft standards published by the Internet Engineering Task Force (IETF) that define a methodology for securing data as it is transmitted over a network. Most of the security protocols that encrypt transmitted data are designed for use on the Internet or for specialized traffic between specific types of clients and servers. Until IPsec was developed, there was no standard to provide comprehensive protection for data as it was transmitted over a LAN.

IPsec protects data by digitally signing and encrypting it before transmission. IPsec encrypts the information in IP datagrams by encapsulating it, so that even if someone was to capture the packets, the information inside would remain uncompromised. Using IPsec can protect a network against a variety of threats, including password penetration, compromise of encryption keys, IP address spoofing, and data modification. An unauthorized user with a protocol analyzer application can still capture packets as they are transmitted over the network, but the user cannot do any of the following:

- Read a packet's contents (because they are encrypted)
- Modify a packet's contents without being detected
- Successfully spoof a recipient by assuming another user's identity
- Discover passwords and keys, or reuse encrypted packets

IPsec operates as an extension to the IP protocol at the network layer, so it provides end-to-end encryption. This means that the source computer encrypts the data, and that data is not decrypted until it reaches its final destination. Intermediate systems, such as routers, treat the encrypted part of the packets purely as payload, so they do not have to perform any decryption; they just forward the encrypted payload as is. The routers do not have to possess the keys needed to decrypt the packets, and they do not have to support the IPsec extensions in any way.

Because IPsec operates at the network layer, it can encrypt any traffic that takes the form of IP datagrams—no matter what kind of information those datagrams contain. IPsec is completely invisible to the transport layer protocols encapsulated in the IP datagrams, such as TCP and UDP, and to the applications generating the traffic, because the data is encrypted at the network layer after it leaves the transmitting application and is packaged in the transport layer protocol. Then, at the destination system, IPsec decrypts the traffic before it arrives at the transport layer or the destination application.

IPSEC PROTOCOLS

The IPsec standards define two protocols that provide different types of security for network communications: the IP Authentication Header (AH) and the IP Encapsulating Security Payload (ESP). The IP Authentication Header protocol does not encrypt the data in IP packets, but it does provide the following services:

- **Mutual authentication** Before two computers can communicate by using IPsec, they must authenticate each other to establish a trust relationship. After the computers have authenticated each other, the cryptographic checksum in each packet functions as a digital signature, preventing anyone from spoofing or impersonating one of the computers.
- **Anti-replay** In some cases, intruders can analyze network traffic patterns, determine the functions of certain packets, and use data from captured packets to wage an attack—even when the data in the packets is encrypted. For example, the first few packets that two computers exchange during a secured transaction are likely to be authentication messages. By retransmitting these same packets, still in their encrypted form, attackers can sometimes use them to gain access to secured resources. IPsec prevents packet replays from being effective by assigning a sequence number to each packet. A system using IPsec will not accept a packet that has an incorrect sequence number.
- **Integrity** IPsec uses cryptographic keys to calculate a checksum called a hash message authentication code (HMAC) for the data in each packet, and it then transmits that HMAC with the data. If anyone modifies the packet while it is in transit, the HMAC calculated by the receiving computer will be different from the one in the packet. This prevents attackers from modifying the information in a packet or adding information to it.

A system using IPsec can use AH by itself or in combination with ESP. Using AH alone provides basic security services, with relatively low overhead. However, AH by itself does not prevent unauthorized users from reading the contents of captured data packets. Using AH does, however, guarantee that no one has modified the packets en route, and that the packets did actually originate at the system identified by the packet's source IP address.

On a TCP/IP network, a normal packet has a format as shown in Figure 5-3. A message generated by an application is encapsulated by a transport layer protocol (TCP or UDP), which is in turn encapsulated by IP at the network layer, and by a protocol such as Ethernet at the data-link layer.

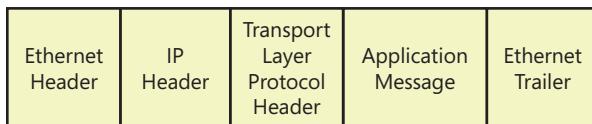


FIGURE 5-3 A typical TCP/IP data packet.

When a computer uses AH to protect its transmissions, the system inserts an AH header into the IP datagram, immediately after the IP header and before the transport layer protocol header, as shown in Figure 5-4.

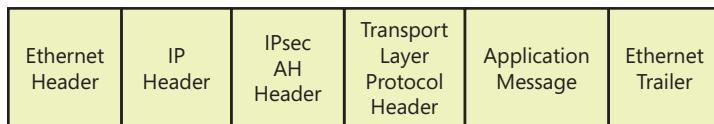


FIGURE 5-4 The AH header location.

The IP ESP protocol encrypts the data in an IP datagram, preventing intruders from reading the information in packets they capture from the network. ESP also provides authentication, integrity, and anti-replay services. Unlike AH, which inserts only a header into the IP datagram, ESP inserts a header and a trailer, which surround the datagram's payload, as shown in Figure 5-5.



FIGURE 5-5 The ESP header and trailer locations.

ESP encrypts all the data following the ESP header—up to and including the ESP trailer. Therefore, someone who captures a packet encrypted with ESP can read the contents of the IP header but cannot read any part of the datagram's payload, including the TCP or UDP header.

An IPsec packet can use ESP by itself or in combination with AH. When a packet uses both protocols, the ESP header follows the AH header, as shown in Figure 5-6. Although AH and ESP perform some of the same functions, using both protocols provides the maximum possible security for a data transmission.

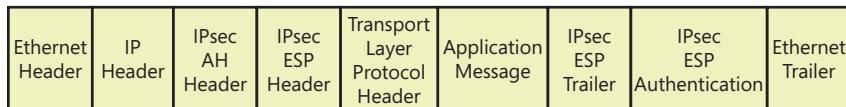


FIGURE 5-6 An IP datagram using AH and ESP.

KEY EXCHANGE PROTOCOLS

As with any security solution that provides encrypted communication between two systems, a preliminary negotiation between IPsec systems is necessary, in which the computers authenticate each other's identities, select appropriate encryption algorithms from those they have in common, and exchange the keys they will use to encrypt subsequent transmissions. A collection of these common settings, which the computers will use to execute the various security services associated with IPsec, is called a security association (SA).

The Internet Security Association and Key Management Protocol (ISAKMP) provides the means for two systems to create and manage the SAs they will need to complete the preliminary negotiation. It's important to understand that ISAKMP does not actually perform the authentication between the systems, nor does it exchange the keys the systems will need; the protocol only provides the framework within which those events take place. ISAKMP operates independently from any other protocols that complete the preliminary negotiation that sets up the IPsec communications between the two computers.

Other protocols that perform the actual key exchanges include Oakley and Simple Password Exponential Key Exchange (SPEKE). Another protocol is the Internet Key Exchange (IKE) protocol, which incorporates part of Oakley and SPEKE, and adds ISAKMP to create a comprehensive, hybrid key exchange solution.

TRANSPORT MODE AND TUNNEL MODE

IPsec can operate in two modes: transport mode and tunnel mode. Transport mode protects communications between computers on a LAN. In transport mode, the two end systems must support IPsec, but intermediate systems, such as routers, do not have to support IPsec. All the discussion of the AH and ESP protocols so far in this chapter applies to transport mode.

Tunnel mode provides security for gateway-to-gateway WAN connections, and particularly, VPN connections, which use the Internet as a communications medium. In tunnel mode, the end systems do not support IPsec; instead, the routers at both ends of the WAN connection use IPsec to secure the data passing over the WAN connection. IPsec, in essence, forms a protected tunnel through an unprotected medium. The internal network traffic between the end systems and the routers uses standard, unprotected TCP/IP communications.

IPsec uses a different packet structure in tunnel mode. In transport mode, IPsec modifies the existing IP datagram by adding its own headers. In tunnel mode, the IPsec implementation creates an entirely new datagram and uses it to encapsulate the existing datagram, as shown in Figure 5-7.

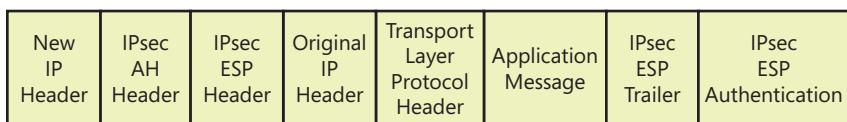


FIGURE 5-7 An IPsec tunnel mode packet.

The “inner” IP header is the header from the original datagram, which remains unchanged. The ESP header and trailer surround the original datagram and are themselves preceded by a new, “outer” IP header. This outer header is designed to get the packet only from one router to the other. Although the source IP address and destination IP address of the inner IP header contain the ultimate source and destination of the packet, the outer header contains the IP addresses of the two gateways that form the endpoints of the tunnel.

Tunnel mode communications proceed as follows:

1. Computers on one of the private networks transmit their data by using standard, unprotected IP datagrams.
2. The packets reach the router that provides access to the WAN, and then the router encapsulates them by using IPsec, encrypting and hashing data as needed.
3. The router transmits the protected packets through the secure tunnel to a second router at the other end of the WAN connection.
4. The second router verifies the packets by calculating and comparing integrity check values, and decrypts the packets if necessary.
5. The second router repackages the information in the packets into standard, unprotected IP datagrams, and transmits them to their destinations on the private network.

EXAM TIP Although candidates for the Network+ exam are not required to know the packet formats for the IPsec protocols, they should be familiar with the names of the component protocols, the concept of tunneling, and the distinction between IPsec's transport and tunnel modes.

SSL and TLS

SSL is a special-purpose security protocol that protects the data transmitted by servers and their clients. Unlike IPsec, SSL operates at the application layer and can protect only the data generated by the specific applications for which it is implemented. Originally designed by Netscape Communications to protect the HTTP data exchanged by web servers and browsers, SSL began to be used by other applications to protect their traffic as well. For example, email applications using the Simple Mail Transfer Protocol (SMTP) and articles exchanged by newsreaders and servers by using the Network News Transfer Protocol (NNTP) made use of SSL.

SSL protects application layer data in the following three ways:

- **Authentication** Clients and servers can exchange credentials to confirm their identities.
- **Encryption** Data exchanged by clients and servers is encrypted by using public key encryption to prevent the data in intercepted packets from being compromised.
- **Data integrity** Packets are signed with HMAC, which the receiver uses to ensure that the data has not been modified in transit.

Netscape revised the SSL standard several times, culminating in SSL 3.0, which the IETF published as a historic document (RFC 6101) in 2011.

TLS is the successor to SSL, and is now the standard cryptographic protocol for web communications. Virtually all current web servers and browsers support TLS, as do many other Internet applications. When you access a secured website on the Internet, your browser points to a Uniform Resource Locator (URL) with an <https://>

prefix instead of the usual *http://*. The initial exchange between the client and the server then proceeds as follows:

1. The *https://* prefix causes the browser to send its HTTP request to TCP port 443 instead of the standard HTTP port 80. The request also includes a list of the cryptographic standards that the client supports.
2. The server responds to the client's HTTP request by selecting the strongest of the offered cryptographic standards the two have in common and by sending the client its digital certificate and public key.
3. The client contacts the certification authority that issued the server's certificate and confirms the server's authenticity.
4. The client, using the agreed-upon cryptographic cipher, encrypts a randomly selected session number by using the server's public key and transmits it. This is called the session key.
5. The server decrypts the client's session key by using its private key, and then uses that number to generate keys for the encryption and decryption of all subsequent data transmissions in the session.

In a similar manner, other applications, such as Voice over IP (VoIP) and virtual private networks use TLS for their cryptography.

Unlike the SSL standard, which the IETF only published after its retirement, TLS has been designed and developed by an IETF working group. The current standard is RFC 5246, "The Transport Layer Security (TLS) Protocol Version 1.2."

True or false: Tunneling is essentially a violation of the OSI model encapsulation rules.

Answer: *True*. In a PPTP packet, you have a data-link layer protocol frame encapsulated within a network layer IP datagram (instead of the other way around), and the datagram encapsulated again within another data-link layer protocol.

EXAM TIP For the purposes of the Network+ exam, SSL and TLS are essentially the same protocol, with TLS 1.0 being the next upgrade after SSL 3.0. In some instances, TLS versions are notated with a corresponding SSL version number, such as TLS 1.0 (SSL 3.1). Later version releases are TLS 1.1 (SSL 3.2) and TLS 1.2 (SSL 3.3). TLS 1.2 is the latest release.

True or false: Many web applications refer to SSL encryption, when they are actually using TLS.

Answer: *True*. SSL has become a generic expression for web encryption, despite the fact that in most cases, today's web browsers and servers use TLS instead.

True or false: The IPsec Authentication Header protocol provides additional security to network links without encrypting the data in the packets.

Answer: *True*. The AH protocol provides message integrity and other services, but it does not encrypt the data in the packets.

Explain the use of remote access

Because it enables outside users to access a network, security must be an integral part of any remote access solution. The following sections examine some of the ways in which networks provide secured remote access.

Remote access server (RAS)

One of the major uses for WAN connections is to enable an individual user to access resources on a remote LAN. Many users today spend part of their time telecommuting, or they require access to email and other network resources while traveling or working at home during off hours.

To establish a remote network connection, the computers involved must have the following elements:

- **Physical layer connection** The computers must be connected by using a WAN technology as a physical medium. In nearly all cases, a remote access solution uses either a direct link to a computer on the target network or an Internet access connection.
- **Common protocols** The two computers to be connected must use the same protocols at the data-link layer and above. This means that you must configure both computers to use a data-link layer protocol suitable for point-to-point connections. The computers must also use the same network and transport layer protocols, such as TCP/IP.
- **Network configuration** To communicate with a remote network by using TCP/IP, for example, a computer must be assigned an IP address and other configuration parameters appropriate for that network. The remote user can configure the TCP/IP settings if someone familiar with the host network supplies them, but most remote networking solutions enable the network server to assign configuration parameters automatically by using the Dynamic Host Configuration Protocol (DHCP) or some other mechanism.
- **Host and remote software** Each of the computers to be connected must be running an application appropriate to its role. The remote (or client) computer needs a client program that can use the physical layer medium to establish a connection, such as by instructing a modem to dial a number. The host (or server) computer—sometimes called a remote access server (RAS)—must have a program that can respond to a connection request from the remote computer and provide access to the network.

PPP

The data-link layer protocol used for most WAN connections is the Point-to-Point Protocol (PPP). PPP is much simpler than Ethernet, the other main data-link layer protocol, because it is designed for use by WAN connections that consist of only two systems. Because there are only two devices involved, there is no need for the protocol to support complex procedures, such as node addressing or media access control. However, PPP does include support for a variety of ancillary protocols that provide authentication and other services.

PPP is designed for use with modems and other direct connections in which there is no need for media access control. Because it connects only two systems, PPP is called a point-to-point or end-to-end protocol. On a system that uses PPP, the TCP/IP protocols define the workings of the entire protocol stack, except for the physical layer.

In most cases, systems use PPP to provide Internet or WAN connectivity, whether or not the system is connected to a LAN. Virtually every standalone PC that ever used a modem to connect to an ISP for Internet access, did so by using a PPP connection. LANs also use PPP connections in their routers to connect to an ISP to provide Internet access to the entire network or to connect to another LAN, forming a WAN connection. Although PPP is commonly associated with modem connections, other physical-layer technologies can also use PPP, including leased lines and various forms of broadband connections.

PPP is a connection-oriented protocol that provides a data link between two systems in the simplest sense of the term. It encapsulates IP datagrams for transport between computers, just as Ethernet does, but the frame it uses is far simpler. This is because the protocol is not subject to the same problems as the LAN protocols.

PPP was created to provide the capability to multiplex different network-layer protocols and support various authentication protocols. Naturally, the cost of these additional features is a larger header, but PPP still only adds a maximum of 8 bytes to a packet (as compared to the 16 bytes needed for an Ethernet frame). Many connections to ISPs, whether by standalone systems or routers, use PPP in some form, because it enables the ISP to implement access control measures that protect their networks from intrusion by unauthorized users.

A typical PPP session consists of several connection establishment and termination procedures, using other protocols in addition to PPP itself. These procedures can include the following:

- **Connection establishment** The system initiating the connection uses the Link Control Protocol (LCP) to negotiate communication parameters that the two machines have in common.
- **Authentication** Although this is not required, the system might use an authentication protocol such as PAP (the Password Authentication Protocol) or CHAP (the Challenge Handshake Authentication Protocol) to negotiate access to the other system.
- **Network-layer protocol connection establishment** For each network-layer protocol that the systems use during the session, they perform a separate connection establishment procedure by using a Network Control Protocol (NCP) such as IPCP (the Internet Protocol Control Protocol).

When the physical-layer connection between the two systems has been established (through a modem handshake or other procedure), the PPP connection establishment process begins. The two systems pass through several distinct phases during the course of the session, as illustrated in Figure 5-8 and discussed in the following sections.

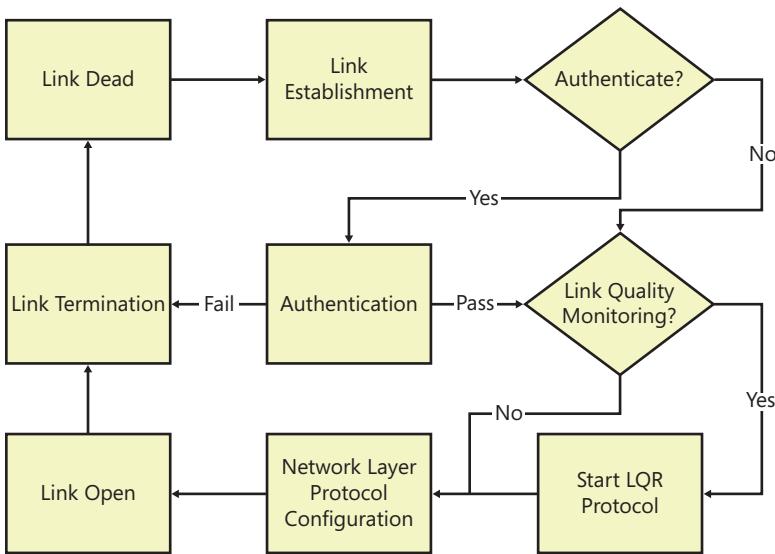


FIGURE 5-8 PPP connection phases.

THE LINK DEAD PHASE

Both systems begin and end the session in the Link Dead phase, which indicates that no physical-layer connection exists between the two machines. In a typical session, an application or service on one system initiates the physical-layer connection by dialing the modem or using some other means. After the hardware connection process is completed, the systems pass into the Link Establishment phase.

THE LINK ESTABLISHMENT PHASE

In the Link Establishment phase, the system initiating the connection transmits an LCP Configure Request message to the destination containing the options it would like to enable, such as the use of specific authentication, link quality monitoring, and network layer protocols (if any). If the receiving system can support all the specified options, it replies with a Configure Ack message containing the same option values, and this phase of the connection process is completed.

THE AUTHENTICATION PHASE

The Authentication phase of the connection process is optional and is triggered by the inclusion of the Authentication Protocol option in the LCP Configure Request message. During the LCP Link Establishment process, the two systems agree on an authentication protocol to use. Use of the PAP and CHAP protocols is common, but other protocols are available. The message format and exchange procedures for the Authentication phase are dictated by the selected protocol.

A successful transaction causes the connection procedure to proceed to the next phase, but the effect of a failure is dictated by the implementation of the protocol.

Some systems proceed directly to the Link Termination phase in the event of an authentication failure, while others might permit retries or limited network access to a help subsystem.

LINK QUALITY MONITORING

The use of a link quality monitoring protocol is also an optional element of the connection process and is triggered by the inclusion of the Quality Protocol option in the LCP Configure Request message. Although the option enables the sending system to specify any protocol for this purpose, only one has been standardized, the Link Quality Report protocol. The negotiation process that occurs at this phase enables the systems to agree on an interval at which they should transmit messages containing link traffic and error statistics throughout the session.

NETWORK-LAYER PROTOCOL CONFIGURATION

PPP supports the multiplexing of network layer protocols over a single connection, and during this phase, the systems perform a separate network-layer connection establishment procedure for each of the network layer protocols that they have agreed to use during the Link Establishment phase. Each network-layer protocol has its own Network Control Protocol (NCP) for this purpose, such as the Internet Protocol Control Protocol (IPCP).

THE LINK OPEN PHASE

When the individual NCP exchanges are completed, the connection is fully established and the systems enter the Link Open phase. Network-layer protocol data can now travel over the link in either direction.

THE LINK TERMINATION PHASE

When one of the systems ends the session or when the session is ended as a result of other conditions such as a physical-layer disconnection, an authentication failure, or an inactivity timeout, the systems enter the Link Termination phase. To sever the link, one system transmits an LCP Terminate Request message, to which the other system replies with a Terminate Ack. Both systems then return to the Link Dead phase.

PPPoE

Point-to-Point Protocol over Ethernet (PPPoE) is a TCP/IP standard, and it is defined in RFC 2516, "A Method for Transmitting PPP Over Ethernet (PPPoE)." PPPoE provides a way to create individual PPP connections between computers on an Ethernet LAN and external services connected to the LAN via a broadband device, such as a cable or DSL modem.

Broadband remote network access devices can easily support multiple computers, and Ethernet is the most common protocol used to join the computers to a network and connect them to the broadband device. However, a shared Ethernet LAN does not enable each computer to access remote services that use individual parameters

for functions such as access control and billing. The object of PPPoE is to connect multiple computers to a remote network by using an Ethernet LAN and broadband technology, while establishing a separate PPP connection between each computer and a specified remote service. Each PPP connection has all of the PPP components, such as LCP negotiation, authentication, and network control protocol configuration.

Remote Terminal Emulation

Although it can use the same WAN technologies as other solutions, remote terminal emulation handles the remote access problem in a different way. VPN solutions enable the client computer to function as a member of the remote network. Any application you launch on the client computer uses the client's processor and memory to run. With terminal emulation, you are taking over a computer on the network and operating it from a remote location. If you execute a program from the client terminal, it actually runs by using the host system's resources.

Terminal emulation has been around since the early days of computing. A program like Telnet enables you to connect to a remote computer and execute commands there. However, graphical user interfaces have complicated the terminal emulation process.

The first graphical terminal emulation program was called WinFrame, designed and marketed by the Citrix Corporation. Citrix created a protocol called Independent Computing Architecture (ICA) that defines how a client and server should exchange graphical terminal emulation traffic. ICA essentially defines how the client sends keyboard and mouse input to the server, and how the server sends a graphical screen display to the client.

Microsoft also pursued graphical terminal emulation technology, first by licensing it from Citrix, and later by creating their own, which they called Windows Terminal Services. As of the Windows Server 2008 R2 release, the name was changed to Remote Desktop Services. The client program, supplied with all Windows versions, is called Remote Desktop Connection (as shown in Figure 5-9) and the protocol that the two computers use to communicate is the Remote Desktop Protocol (RDP).

Remote terminal connections are based on IP, so they can theoretically use any type of WAN connection, as long as it provides access to the computer the client will control.

Secure Shell

Secure Shell (SSH) is a client/server program that provides character-based remote access with full encryption. SSH was designed as a replacement for Telnet and similar UNIX/Linux programs that provide similar services, but which transmit data, including passwords, in clear text. SSH uses public key encryption to authenticate users who, once connected, can execute commands at the remote command prompt. There are also associated protocols that provide other capabilities, such as SSH File Transfer (SFTP).

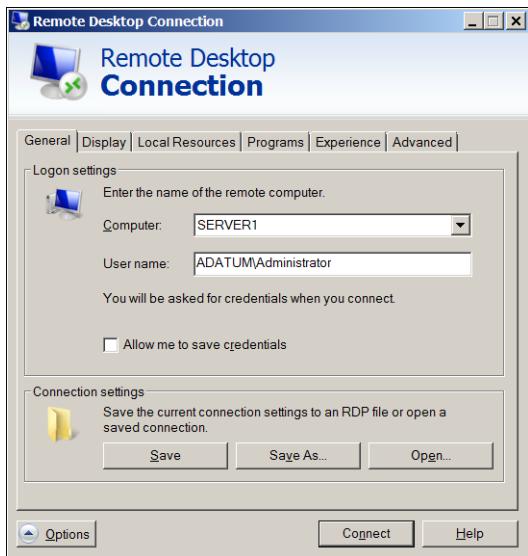


FIGURE 5-9 The Remote Desktop Connection client.

True or false: RAS connections enable users to execute applications on a remote network computer.

Answer: *False*. A RAS connection provides users with access to network resources, but any applications they execute run on their local systems.

EXAM TIP The Network+ objectives refer to PPP primarily as a network access security mechanism. Candidates should be particularly aware of how PPP integrates various authentication protocols into its connection establishment sequence.

True or false: PPP frames always include the MAC address of the destination computer.

Answer: *False*. Because PPP is used only for connections between two systems, there is no need to include addresses in the frames.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What protocol does PPP use to negotiate the communication parameters that the two connecting machines have in common?
2. Which protocol has replaced SSL as the dominant means of encrypting web client/server traffic at the application layer?
3. Which IPsec protocol inserts a header, providing mutual authentication and data integrity, but does not provide data encryption?

4. What protocol enables the computers on a home network to establish individual connections to remote services accessible through a broadband router?
5. SSH provides a secure alternative to what traditional remote access protocol?

Objective 5.3: Explain methods of user authentication

Authentication is the process of verifying a user's identity, for the purpose of distinguishing legitimate users from uninvited guests. As such, it is one of the most prominent and one of the most visible concepts in security. From driver's licenses to user names and passwords, authentication is a regular part of everyone's daily life. Without authentication, administrators could not control access to network resources.

If a network's authentication strategy is not strong enough, viruses, worms, and malicious attackers can gain access to it. Password guessing, password cracking, and man-in-the-middle attacks (in which a third party eavesdrops on a protected communication to gain access to passwords or encryption keys) are all attempts to exploit weaknesses in an authentication strategy. However, an authentication strategy can also be too restrictive, keeping attackers out, but also preventing legitimate users from doing their jobs. This objectives calls for Network+ exam candidates to be familiar with some of the most common authentication protocols and techniques currently used on data networks.

Exam need to know

- Explain PKI authentication
For example: How many encryption keys are required for PKI?
- Explain Kerberos user authentication
For example: How does Kerberos authenticate users without transmitting passwords over the network?
- Explain AAA authentication
For example: How does an AAA server provide services to other servers?
- Explain network access control
For example: How does a network use IEEE 802.1X to perform a posture assessment?
- Explain CHAP authentication
For example: Why is CHAP rarely used for authentication today?
- Explain MS-CHAP authentication
For example: What are the differences between MS-CHAP v1 and MS-CHAP v2?
- Explain EAP authentication
For example: Why is EAP able to support multifactor authentication when the other authentication protocols cannot?
- Explain two-factor authentication
For example: Why do smartcards require a password?

- Explain multifactor authentication
For example: Under what conditions is multifactor authentication necessary?
- Explain single sign-on authentication
For example: How is single sign-on an advantage to the network administrator?

Explain PKI authentication

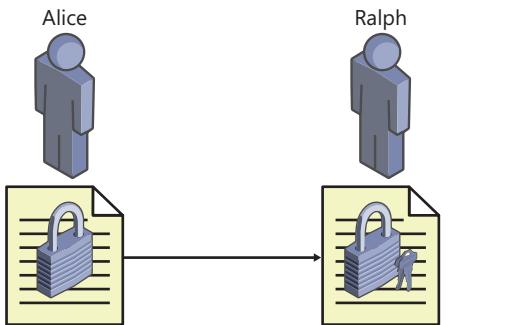
At the most basic level, encryption is a system in which new characters are substituted for those in the original content. For example, if you create a key specifying that the letter A should be replaced by O, the letter B by F, the letter C by S, and so forth, any message you encode using that key can be decoded by anyone else possessing that key. This is called secret key encryption, because you must protect the key from being compromised.

For computer transactions, this type of encryption is useless, because there is no practical and secure way to distribute the secret key to all of the parties involved. After all, if the object is to send an encrypted message to a recipient over the network, it would hardly be appropriate to first send the secret encryption key in an unsecured message.

For encryption on a data network to be both possible and practical, computers typically use a form of public key encryption. In a public key infrastructure (PKI), every user has two keys: a public key and a private key. As the names imply, the public key is freely available to anyone, whereas the private key is carefully secured and never transmitted over the network. The way the system works is that data encrypted with the public key can only be decrypted with the private key, and conversely, data encrypted with the private key can only be decrypted by using the public key. Data encrypted with a public key cannot be decrypted with that public key, nor can data encrypted with a private key be decrypted by using that private key. It is the protection of the private key that guarantees the security of messages encrypted in this system.

If Alice wants to send Ralph a message, making sure that no one but Ralph can read it, then Alice must obtain Ralph's public key and use it to encrypt the message, as shown in Figure 5-10. Alice can then transmit the message to Ralph over the network, secure in the knowledge that only Ralph possesses the private key needed to decrypt it. Even if an intruder were to intercept the message during transmission, it would still be in its encrypted form, and therefore impenetrable. When Ralph receives the message and decrypts it using his private key, he can reply to it by using Alice's own public key to encrypt his response, which only Alice can decrypt, by using her private key.

PKI also works in reverse, to provide a service called nonrepudiation. Nonrepudiation is essentially a method for ensuring that the communication you receive from a specific person actually originates from that person. The sender encrypts the data by using his or her private key, and the receiver decrypts it by using the sender's public key, as shown in Figure 5-11. Because no one but the sender has access to the private key, the ability to decrypt it with the public key proves that the sender actually encrypted it. This is sometimes referred to as a digital signature.

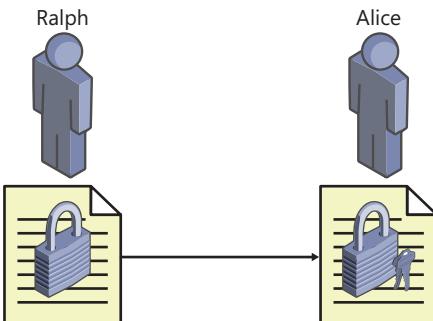


Document encrypted by Alice
with Ralph's public key...

...can only be decrypted by
Ralph with his private key

Therefore, no one can read the document except Ralph

FIGURE 5-10 Document encryption using PKI.



Document encrypted by Ralph
with his private key...

...can be decrypted by anyone
using Ralph's public key

Therefore, the document must have come from Ralph

FIGURE 5-11 Document nonrepudiation using PKI.

True or false: Single key encryption is inefficient for data networks.

Answer: *True*. Single key encryption uses one key to encrypt and decrypt data. This is not a practical solution for data networks because there is no secure method for distributing the keys.

True or false: In a secured PKI transaction, the sender encrypts the data using the recipient's public key and the recipient replies using his or her private key.

Answer: *False*. In a secured PKI transaction, all messages must be encrypted using public keys and decrypted using private keys.

Explain Kerberos authentication

The Kerberos protocol is an authentication protocol that uses tickets to coordinate the authentication of network clients and servers. Named for the three-headed dog guarding the entrance to Hades in Greek mythology, a Kerberos implementation consists of the following three components:

- **Clients** Users or applications that must be authenticated before they can access network resources.
- **Servers** Systems hosting resources that clients need to access.
- **Key Distribution Center (KDC)** An authentication server that functions as an intermediary between clients and servers by issuing tickets to clients, which they can use to access server resources.

The Kerberos protocol was developed at the Massachusetts Institute of Technology in the early 1980s, and it is now standardized by the IETF in RFC 1510, "The Kerberos Network Authentication Service (V5)." The Kerberos protocol has been implemented in various UNIX and Linux distributions for years, and it is the default authentication protocol used by the Active Directory Domain Services (AD DS) directory service included in all Windows Server versions from Windows 2000 through Windows 7.

When a client logs on to a network that uses the Kerberos protocol, it sends a request message to the authentication service on a KDC, which already possesses the account name and password associated with that client. The KDC responds by sending a ticket-granting ticket (TGT) to the client, as shown in Figure 5-12. The TGT is encrypted by using a key based on the client's password. When the client receives the TGT, it prompts the user for the password and uses it to decrypt the TGT. Because only that user (presumably) has the password, this process serves as an authentication.

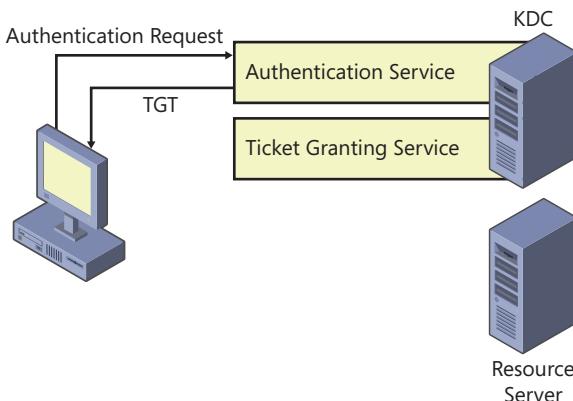


FIGURE 5-12 The Kerberos authentication service exchange.

Now that the client possesses the TGT, it can access network resources by sending a request to a Ticket-Granting Service (TGS), which might or might not be running on the same KDC. The request contains an encrypted copy of the TGT. The TGS, after

it has decrypted the TGT and verified the user's status, creates a server ticket and transmits it to the client, as shown in Figure 5-13.

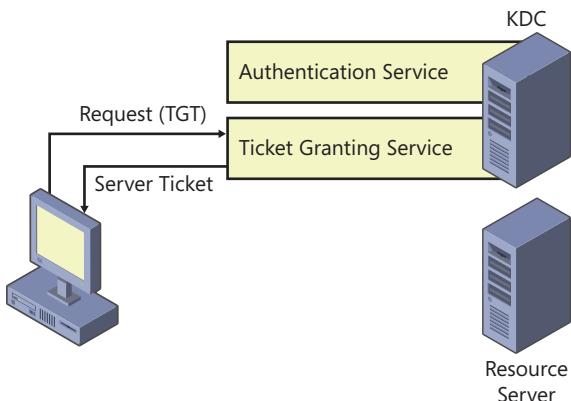


FIGURE 5-13 The Kerberos TGS exchange.

The server ticket allows a specific client to access a specific server for a limited length of time. The ticket also includes a session key, which the client and the server can use to encrypt the data transmitted between them, if necessary. The client transmits the server ticket (which was encrypted by the TGS by using a key that the server already possesses) to that server. After the server decrypts the server ticket, the server grants the client access to the requested resource, as shown in Figure 5-14.

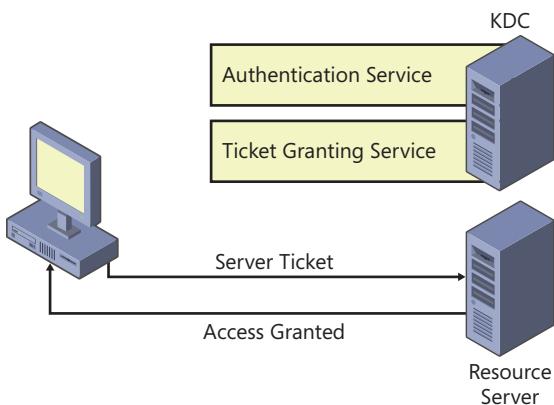


FIGURE 5-14 The Kerberos server ticket exchange.

The client is now authenticated to the network and authorized to access the requested resource. The next exchange between the client and the server will consist of messages specific to the application that requires access to the network resource.

True or false: Kerberos requires users to have continuous access to the KDC.

Answer: *True*. After the initial authentication, the Kerberos KDC is involved each time a user is authorized to access a particular network resource.

EXAM TIP The Network+ exam requires candidates to know that Kerberos is the authentication protocol for AD DS, and that it works without transmitting passwords over the network. However, candidates do not have to memorize the details of the Kerberos authentication transaction.

True or false: Today, the primary application for Kerberos on networks is Active Directory Domain Networks.

Answer: *True*. Most editions of Microsoft Windows Server 2008 R2 include AD DS. Using AD DS, you can promote a server to a domain controller, which functions as a Kerberos KDC. AD DS users then authenticate to the domain controller, which also grants them access to specific network resources.

Explain AAA authentication

Directory services such as AD DS are designed to support LANs on which every user has a permanent account created by an administrator and continuous access to the servers providing authentication and authorization services. There are other situations, however, in which these stipulations do not apply. Remote access environments, for example, have users connecting to a server by using some form of WAN or VPN link. These users require authentication and authorization also, but they present different challenges.

In the past, these users connected to the servers by using PPP, which offered a framework providing several authentication options. However, authenticating all of these users was a problem in cases where a constantly changing list of user names and passwords had to be maintained on dozens of servers. This led to the creation of a different kind of centralized service called Authentication, Authorization, and Accounting, or AAA.

MORE INFO For more information on PPP, see “Objective 5.2: Explain the methods of network access security.”

An AAA server is a computer with the account information needed to grant or deny any user access to the network, as well as the information needed to authorize access to specific resources. Finally, the server includes the capability to maintain accounting information about its activities, such as how often a particular user logs on to the network.

RADIUS

The most common AAA server implementation is called Remote Authentication Dial In User Service (RADIUS). Devised in the early 1990s as a tool for ISPs and other remote access providers, RADIUS is a standard that defines a client/server protocol running on the application layer and using UDP for transport services on ports 1812 and 1813, for authentication and accounting, respectively. Some implementations

use ports 1645 and 1646, which are the unofficial ports the service used prior to standardization.

NOTE RADIUS was originally designed by a private company, and only later standardized by the IETF in RFC 2865, “Remote Authentication Dial In User Service (RADIUS),” and RFC 2866, “RADIUS Accounting.”

The RADIUS server provides authentication, authorization, and accounting services to the remote access servers that receive incoming network access requests. The remote access servers contain a RADIUS client, as shown in Figure 5-15, which, after prompting the user for credentials, generates a RADIUS Access-Request message and sends it to the server. The Access-Request message contains one or more attributes, which might specify the user’s account name and password, contain a digital certificate, and/or provide other information to help identify the user.

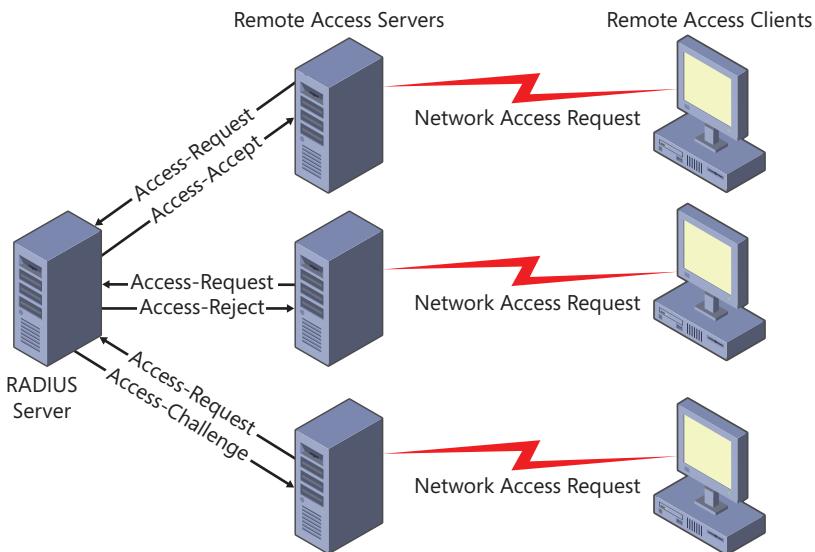


FIGURE 5-15 A RADIUS server and its clients.

The RADIUS server has access to the information needed to verify the user’s identity, which it does by using any one of several remote access authentication protocols. Some RADIUS servers maintain the user’s credentials in an internal flat-file database, whereas others access an external information store, such as a SQL or AD DS database. After processing the user’s credentials, the RADIUS server generates a reply and sends it back to the client on the remote access server. The reply message can take one of the following three forms:

- **Access-Accept** Grants the user access to the network, as requested.
- **Access-Challenge** Requests additional credentials from the user to complete the authentication, such as a smart card, PIN, biometric scan, or secondary password.
- **Access-Reject** Denies the user’s requested access to the network.

As with other centralized authentication systems, RADIUS simplifies the task of providing updated account information to a large fleet of network access servers. The Windows Server operating systems include a RADIUS implementation, which in Windows Server 2008 R2 is part of the Network Policy and Access Services role. In earlier versions, the software was known as Internet Authentication Server (IAS). There are other implementations for UNIX/Linux, including open source options such as FreeRADIUS.

TACACS+

Terminal Access Controller Access-Control System (TACACS) is another centralized logon solution that enables users that successfully authenticate to one system to access other systems as well. The primary differences between TACACS and RADIUS are that TACACS separates the authentication, authorization, and accounting services into separate processes and uses TCP port 49 for its transport-layer communications, rather than UDP.

Cisco Systems developed the original TACACS into successive standards called extended TACACS (XTACACS) and later TACACS+, which they used to control administrative access to their router and switch products.

True or false: RADIUS clients are remote access servers, not remote users.

Answer: *True*. The RADIUS client/server relationship is between remote access servers requesting services and RADIUS servers granting them. The remote access users have no direct contact with the RADIUS server.

EXAM TIP Candidates for the Network+ exam should be familiar with how RADIUS provides centralized authentication, authorization, and accounting services for remote access installations.

True or false: ISPs often use RADIUS for user authentication.

Answer: *True*. ISPs use RADIUS to centralize authentication, eliminating the need to replicate a large customer database to many different remote access servers.

True or false: RADIUS uses TCP for its transport services, while TACACS+ uses UDP.

Answer: False. RADIUS uses UDP and TACACS+ uses TCP.

Explain network access control

Network access control (NAC) is a means of implementing a set of policies that specify whether a client is to be granted access to a network. In most cases, these policies include a set of software requirements. For example, clients might be required to have anti-virus software installed and be running the latest operating system updates. When a client attempts to connect to the network, the server performs a posture assessment of the client. If the client does not meet the standards specified by the NAC policies, a remediation might take place, in which the client can install the required software.

NAC products implement their policies in various ways, but one of the most common uses the IEEE 802.1X protocol. The IEEE 802.1X standard defines a centralized authentication mechanism for port-based network access control. Although designed for use on wired or wireless LANs, it is primarily wireless networks that use it.

The 802.1X standard defines three roles, as follows:

- **Suplicant** The client seeking access to the network.
- **Authenticator** A switch or access point that will provide the client with access to the network.
- **Authentication server** A RADIUS server that will perform the actual client authentications.

These three entities communicate using encapsulated Extensible Authentication Protocol (EAP) messages. The supplicant sends its authentication credentials to the authenticator, which forwards them to the authentication server, which checks them against an appropriate user account store. These messages can also include information about the supplicant that enables the authentication server to perform the posture assessment.

As part of the 802.1X authentication process, EAP generates a unique encryption key for each client. The use of RADIUS forces the client to regularly generate a new encryption key, which makes it far more difficult for attackers to gather enough traffic and enough time to penetrate a key.

The 801.X standard divides the port in the switch or access point into two logical ports: an uncontrolled port, which is used to exchange the EAP messages, and a controlled port, which is opened to other traffic once the authentication and authorization processes are completed.

True or false: The IEEE 802.1X standard defined a means by which wireless networks can generate new shared secrets.

Answer: *True*. In addition to a simple yes or no response to an authentication request, the authentication server can also provide other applicable connection parameters for the supplicant, including a dynamic shared secret. The authentication server also relays the positive access decision to the applicant, transmitting the dynamic shared secret to it as well. The supplicant and the authenticator, that is, the client and the access point, now share common key material that they can use to encrypt and decrypt the traffic that they will exchange. Each time the shared secret changes, potential attackers must restart the process of cracking the encryption key.

EXAM TIP Candidates for the Network+ exam should be aware that the use of IEEE 802.1X is synonymous with WPA2-Enterprise.

True or false: The use of IEEE 802.1X can make WEP suitably secure for use on a wireless LAN

Answer: *True*. IEEE 802.1X can increase the security of a wireless network using WEP by authenticating users connecting to a wireless LAN, authorizing user access to the network, and dynamically changing WEP encryption keys.

Explain CHAP authentication

CHAP is a challenge-response authentication protocol that uses the industry-standard Message Digest 5 (MD5) hashing scheme to encrypt the response—that is, the user's password. This means that anyone capturing the authentication packets will not be able to read the passwords in them. CHAP, however, provides relatively weak protection compared to some of the other authentication protocols. CHAP does not support encryption of the connection data, and the passwords it uses must be stored in a reversibly encrypted format. This means that if users are establishing connections with their standard operating system user accounts, the network administrator must weaken the encryption of those passwords to accommodate CHAP.

True or false: CHAP authentication is disabled by default in Windows.

Answer: *True.* CHAP is recommended over the PAP, which transmits passwords in clear text, but there are much more capable authentication protocols available. Most networks that use CHAP do so because they have users running operating systems other than Windows that do not support anything else.

Explain MS-CHAP authentication

MS-CHAP is a variation of the CHAP authentication protocol that was designed by Microsoft, specifically for its Windows operating systems. There are two versions of MS-CHAP, as follows:

- **MS-CHAP v1** Version 1 of MS-CHAP is a one-way authentication protocol, meaning that the client computer is authenticated to the server, but the server is not authenticated to the client. MS-CHAP v1 provides both authentication encryption and data encryption, but the encryption is relatively weak because the protocol uses the account password to create the encryption key. The key remains the same as long as the user retains the same password. This provides a potential intruder with more time to crack the encryption, weakening the cryptography. At this time, only older versions of Windows, such as Windows 95 and Windows 98, use MS-CHAP v1.
- **MS-CHAP v2** Version 2 of MS-CHAP improves on version 1 by adding two-way authentication and increasing the strength of the encryption. Version 2 uses a new encryption key for each connection and for each direction in which data is transmitted. This makes the encryption far more difficult to crack. MS-CHAP v2 is the preferred authentication method for systems that do not use smart cards or public key certificates for authentication. By default, Windows VPN remote access clients use MS-CHAP v2 to authenticate.

True or false: MS-CHAP v2 is the strongest password-based authentication protocol supported by the Microsoft Windows remote access client and server.

Answer: *True.* MS-CHAP v2 provides stronger password protection than MS-CHAP v1, CHAP, or PAP. EAP can provide more protection, but only when using authentication mechanisms other than passwords, such as smart cards.

True or false: MS-CHAP v1 and MS-CHAP v2 are the only Windows remote access authentication protocols that support password change during the authentication process.

Answer: *True*. If an organization has users who always work remotely, the inability to change a password during authentication can be a real problem because they cannot simply change the password the next time they are in the office. If you use an authentication method other than MS-CHAP, users will have to connect by using a mechanism other than a VPN to change their passwords.

Explain EAP authentication

EAP is a shell protocol that enables systems to use various types of authentication mechanisms. The primary advantage of EAP is that it enables a computer to use mechanisms other than passwords for authentication, including public key certificates, smart cards, and biometric devices such as fingerprint scanners. Some of the EAP variations supported by Windows are as follows:

- **Protected EAP (PEAP)** The primary application for PEAP is the authentication of wireless users with an account name and password.
- **Message Digest 5 Challenge (MD5-Challenge)** MD5-Challenge uses the same challenge handshake protocol as PPP-based CHAP, but the challenges and responses are sent as EAP messages. A typical use for MD5-Challenge is to authenticate remote access clients other than those created by Microsoft, such as those running Mac OSX. You can also use MD5-Challenge to test EAP interoperability. EAP with MD5-Challenge does not support encryption of connection data.
- **Smart Card or Other Certificate** This authentication method, also known as EAP-Transport Layer Security (EAP-TLS), enables clients to authenticate by using a smart card or a public key certificate.

It is also possible for third-party applications to add other authentication methods to EAP.

True or false: EAP is the strongest authentication protocol supported by Windows.

Answer: *False*. EAP has the potential to provide the strongest authentication, but it is an extensible protocol that is dependent on other protocols for its capabilities.

EXAM TIP Candidates for the Network+ exam should be familiar with the various authentication protocols supported by Windows remote access and when they are used.

True or false: EAP authentication requires users to have a smart card.

Answer: *False*. EAP is an extensible protocol that can support various authentication methods, including passwords, digital certificates, and biometrics.

Explain multifactor authentication

Passwords can be guessed or otherwise compromised, and smart cards can be stolen. A determined intruder can even evade a biometric scan, with sufficient effort. In some cases, one form of authentication alone might not meet an organization's security requirements. Multifactor authentication combines two or more of these authentication methods and reduces the likelihood that an intruder will be able to successfully impersonate a user during the authentication process.

True or false: A system requiring users to supply two different passwords is an example of multifactor authentication.

Answer: *False*. Multifactor authentication requires the use of two or more different authentication mechanisms, such as a smart card and a password.

Explain two-factor authentication

The most common example of multifactor authentication is the combination of a smart card with a password. This is a form of two-factor authentication. Typically, smart cards require the user to specify a password to retrieve the key stored on the card. Before you can authenticate to such a system, you must provide a password (something you know) and a smart card (something you have).

True or false: Requiring a password for smart card authentication lessens the chance of the card being stolen and misused.

Answer: *True*. Any authentication mechanism requiring something the user has invites the risk of loss or theft. Requiring a second authentication factor, in addition to the possession, reduces the potential consequences of that loss or theft.

Explain single sign-on authentication

Centralized authentication models provide much simpler management for larger networks, which lowers help desk costs related to account management. In a centralized model, network resources rely on a single system to authenticate users. When users attempting to access a particular network resource supply their credentials in the normal manner, the server hosting the resource relays those credentials to a separate authentication server, which either grants or denies the users access.

Centralized authentication is required to create an environment in which users can access all of their allotted network resources with a single set of credentials, a convenient arrangement known as single sign-on. The main drawback of the centralized model is that systems must transmit user credentials over the network, which presents an additional security hazard. Most centralized authentication systems use elaborate forms of encryption to protect this data as it travels over the network.

True or false: Single sign-on is a definite advantage for users, but not for network administrators.

Answer: False. Single sign-on is an advantage for administrators as well as users, because the centralized authentication mechanism eliminates the need to maintain user account databases on multiple servers.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What authentication protocol do Windows networks use for AD DS authentication of internal clients?
2. In a PKI system, which key is needed to decrypt data encrypted with a user's public key?
3. Which of the Windows remote access authentication protocols must you use to authenticate users with smart cards?
4. Which of the roles specified by the IEEE 802.1X standard is typically filled by a RADIUS server?

Objective 5.4: Explain common threats, vulnerabilities, and mitigation techniques

To adequately protect a network, administrators must be familiar with the technologies that outsiders use to attack it. Most of the attacks leveled at networks are variations on a few basic types of threat, and familiarizing yourself with these threats enables you to erect better defenses against them. This objective describes some of the most common threats and the most popular techniques for defeating them.

Exam need to know

- Explain common threats to wireless networks
For example: How are the following threats to a wireless network: war chalking, WEP cracking, WPA cracking, evil twin, and rogue access point?
- Explain common attacks
For example: How familiar are you with the following attacks: DoS, DDoS, man in the middle, social engineering, virus, worms, buffer overflow, packet sniffing, FTP bounce, and smurf?
- Explain common mitigation techniques
For example: What types of attacks can you prevent with training and awareness, patch management, policies and procedures, and incident response?

Explain common threats to wireless networks

Wireless networks have their own specialized security protocols, due to the specialized nature of the threats against them. As with other threats, the continued development of greater security technologies drives certain people to constantly search for new weaknesses they can exploit.

Some of the most common threats against wireless networks are as follows:

- **War driving** Today, many wireless LANs protect against intruders by using some form of encryption, but that was not always the case. In the early days of wireless LANs, many people left their networks unprotected, making it possible for unauthorized users to connect to them and access their files or use their Internet connection. War driving is the process of cruising around a neighborhood with a portable computer, looking for unprotected wireless networks available for connection.
- **War chalking** Associated with war driving, war chalking is a practice in which the people discovering an unprotected network leave a mark on a wall or gatepost indicating its presence, so that future war drivers can find it.
- **WEP/WPA cracking** The process of penetrating an encryption protocol by discovering its cryptographic key. All of the encryption protocols that wireless LANs use, WEP, WPA, and WPA2, are crackable with enough time and effort. The process basically consists of locating a wireless network, using a packet-sniffer program to capture some of its traffic, and then analyzing the contents of the packets to discover the keys used to encrypt them. Wireless network cracking tools are freely available on the Internet, so an attacker does not even have to possess the expertise needed to write them.
- **Evil twin** An unauthorized wireless access point deliberately configured to closely mimic an authorized one. Users fooled by the impersonation connect to the access point, which provides the attacker with access to their data.
- **Rogue access point** An unauthorized wireless access point connected to a network. This is arguably the greatest possible security hazard for a wireless network administrator, because its perpetrators are often innocents. An authorized user wanting the convenience of wireless laptop access in the office purchases an inexpensive AP and plugs it into the network with no security enabled.

True or false: There is no such thing as an encryption protocol that cannot be cracked.

Answer: *True*. All codes can be cracked eventually; it is just a question of how long it will take and how much effort, or computing power, can be devoted to it.

EXAM TIP Candidates for the Network+ exam should be familiar with the various attacks common to wireless networks and be prepared to explain how they work.

True or false: A rogue wireless access point can jeopardize the security of a wired network.

Answer: *True*. A rogue access point with no protection can enable unauthorized users to access the wireless network, and through it, the wired network as well.

Explain common attacks

Understanding the nature of the threats against a network is an essential part of building an effective strategy against them. The relationship between the attackers and the defenders is a type of arms race, with most of the attackers creating ever-different variations on a few familiar themes. The following sections describe some of the most basic attack concepts.

Denial of service

The busier a service is, the longer it takes for customers to get their orders. This axiom applies as easily to networking as it does to retail. When a server is busy processing thousands of incoming requests, performance degrades and all of the clients suffer. A denial of service (DoS) attack is an attempt to provoke this very situation by flooding a server or an application with incoming traffic. In the simplest form of DoS attack, an attacker can use the Ping utility to send an endlessly repeating stream of ICMP messages to a server and bring it to a near halt.

A DoS attack does not have to use ICMP messages, however. All a determined attacker has to do is find an open port on a server, after which it is not difficult to bombard it with some kind of traffic. Some attacks even use the type of traffic the server is designed to accept. For example, an attacker can flood a web server with incoming web requests, excluding the requests generated by legitimate users. In this case, the server administrator can't just close off port 80 in the firewall, because that would exclude all of the incoming traffic. The counter for this type of attack would be to discover where the flood is coming from and block it by its IP address.

Smurf

One particularly sneaky form of ICMP-based DoS attack—called a smurf attack— involves flooding a network with Ping messages sent to the network's broadcast address. These messages are also spoofed; the source address field contains the IP address of the computer that is the intended victim. This way all of the computers receiving the broadcast will send their responses to the victim, flooding its incoming buffers.

Distributed denial of service

A single computer can only transmit so many packets, limiting the potential effectiveness of an individual DoS attack against a large server farm. However, attackers who distribute a type of malware called a Trojan can take control of other people's computers without them knowing it. These remote-controlled computers are called bots or zombies, and the horde of zombies under one attacker's control is called a botnet. The attacker can use the zombies to generate DoS messages from hundreds of systems at once, overwhelming even the most robust application. This is called a distributed denial of service (DDoS) attack.

Man in the middle

A man in the middle (MITM) attack is one in which the attacker interposes himself or herself between two individuals who think they are communicating with each other. The attacker receives the messages from each party in the transaction and relays them to the other party, but not without reading them, or even modifying them, first.

The attacker in a man in the middle attack must be able to receive all of the communications generated by both of the other parties to remain a convincing intermediary. While undetected, MITM attackers can use their access to obtain sensitive information, such as passwords and shared keys.

Detection of an MITM attack is sometimes possible through analysis of the latency periods between message transmissions and receipts. The sudden appearance of communication delays between two parties not attributable to other causes can indicate the presence of an intermediary. The best defense against MITM attacks is mutual authentication, such as that performed by PKI systems. In fact, these authentication systems were invented, in part, to counter this type of attack.

FTP bounce

A variation on the MITM attack, called an FTP bounce attack, involves the use of the Port command in the FTP protocol to gain access to ports in another computer that are otherwise blocked. After connecting to an FTP server, the attacker uses the commands within the FTP protocol, but directs them at a different computer. At one time, many FTP implementations could do this, but most developers have since closed this potentially exploitable opening.

Social engineering

Sometimes the easiest way for an attacker to obtain sensitive information is simply to ask for it. Social engineering is the term used to describe a practice in which a seemingly friendly attacker contacts an authorized user by telephone, mail, or email; pretends to be an official of some sort; and gives some excuse for needing the user's password or other confidential information.

Most people are reasonably helpful by nature, and when someone asks for a favor they can easily perform, they do it. This is particularly true in a corporate environment, where a call from someone you don't know in another department is not unusual. An attacker with a friendly nature and a convincing story can often compel users to give up all sorts of valuable information.

Another, more refined form of this tactic, called phishing, consists of sending out an official-looking email message or letter to users that points them to a website containing a form asking for personal information. The email message might appear to be from a bank, a credit card company, or a government agency, and it might contain a request for help or a threat of some inconvenient action if the user does not comply with the instructions provided. The letter and website are, of course, bogus, and the confidential information the users supply goes right to the attacker.

Malware

Malware is a generic term referring to any software that has a malicious intent, whether obvious or obscure. Some types of malware are relatively benign and are intended only to generate business for the distributor, whereas others are deliberate attempts to cause damage with no rational motive. The most common types of malware are as follows:

- **Virus** A type of program that replicates by attaching itself to an executable file or a computer's boot sector and performs a specified action—usually some form of damage—at a prearranged time. Viruses do not spread through a network by themselves; they require a user to run the infected program and load it into memory. Viruses might also replicate themselves through removable media, such as USB flash drives.
- **Worm** A program that replicates itself across a network by taking advantage of weaknesses in computer operating systems. Unlike a virus, a worm can replicate across a network without any user activity. Some worms do nothing more than consume network bandwidth, whereas others can damage files, generate spam email messages, or install a backdoor program, turning the target computer into a zombie.
- **Trojan** A nonreplicating program that appears to perform an innocent function but that in reality has another, more malicious, purpose. One common tactic is to insert code into a free game or other application, which turns the computer into a server by opening up specific ports to incoming traffic. This enables an attacker on the Internet to take control of the computer without the owner's knowledge and use it as a zombie for any purpose, including initiating distributed attacks against other targets.
- **Spyware** A hidden program that gathers information about your computer activities and sends it to someone on the Internet. Types of spyware can include adware, a relatively harmless program that tracks the Internet sites you visit for the purpose of sending you targeted advertisements. Other types of spyware are more dangerous, such as those that record your keystrokes (and other usage data) to capture your passwords and other sensitive information. Spyware is usually something the user downloads unknowingly, by clicking on a link in an email message or on a webpage.
- **Macro** Macros are application-specific scripts that users can write themselves to automate repetitive tasks. As an application iterates through numerous versions, its developers often strengthen the macro language, sometimes to the point at which it becomes a programming language in its own right. This increased capability enables attackers to exploit the application's macro capabilities and use them to create viruses that replicate and spread through the application.

Buffer overflow

A buffer is an area of computer memory designed to hold incoming data as it is being processed. A buffer overflow is a condition in which a program sends too much data to a buffer and it spills over into an area of memory intended for another

purpose. The results of a buffer overflow depend on the application and the operating system, but they can include error messages, data corruption, or even a system crash.

Ordinarily the result of a programming error, buffer overflows are sometimes the result of people taking advantage of inherent weaknesses in operating systems or applications, and writing code designed to deliberately cause buffer overflows to occur. These attackers can deliver the code to the target system in many different ways, including by viruses, worms, and trojans.

Packet sniffing

A packet sniffer, also called a protocol analyzer, is an application that intercepts and captures packets as they are transmitted over a network. Sniffers are legitimate tools for network administrators, but they are also just as valuable as weapons for attackers, as they enable individuals to capture passwords and other confidential information. For more information, see “Objective 4.2: Given a scenario, use appropriate hardware tools to troubleshoot connectivity issues.”

True or false: Many firewalls are configured to prevent DoS attacks by default.

Answer: *True*. The default configuration on many firewalls blocks the ICMP Echo Request messages that Ping generates, to prevent this simple type of DoS attack.

EXAM TIP The term “virus” is often inaccurately used as a catchall for various types of malware, including worms, trojan horses, and spyware. The primary characteristic that distinguishes a virus is its ability to replicate itself. The Network+ exam expects candidates to know the differences between the various malware types.

True or false: Administrators can configure firewalls to prevent social engineering attacks.

Answer: *False*. Social engineering is not a technical form of attack, and is therefore not preventable using a firewall or other technology. The only protection against social engineering is user training and preparation.

Explain common mitigation techniques

Learning about the threats might be the first step in combatting them, but then the administrator has to devise a strategy for fighting back or preventing them in the first place. Threat mitigation is an ongoing process in which both sides continue to learn. However, there are several standard mitigation techniques that all network administrators should keep in mind, including the following:

- **Training and awareness** In many cases, successful attacks are the result of user error. A person clicks the wrong link, opens the wrong email message, or executes the wrong file, and the door admitting the intruder to the network is opened. Educating users about the potential dangers and what not to do when confronted with them is the best way to protect the network from intrusion.

- **Patch management** Attacks are often possible due to weaknesses in applications or operating systems, which intruders have learned to exploit. Software developers are constantly discovering new weaknesses, and they release updates to close the security holes that result. Keeping all of the computers on the network updated must be an essential part of an administrator's security regimen. In addition, all computers should be equipped with appropriate anti-malware software, which also must be updated on a regular basis.
- **Policies and procedures** For a network to run efficiently, there must be policies in place that govern what administrators and users should and should not do. Published policies can prevent many security breaches before they happen, and proper procedures can enable users to recognize security problems when they happen and take appropriate action.
- **Incident response** Policy should dictate how administrators respond to security-related events, and all threats and attacks should be carefully documented. A history of occurrences can provide evidence of an escalation of tactics, indicating that attackers are targeting the organization specifically.

True or false: Keeping software and operating systems updated is one of the best ways to combat DoS attacks.

Answer: *False*. DoS attacks are not exploitations of software weaknesses. Rather, they are attacks that take advantage of a system's normal operations. Preventing them is usually a matter of careful firewall configuration, rather than patch management.

EXAM TIP Candidates for the Network+ exam should be familiar with common mitigation techniques and the types of attacks they can prevent.

True or false: Some threats can only be mitigated by educating users.

Answer: *True*. Some attacks, such as social engineering, are psychological, not technological. The only way to prevent them is to teach users to recognize the signs of a person trying to obtain confidential information.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. What is the term used to describe a flood of useless packets deliberately directed at a particular computer?
2. What term describes an attack that consists of nothing more than a person calling users on the phone and tricking them into supplying sensitive information?
3. Which of the attacks described in this objective uses a botnet to bombard a target with traffic?
4. What tools are needed by an individual performing a war driving attack?

Objective 5.5: Given a scenario, install and configure a basic firewall

A firewall is a hardware or software product that protects a network from unauthorized access by outside parties, while letting appropriate traffic through, as shown in Figure 5-16. If your network is connected to the Internet, you must have some sort of firewall to protect it, because otherwise intruders can wreak havoc on the network that you have so carefully designed and constructed.

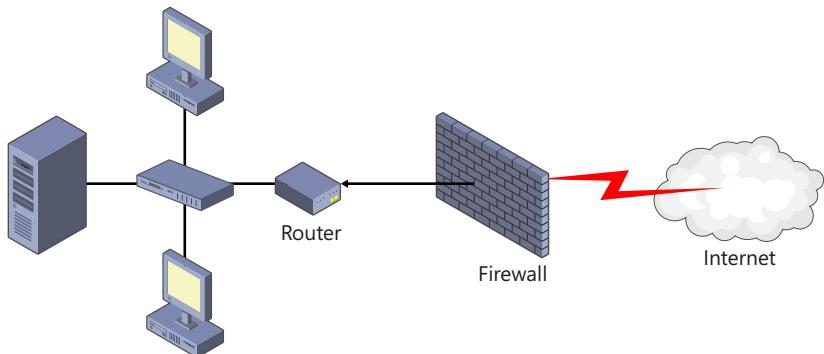


FIGURE 5-16 A firewall blocking unauthorized network traffic.

For this objective, you must be familiar with the various types of firewalls and how they operate.

Exam need to know

- Install firewalls
 - For example:* How is the installation of software and hardware firewalls different?
- Configure port security
 - For example:* How does a firewall protect a system's open ports?
- Configure stateful inspection vs. packet filtering
 - For example:* How does stateful packet inspection provide more protection than packet filtering?
- Configure firewall rules
 - For example:* Under what conditions would you use block/allow rules, implicit deny rules, and ACLs?
- Configure NAT/PAT
 - For example:* How does network address translation enable privately addressed computers to access the Internet?
- Configure a DMZ
 - For example:* How do you construct a DMZ?

Install software and hardware firewalls

A firewall is essentially a barrier between two networks that evaluates all incoming or outgoing traffic to determine whether it should be permitted to pass through to the other network, a process called packet filtering. A firewall can take many different forms and can use different criteria to evaluate the network traffic it receives. Some firewalls are dedicated hardware devices, essentially routers with additional software that monitors incoming and outgoing traffic. In other cases, firewalls are software products that run on a standard computer.

Firewalls can use several methods to examine network traffic and detect potential threats. Most firewall products use more than one of these methods and often provide other services as well. For example, one firewall product, called a proxy server, not only allows users to access webpages with complete safety, but also can cache frequently used pages for quicker retrieval by other systems.

Firewalls typically take two forms, as follows:

- **Hardware firewall** Also called a network-based firewall or a router firewall, this type of firewall is implemented as part of a router and protects the entire network to which it is connected from intrusions originating on another network.
- **Software firewall** Also called a host-based firewall or a personal firewall, this type is implemented on a computer and protects only that computer from intrusions originating on the network.

True or false: Every computer on a network must have a software firewall.

Answer: *False*. If a network is adequately protected by hardware firewalls on Internet access routers, it is not imperative that each computer have a software firewall.

EXAM TIP With respect to firewalls, the Network+ exam uses the terms “software firewalls” and “hardware firewalls,” which can sometimes be misleading. For example, a computer running Windows Server can function as a router, meaning that the Windows Firewall program is technically functioning as a hardware firewall—because it provides protection for an entire network—even though it is implemented as software.

True or false: In addition to protecting a LAN from the Internet, you can also use firewalls internally to protect one section of the network from the rest.

Answer: *True*. You can use internal firewalls to erect barriers between departments. For example, you can use a firewall to isolate the LAN used by your company’s accounting department to prevent other users from accessing confidential financial records.

Configure port security

One of the chief functions of a firewall is to protect systems from intrusion through their open ports. Filtering by port numbers, also known as *service-dependent filtering*, is the most common and flexible type of packet filtering performed by firewalls.

Because port numbers represent specific applications, administrators can use them to prevent traffic destined for these applications from reaching a system or a network.

For example, to protect a perimeter network containing your company's web servers, you can create filters that allow only traffic using port 80 to enter from the Internet, blocking all other application ports.

MORE INFO For more information on the well-known port numbers assigned to specific applications and services, see "Objective 1.5: Identify common TCP and UDP default ports."

Configure stateful inspection vs. packet filtering

A packet filter is the most basic type of firewall, one in which the system implementing the filter examines each packet as it arrives and decides if it meets the criteria for admission to the network. Packets that do meet the admission criteria are processed by the system in the normal manner; those that do not are silently discarded.

For example, many Internet email servers SMTP for outgoing traffic and the Post Office Protocol 3 (POP3) for incoming traffic. These protocols use the well-known port numbers 25 and 110, respectively. You can configure a firewall with a packet filter that permits only packets destined for port numbers 25 and 110 to pass through, as shown in Figure 5-17. Packets destined for any other port numbers are discarded before they can reach the network and do any damage.

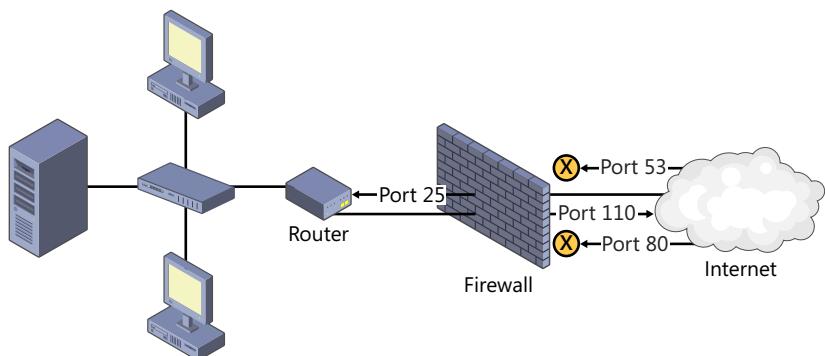


FIGURE 5-17 A firewall filtering out all packets for ports other than 25 and 110.

In addition to port numbers, the criteria most commonly used in packet filtering are protocol identifiers, IP addresses, and MAC addresses.

- **Protocol identifiers** The Protocol field in every packet's IP header contains a code that identifies the next protocol that should receive the packet in the networking stack of the destination. In most cases, the code represents a transport layer protocol, such as TCP or UDP. However, IP datagrams frequently carry ICMP messages as well. Filtering by using protocol identifiers is not very precise because it blocks or allows all the traffic that uses a particular protocol. However, for certain applications, blocking an entire

protocol is necessary, and it is easier than anticipating the specific applications an attacker might use. For example, if you have a network that contains only Internet web and FTP servers, you could use protocol filters to limit incoming traffic to TCP packets. Because these servers rely on TCP for their primary functions, you could block all UDP and ICMP traffic, preventing attacks from using any applications that rely on these protocols.

- **IP addresses** IP address filtering lets you limit network access to specific computers. For example, if you have an Internet web server on a LAN with other computers, and you want Internet clients to be able to access only the web server, you can create a filter permitting only those packets addressed to the web server to enter the network from the Internet. You can also use IP address filtering to protect part of a private network. You can create filters that give only certain computers access to the protected LAN, while preventing all others from accessing it.
- **MAC addresses** Filtering based on MAC addresses (or hardware addresses) provides the same basic functionality as IP address filtering. However, it is more difficult to spoof a hardware address than it is an IP address, so MAC address filters are inherently more secure than IP address filters. MAC address filtering is rarely used on Internet routers or firewalls, but for wireless access points and internal filtering, MAC address filtering is a useful means of restricting access to specific resources.

There are two main drawbacks to using packet filtering as a security mechanism:

- Packet filtering requires a detailed understanding of TCP/IP communications and the ways of the criminal mind. Using packet filters to protect your network means participating in an ongoing battle with attackers. Intruders are constantly inventing new techniques to defeat standard packet filter configurations, and you must be ready to modify your filters to counteract these techniques.
- Packet filters can only detect attacks implemented in the packet headers; they do not examine the application data inside the packets. For example, you might configure the packet filters on your firewall to allow all port 80 traffic into the network so that Internet users can access your web server, but at the same time, you could be admitting packets that are designed to attack the web service itself. To examine the application layer data in the packets, you must use a proxy server.

Some packet-filtering firewalls include additional security capabilities, typically in the form of a technique called stateful packet inspection. Stateful packet inspection is a generic term for a process in which a router examines incoming packets more carefully than usual.

In a typical packet-filtering firewall, when the router examines packets it is concerned only with the basic criteria listed earlier, such as port and protocol numbers and IP addresses. A firewall that supports stateful packet inspection examines other network and transport layer header fields as well, looking for patterns that indicate damaging behaviors, such as IP spoofing, SYN floods, and teardrop attacks. SYN floods and teardrop attacks are two forms of DoS attacks in

which a target is bombarded with large numbers of packets containing TCP SYN flags or datagrams requiring fragmentation, respectively.

The router also tracks the connections between the systems generating packets by examining the Sequence Number values in the TCP headers. This enables the router to determine the current state of each connection. To gain admittance to the network, packets not only must meet the requirements of the packet filters, but they must also be part of a connection listed in the router's state table.

Because the router must examine each packet more carefully, firewalls that use stateful packet inspection are necessarily slower than simple packet-filtering firewalls. Firewalls with stateful packet inspection are usually also more expensive than firewalls without it.

Though there are some free stateful packet inspection firewalls, such as the Linux Netfilter module, most commercial products are quite costly. However, most commercial products include a graphical configuration interface and better documentation, which makes it easier to set up and maintain the firewall. Different manufacturers implement stateful packet inspection in different ways, so not all routers with this capability offer the same degree of protection.

True or false: When configuring a firewall, you must select one type of packet filter; you cannot combine filters of different types.

Answer: *False*. The real strength of using packet filtering as a security mechanism comes when you combine different types of filters to create a composite solution. For example, you might want to open up the Telnet port (port 23) so administrators can remotely manage the company web servers from home, using the Internet. However, leaving this port open is an invitation for unauthorized Internet users to access your servers. You could add an additional filter that limits port 23 access to only your administrators' IP addresses. This would protect the network without compromising the functionality that the administrators need.

EXAM TIP Candidates for the Network+ exam should be familiar with the four criteria used for packet filtering, as well as the differences between packet filtering and stateful packet inspection.

True or false: Configuring packet filters in the real world is often a complex trial and error process.

Answer: *True*. Though packet filtering might sound simple, that is rarely the case. For example, an email server in reality almost certainly receives legitimate traffic that uses ports other than 25 and 110. This traffic uses different port numbers, and for the server to function properly, its packet filters must admit all legitimate traffic while blocking everything else. Manually implementing packet filters, therefore, requires a careful analysis of a computer's traffic patterns, as well as a certain amount of trial and error, before you can arrive at a workable filtering solution. For this reason, most firewall products have preconfigured filter patterns that let you select an application or a protocol, after which the software configures the packet filters for the appropriate ports automatically.

Configure firewall rules

Creating packet filters is a matter of selecting the specific criteria you want the system to examine and specifying the values that it will allow or deny passage through the filter. Administrators can configure a firewall to filter packets in two basic ways, as follows:

- **Inclusive** The network interface is completely blocked and packet filters specify what traffic can pass through.
- **Exclusive** The network interface is left completely open and packet filters specify what traffic to block.

These two approaches are basic philosophies that pertain to most, if not all, security technologies. An inclusive packet filter implementation (an application of what is sometimes called the implicit deny rule) is inherently more secure; however, it can be more difficult to debug because you must make sure that all the traffic that needs to pass through the filters is getting through. Exclusive filtering assumes that the administrator is familiar with all possible threats to the network and knows what filtering criteria the firewall must use to stop them. This is a dangerous presumption.

Most operating systems have a firewall with packet-filtering capabilities built into it, but the Windows Server operating systems actually have two. The Routing and Remote Access Service (RRAS) that the operating system uses to function as a router includes a packet-filtering mechanism with which you can create different filters for each network interface on the computer. With RRAS packet filtering, you can apply the following additional capabilities:

- Create filters based on the IP addresses, protocols, and port numbers of a packet's source or destination
- Create inclusive or exclusive filters
- Create filters for inbound or outbound traffic
- Create filters for ICMP messages, specified by the message type and code values
- Create multiple filters of the same type

The interface you use to create packet filters in the Routing and Remote Access console is shown in Figure 5-18.

All of the Windows operating systems also include Windows Firewall, which is an automated packet-filtering implementation that can control what network traffic enters and leaves the computer. Windows Firewall has two interfaces, a simpler one that is designed more for the user, and a more comprehensive interface for administrators.

In the Windows Firewall control panel interface, as shown in Figure 5-19, users can select the programs and features that they want to admit through the firewall. The program then adjusts the packet filters accordingly, so the users do not have to be concerned with port and protocol numbers and other technical details.

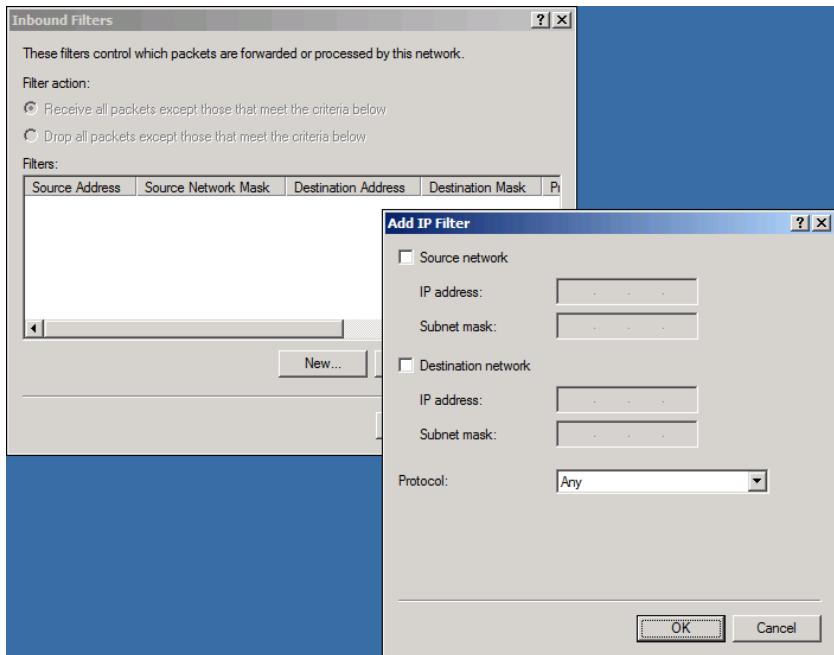


FIGURE 5-18 The Routing and Remote Access console's Inbound Filters and Add IP Filter dialog boxes.

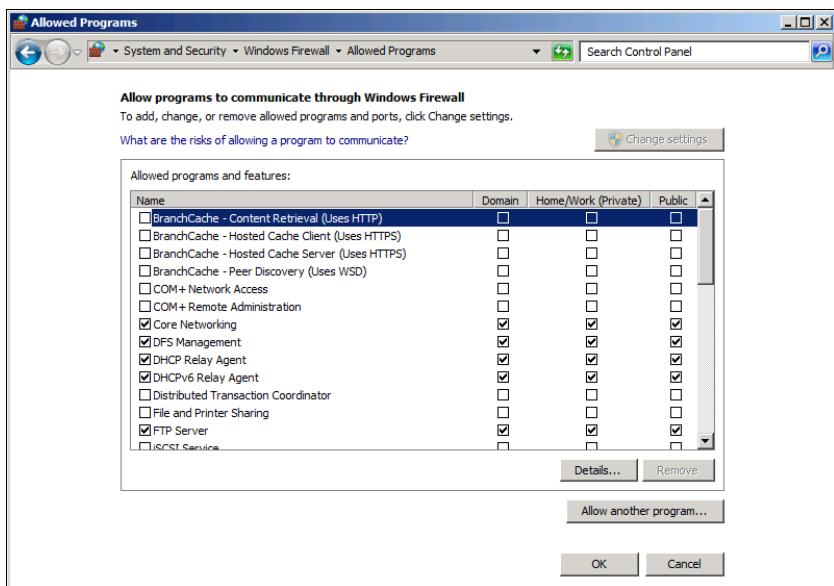


FIGURE 5-19 The Windows Firewall control panel interface.

The other interface is the Windows Firewall with Advanced Security snap-in for Microsoft Management Console (MMC). This snap-in provides a complete list of all the packet filters on the system, as shown in Figure 5-20, and with it you can activate and deactivate them, or modify them, or create new filters of your own. With this interface, the administrator is essentially working directly with the ACL containing the filters.

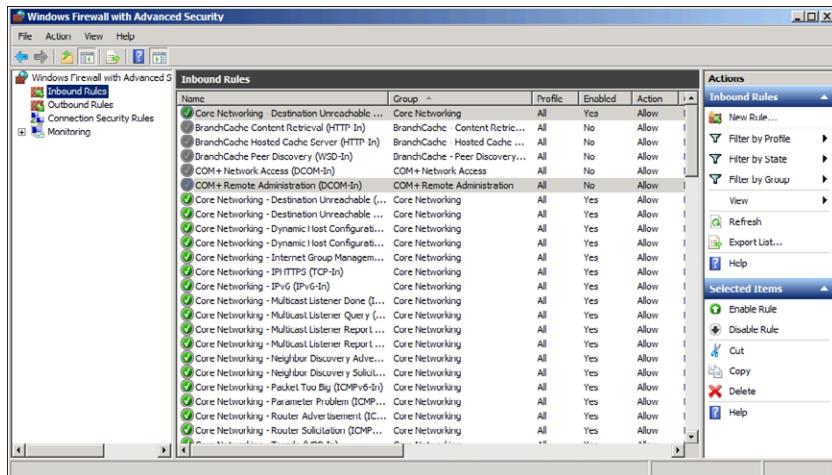


FIGURE 5-20 The Windows Firewall with Advanced Security snap-in.

Most UNIX and Linux distributions have packet filtering built into the kernel, although the various operating systems have different tools for managing the filtering rules. For example, Berkeley Software Distribution (BSD) UNIX uses a tool called *ipfw*, whereas Linux uses *iptables*. Both of these are command-line tools with extensive sets of arguments and parameters.

Hardware routers have interfaces with similar capabilities to the ones provided by operating systems, but their ease of use can vary from product to product. The broadband routers designed for home and small business users have relatively simple interfaces, like that shown in Figure 5-21. More advanced routers provide more comprehensive access and often require more knowledge from the administrator.

In addition to routers and operating systems, there are standalone firewall products you can run on any computer that also have packet-filtering capabilities. Firewall-based filters have the following two advantages:

- **Better performance** By implementing the routing and filtering functions on different systems, you are less likely to experience degraded network performance.
- **Increased filtering capabilities** Dedicated firewall products are likely to have more advanced packet-filtering capabilities, such as preset filter configurations designed to protect against specific types of attacks.

Application	Start	End	Protocol	IP Address	Enable
smtp	25	25	Both	192.168.2.1	<input type="checkbox"/>
http	80	80	Both	192.168.2.1	<input type="checkbox"/>
ssl	443	443	Both	192.168.2.1	<input type="checkbox"/>
sbs	987	987	Both	192.168.2.1	<input type="checkbox"/>
rvw	3389	3389	Both	192.168.2.1	<input type="checkbox"/>
smtp	25	25	TCP	192.168.2.2	<input type="checkbox"/>
http	80	80	TCP	192.168.2.2	<input type="checkbox"/>
ssl	443	443	TCP	192.168.2.2	<input type="checkbox"/>
smtp	25	25	Both	192.168.2.13	<input checked="" type="checkbox"/>
vlab	5800	5921	Both	192.168.2.10	<input type="checkbox"/>

FIGURE 5-21 Packet-filtering configuration in a broadband router.

NOTE The basic capabilities of most packet-filtering implementations are roughly the same; the differences are in the interface and the configurability of the filters. Two products might have the same packet-filtering capabilities, but one with preset configurations and detailed documentation will be easier to use than one that requires you to design filter configurations yourself and fully understand the TCP/IP communications processes that are affected by the filters you are creating.

Packet filtering is not a perfect security solution. Intruders can still attack a server by using the ports and protocols that the firewall lets through, or find a clever new way to bypass the filters you have in place. The trick to using packet filters effectively is to strike a balance between providing sufficient access to legitimate users and blocking enough traffic to provide protection.

In some cases, the creation of packet filters can be an ongoing battle of wits between the protector and a determined attacker. Every time the attacker finds a way to penetrate the filters, the system administrator modifies them to close the opening that is being exploited. Advanced packet filtering requires a detailed understanding of the TCP/IP protocols and the applications that use them.

True or false: Packet filtering can work in either direction.

Answer: *True*. You can use filters to prevent users on the Internet from accessing your private network, or you can use them to limit the types of Internet access granted to your internal users.

True or false: Routers can filter packets with no appreciable effect to their performance.

Answer: *False.* Packet filters can introduce a large amount of overhead, slowing down a router's performance. The router must examine each incoming packet, compare it against all of the filters, and then decide whether to admit the packet to the network. If you have a large, complex system of filters, the amount of time needed for the router to process each packet can become a major network performance bottleneck.

Configure NAT/PAT

Network address translation (NAT) is a routing technique that enables computers with private (that is, unregistered) IP addresses to access the Internet. If you connect a network to the Internet without firewall protection of any kind, you must use registered IP addresses for your computers so that they can communicate with other systems. However, registered IP addresses are visible from the Internet. This means that any user on the Internet can access your network's computers and, with a little ingenuity, wreak havoc on your network. NAT prevents this from happening by enabling you to assign unregistered IP addresses to your computers.

The Internet Assigned Numbers Authority (IANA) has designated three address ranges for use on private networks. These address ranges are not registered to any Internet user and are not visible from the Internet. You can safely deploy them on your computers without the danger of exposing them to Internet intruders.

MORE INFO For more information on private network addressing, see "Objective 1.3: Explain the purpose and properties of IP addressing."

However, this also means that Internet servers, when they receive requests from the private network computers, cannot send replies to them, because they do not have viable addresses. NAT solves this problem by functioning as an intermediary between the Internet and a client computer on an unregistered network. For each packet generated by a client, the NAT router substitutes a registered address for the client's unregistered address.

Under normal conditions, routers do not modify datagrams any more than the postal service modifies envelopes. A NAT router, however, modifies each outgoing datagram it receives from an unregistered client computer by changing the value of the Source IP Address field in its IP header. The following steps explain this process:

1. When a client sends a request message to an Internet server, the datagram containing the request first goes to a NAT router.
2. NAT substitutes a registered IP address for the client computer's unregistered IP address in the datagram and then forwards it to the destination server on the Internet. The NAT router also maintains a table of those unregistered addresses and the public address assigned to them, in order to keep track of the datagrams it has processed.
3. When the destination server on the Internet receives the request, it processes it in the normal manner and generates its reply. However, because the

Source IP Address value in the request datagram is the NAT router's registered address, the destination server addresses its reply to the NAT router, not to the original client.

- When the NAT router receives the reply from the Internet server, it modifies the datagram again, substituting the client's original, unregistered address for the Destination IP Address in the datagram's IP header, and forwards the packet to the client on the private network.

The NAT router's processes are invisible both to the client on the private network and to the server on the Internet. The client generates a request and sends it to a server, and the client eventually receives a reply from that server. The server receives a request from the NAT router and transmits its reply to the same router. Both the client and the server function normally, unaware of the NAT router's intervention. More importantly, the client computer remains invisible to the Internet and is protected from most types of unauthorized access originating from outside the private network.

There are several different types of NAT, some of which share the same acronyms, as follows:

- Static NAT** Static NAT (SNAT) translates multiple unregistered IP addresses to an equal number of registered addresses. This enables each client to always use the same registered address. This type of NAT does not conserve IP address space, because you need the same number of registered addresses as unregistered addresses. Static NAT is also not as secure as the other NAT types, because each computer is permanently associated with a particular registered address. This makes it possible for Internet intruders to direct traffic to a particular computer on your network by using that registered address.
- Dynamic NAT** Dynamic NAT is for situations when you have fewer registered IP addresses than unregistered computers. Dynamic NAT translates each unregistered address to one of the available registered addresses. Because the registered address assigned to each client changes frequently, it is more difficult for intruders on the Internet to associate a registered address with a particular computer, as in static NAT. The main drawback of dynamic NAT is that it can support only the same number of simultaneous users as the number of available registered IP addresses. If all the registered addresses are in use, a client attempting to access the Internet receives an error message.
- Port address translation (PAT)** Also known as masquerading, this method translates all the unregistered IP addresses on a network by using a single registered IP address, as shown in Figure 5-22. The NAT router uses port numbers to differentiate between packets generated by and destined for different computers, so that multiple clients can access the Internet simultaneously. Masquerading provides the best security of the NAT types because the association between the unregistered client and the registered IP address/port number combination in the NAT router lasts only for a single connection.

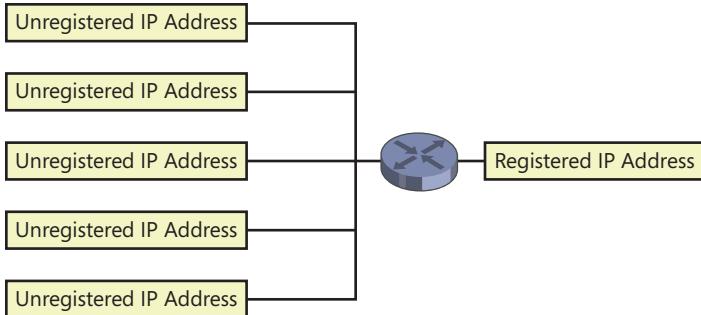


FIGURE 5-22 Port address translation.

Today, most NAT implementations rely on masquerading because it minimizes the number of registered IP addresses needed, and because it maximizes the security provided by NAT. However, NAT by itself, even if it uses masquerading, is not a true firewall, and it does not provide ironclad security for high-risk environments. NAT effectively blocks unsolicited requests and other probes from the Internet, meaning that it prevents intruders from searching for unprotected file shares, open ports, and private web or FTP servers on the private network. However, NAT does not prevent users on the Internet from launching directed DoS attacks against specific computers or from using other more complex tactics to compromise your network security.

NAT is implemented in a variety of products, both hardware and software. Most of the hardware routers that provide shared Internet access support NAT, particularly the low-end devices intended for home or small business use.

Many operating systems support NAT as well. The Windows operating systems have two different NAT implementations. In the Windows Server operating systems, including Windows Server 2008 R2, NAT is integrated into RRAS, enabling a computer that is functioning as a router to translate IP addresses between any two network interfaces. The workstation operating systems, including Windows 7, have a feature called Internet Connection Sharing (ICS), which is similar to the NAT DHCP implementations in stand-alone router products. With ICS, a computer can share a network connection to an ISP with other computers on the LAN, allowing them to access the Internet by using the ISP-connected computer as a router.

UNIX and Linux typically implement NAT as part of the operating system kernel. In Linux, NAT is integrated into the same Netfilter component that provides packet filtering capabilities. You use the same iptables tool to configure both the NAT process and packet filtering.

True or false: NAT works with any IP application running on a protected computer.

Answer: *True*. Because NAT functions at the network layer of the OSI model, it works with any application that communicates by using IP. Client computers on the private network can run Internet email clients, web browsers, FTP clients, or any other Internet application, and NAT provides protection against intruders.

EXAM TIP There are several NAT product manufacturers who use the acronym SNAT to refer to different technologies, including secure NAT and stateful NAT. However, static NAT, referring to the use of one registered address for each unregistered address, is the most common spelling out of the acronym SNAT.

True or false: PAT requires an equal number of private and public IP addresses.

Answer: *False*. PAT requires only one public IP address, which it uses to serve any number of private IP addresses.

Configure a DMZ

Router-based firewalls are an integral part of building a network with concentric layers of security. The outermost layer, the one nearest to the Internet and the dangers it represents, is called a peripheral network or a demilitarized zone (DMZ). The peripheral network is where you place the servers that must be accessible from the Internet, such as web, FTP, and SMTP servers, as shown in Figure 5-23. The router connecting this network to the Internet (or to an ISP's network) contains a firewall that you configure to admit the traffic that must reach these servers.

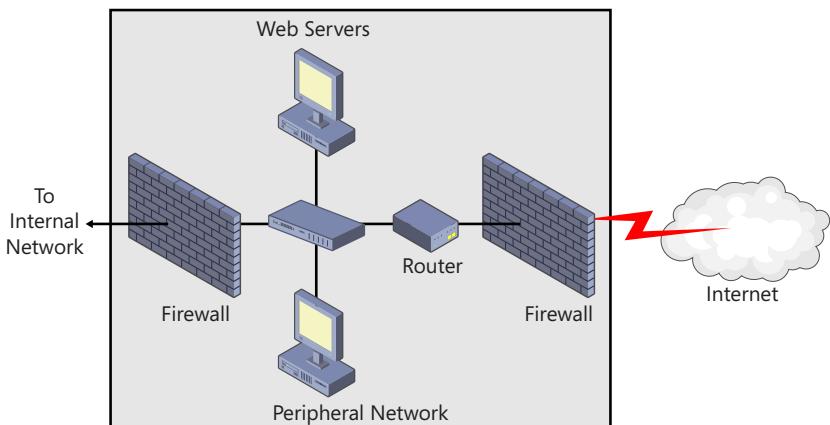


FIGURE 5-23 A peripheral network and the firewall between it and the Internet.

The peripheral network also has another router, which connects it to the innermost network, the one that contains the servers requiring the most protection, such as domain controllers, intranet servers, and file servers hosting confidential files. The firewall on this inner router blocks the traffic that the other firewall admits to the peripheral network. Therefore, users on the Internet can send request messages to the web servers on the peripheral network, but not to the intranet web servers on the inner network.

To create an effective peripheral network, administrators must be aware of the traffic that needs to pass between the outer network and the inner one. For example, if your web server runs an application that requires access to a database, you might want to locate the database server on the more secure inner network.

In this case, you must configure the inner firewall to pass the database traffic, but exclude the web server traffic from the Internet.

True or false: Creating a DMZ requires two firewalls.

Answer: *True*. One firewall separates the DMZ from the Internet, and the other separates the DMZ from the internal network.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Service dependent filtering blocks traffic using what element of the incoming packets?
2. What must you do to configure a firewall to admit FTP traffic using the default port settings to the internal network?
3. A firewall that scans transport layer header fields for evidence of SYN floods uses what type of scanning?
4. How does NAT protect a network from Internet intrusion?

Objective 5.6: Categorize different types of network security appliances and methods

Firewalls are not the only devices that administrators can use to increase the security of a network. There are a variety of other tools and techniques that help to keep the sides even in the network security arms race. Some of these tools and techniques are discussed in the following sections.

Exam need to know

- Categorize IDSe and IPSes
For example: What are the differences between behavior-based, signature-based, network-based, and host-based systems?
- Categorize vulnerability scanners
For example: What are the capabilities of the Nessus and Nmap products?
- Categorize honeypots and honeynets
For example: How do honeypots protect a network?

Categorize IDSe and IPSes

A firewall is a static device, a filter that you configure and that admits or blocks traffic based on the criteria you select. A firewall of this type is not smart, in that the device does not know whether packets contain an attack or not. The firewall simply looks for specific bit patterns in specific areas of incoming packets and reacts accordingly.

An intrusion detection system (IDS) is a network protection product, implemented as hardware or software, which operates at a somewhat higher level on the scale of artificial intelligence. An IDS inspects incoming packets or system processes and

examines them for evidence of malicious activity. When the IDS finds such evidence, it logs the activity, gathers information about the circumstances, and usually notifies an administrator by email or some other means.

There are two basic types of IDSes, as follows:

- **Network-based IDS (NIDS)** A separate device, typically connected to a switch or router at a strategic network interface point, such as a peripheral network. A NIDS accesses network traffic, often by using a technique such as port mirroring, and examines the contents of packets, looking for patterns and data types typical of well-known attacks.
- **Host-based IDS (HIDS)** An application or agent running on a computer that analyzes system and application event logs, performs file integrity checks, and monitors registry modifications for evidence of malicious activity.

In both cases, the IDS includes an interface to a network management program or provides management capabilities of its own. The primary function of the product is to notify administrators when the system detects suspicious activity. Some IDS products are standalone applications, whereas others are distributed, meaning that they can have multiple information-gathering sensors or agents scattered around the network, all of which report their findings to a central management console.

NIDS products are signature based, meaning that they compare the data in the packets they scan to a library of known attack patterns, looking for matches that indicate problems. As new threats are discovered, the manufacturer of the product must update the signature files and release them in the form of updates. An NIDS must be current to be effective.

HIDS products work in a different manner, by tracking system performance patterns and establishing their own baselines. They then continue to monitor performance and note behavior that deviates substantially from those baselines. These types of products are said to be behavior based.

An IDS is only capable of detecting anomalous behavior on a network or a host; it can't do anything about the potential attacks it discovers. Other products, called intrusion prevention systems (IPSes) can, however. An IPS functions in much the same way as an IDS, except that it is also designed to take specific actions to prevent an attack when it detects one. Among other things, an IPS can discard packets, terminate connections, or create filters on the fly that prevent packets using certain ports or addresses from entering the network.

True or false: Network-based IDS products require a switch with port mirroring because it is the only way they can monitor all of the traffic on the network.

Answer: *True*. Port mirroring is a feature of some switches that copies all incoming traffic to a special port, as well as to the dedicated destination ports. By connecting an IDS or protocol analyzer to this port, an administrator can access all of the network traffic. For more information on port mirroring, see “Objective 1.4: Explain the purpose and properties of routing and switching.”

EXAM TIP The Network+ exam distinguishes between IDS and IPS types in two ways: by location, using the terms “network-based” and “host-based,” and by method,

using the terms “behavior-based” and “signature-based.” Candidates for the exam should be familiar with all four of these distinctions.

Categorize vulnerability scanners

A server, by definition, is a software program that stands ready to receive and service requests from clients. For those clients to be able to communicate with the server over a network, the server must open a port and listen for incoming traffic over that port. For example, a web server opens TCP port 80, because that is the well-known port for HTTP traffic, the port that web browsers use to communicate by default.

However, every port left open on a server computer is essentially an unlocked door through which intruders can conceivably enter. A port scanner is a software product that displays all of the open ports on a computer or on a network’s computers. For network administrators, this is a type of vulnerability scanner that they can use to detect potential security breaches. For the attacker, it can also be a means of finding unprotected entrances into a network.

There are many port scanners available, including Netstat.exe, a relatively rudimentary one provided with all versions of Windows. Arguably the most popular one, however, is Nmap, an open-source, command line program originally designed for UNIX systems, but that is now available for Windows with command-line and graphical interfaces.

When you initiate a scan with Nmap by supplying it with a network IP address, the program searches the available IP addresses for functioning computers and then runs a series of scripts against each one. The GUI version of the program, called Zenmap, displays the information in a series of screens for each system it finds, as shown in Figure 5-24.

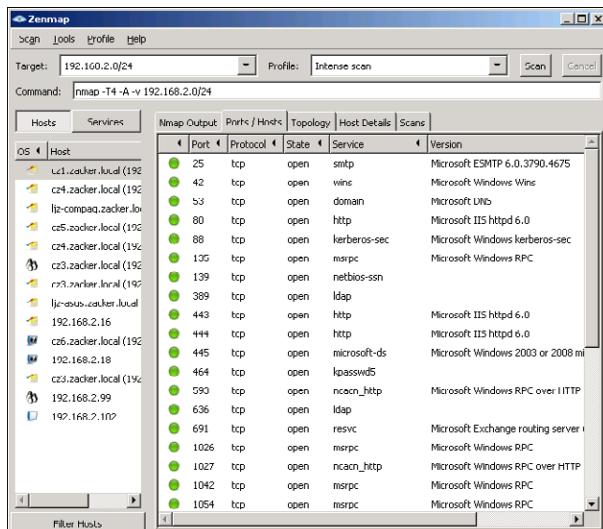


FIGURE 5-24 A graphical Nmap display in Zenmap.

A vulnerability scanner is a software program that attempts to discover weaknesses in the security of a network and its computers. The differences between protocol analyzers and port scanners and vulnerability scanners can sometimes be subtle. For example, a vulnerability scan typically begins with a port scan, in which the software discovers the computers on the network and locates any open ports they might have.

However, a vulnerability scanner then proceeds to launch a variety of attacks against the open ports, attempting to exploit their vulnerabilities. Although the attacks use the same techniques that actual intruders might use, they are designed not to cause any real damage. However, some products enable you to switch off this safe mode of operation so that you can launch actual attacks against a test platform, to ascertain its vulnerability.

One of the most popular vulnerability scanners is called Nessus, produced by Tenable Network Security. Originally released as open-source software in 1998, Nessus is now a commercial product. Over the years, many additional features have been added, so that in addition to port scanning and a constantly expanding library of exploits, Nessus can now perform the following additional tasks:

- Scan accounts for weak or missing passwords
- Audit anti-virus software configurations
- Scan for missing operating system patches and updates
- Check for compliance with system configuration policies
- Locate systems participating in peer-to-peer networks
- Scan for various forms of malware

After completing its scan, Nessus can display a report listing all of the exploits to which computers on the network are vulnerable, as shown in Figure 5-25.

Plugin ID	Count	Severity	Name	Family
47709	1	Critical	Microsoft Windows 2000 Unsupported Installation Detection	Windows
34460	2	High	Obsolete Web Server Detection	Web Servers
41028	1	High	SNMP Agent Default Community Name (public)	SNMP
26920	6	Medium	Microsoft Windows SMB NULL Session Authentication	Windows
57608	6	Medium	SMB Signing Disabled	Misc.
51192	5	Medium	SSL Certificate Cannot Be Trusted	General
57690	4	Medium	Terminal Services Encryption Level is Medium or Low	Misc.
12217	3	Medium	DNS Server Cache Snooping Remote Information Disclosure	DNS
18405	3	Medium	Microsoft Windows Remote Desktop Protocol Server Man-in-the-Middle Weakness	Windows
20007	3	Medium	SSL Version 2 (v2) Protocol Detection	Service detection

FIGURE 5-25 A Nessus vulnerability scan report.

True or false: A computer with no open ports is unable to function as a server.

Answer: *True*. A server, by definition, listens for incoming messages from clients. A computer can only receive unsolicited incoming traffic through a port that it has left open.

EXAM TIP Most of the network administration utilities mentioned in the Network+ exam objectives are included with operating systems. Nmap and Nessus are two that are not. Both are third-party products, available in free evaluation versions. Candidates for the exam should be familiar with both.

True or false: A port scanner works by launching attacks against the open ports in a computer.

Answer: *False*. A port scanner only lists the open ports; it does not launch attacks against them. A vulnerability scanner has the ability to launch attacks.

Categorize honeypots and honeynets

Learning about your adversary is one of the best ways to gain an advantage in any type of conflict. One of the tools that network administrators use to learn about the attacks that threaten them is the honeypot, a system designed to function as a lure for attackers.

Although the honeypot contains no actual data of any value, the administrator configures it with the correct applications and settings so that it looks as though it might. After it is in place, the honeypot sits on the network looking attractive to potential intruders.

In this role, the honeypot can be valuable in two ways. First, it functions as a decoy, drawing attacks that could otherwise be directed at the network's actual resources. Second, the honeypot can gather information about the exact nature of the attacks, which can help the administrators bolster the network's defenses.

On a larger scale, administrators might set up several honeypot computers, to create a decoy server farm or workstation LAN, for the same purpose. This configuration is known as a honeynet.

True or false: As security tools, the primary drawback of honeypots and honeynets is their high cost.

Answer: *True*. Honeypots and honeynets are fully functional servers or networks that are set out as decoys to attract attacks. As a security measure, they are much more expensive than virtually any of the standard diagnostic tools.

Can you answer these questions?

Find the answers to these questions at the end of this chapter.

1. Which of the products discussed in this objective can identify a network's security weaknesses by mounting attacks against the systems?
2. What feature does an IDS need to monitor all the traffic on a switched network?
3. Which type of IDS uses signatures that must be updated regularly?

Answers

Objective 5.1: Given a scenario, implement appropriate wireless security measures

1. By permitting only devices with specified MAC addresses to connect to an access point.
2. Advanced Encryption System (AES).
3. WPA2-Enterprise calls for the use of RADIUS, while WPA2 does not.
4. The use of 40-bit encryption keys, 24-bit initialization vectors, and static shared secrets.

Objective 5.2: Explain the methods of network access security

1. Link Control Protocol.
2. Transport Layer Security (TLS).
3. IP Authentication Header.
4. PPPoE.
5. Telnet.

Objective 5.3: Explain methods of user authentication

1. Kerberos.
2. Only the user's private key can decrypt data encrypted with that user's public key.
3. EAP.
4. The authentication server role.

Objective 5.4: Explain common threats, vulnerabilities, and mitigation techniques

1. Denial of service (DoS) attack.
2. Social engineering.
3. A smurf attack.
4. A portable computer and a vehicle.

Objective 5.5: Given a scenario, install and configure a basic firewall

1. Port numbers.
2. You must open TCP ports 20 and 21.
3. Stateful packet inspection.

4. NAT enables a network with private IP addresses to send messages to the Internet, but systems on the Internet cannot send messages directly to the systems on the protected network.

Objective 5.6: Categorize different types of network security appliances and methods

1. Nessus.
2. Port mirroring.
3. Network-based IDS products are signature based.

APPENDIX

Network+ Acronyms

ACRONYM	STANDS FOR
AAA	Authentication, Authorization, and Accounting
ACL	Access Control List
ADSL	Asymmetric Digital Subscriber Line
AES	Advanced Encryption Standard
AH	Authentication Header
AM	Amplitude Modulation
APIPA	Automatic Private Internet Protocol Addressing
ARIN	American Registry for Internet Numbers
ARP	Address Resolution Protocol
ASP	Application Service Provider
ATM	Asynchronous Transfer Mode
BERT	Bit-Error Rate Test
BGP	Border Gateway Protocol
BNC	British Naval Connector / Bayonet Niell-Concelman
BootP	Boot Protocol /Bootstrap Protocol
BPDU	Bridge Protocol Data Unit
BRI	Basic Rate Interface
CARP	Common Address Redundancy Protocol
CHAP	Challenge Handshake Authentication Protocol
CIDR	Classless Inter-Domain Routing
CNAME	Canonical Name
CRAM-MD5	Challenge-Response Authentication Mechanism – Message Digest 5
CSMA / CA	Carrier Sense Multiple Access / Collision Avoidance
CSMA / CD	Carrier Sense Multiple Access / Collision Detection
CSU	Channel Service Unit
dB	decibels
DHCP	Dynamic Host Configuration Protocol

ACRONYM	STANDS FOR
DLC	Data Link Control
DMZ	Demilitarized Zone
DNS	Domain Name Service / Domain Name Server / Domain Name System
DOCSIS	Data-Over-Cable Service Interface Specification
DoS	Denial of Service
DDoS	Distributed Denial of Service
DSL	Digital Subscriber Line
DSU	Data Service Unit
DWDM	Dense Wavelength Division Multiplexing
E1	E-Carrier Level 1
EAP	Extensible Authentication Protocol
EDNS	Extension Mechanisms for DNS
EGP	Exterior Gateway Protocol
EIGRP	Enhanced Interior Gateway Routing Protocol
EMI	Electromagnetic Interference
ESD	Electrostatic Discharge
ESSID	Extended Service Set Identifier
ESP	Encapsulated Security Packets
FDDI	Fiber Distributed Data Interface
FDM	Frequency Division Multiplexing
FHSS	Frequency Hopping Spread Spectrum
FM	Frequency Modulation
FQDN	Fully Qualified Domain Name / Fully Qualified Distinguished Name
FTP	File Transfer Protocol
GBIC	Gigabit Interface Converter
Gbps	Giga bits per second
PGP	GNU Privacy Guard
HDLC	High-Level Data Link Control
HSRP	Hot Standby Router Protocol
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
Hz	Hertz
IANA	Internet Assigned Numbers Authority
ICA	Independent Computer Architecture

ACRONYM	STANDS FOR
ICANN	Internet Corporation for Assigned Names and Numbers
ICMP	Internet Control Message Protocol
ICS	Internet Connection Sharing
IDF	Intermediate Distribution Frame
IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers
IGMP	Internet Group Multicast Protocol
IGP	Interior Gateway Protocol
IIS	Internet Information Services
IKE	Internet Key Exchange
IMAP4	Internet Message Access Protocol version 4
InterNIC	Internet Network Information Center
IP	Internet Protocol
IPS	Intrusion Prevention System
IPSec	Internet Protocol Security
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
ISAKMP	Internet Security Association and Key Management Protocol
ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
IT	Information Technology
IV	Initialization Vector
Kbps	Kilobits per second
L2F	Layer 2 Forwarding
L2TP	Layer 2 Tunneling Protocol
LACP	Link Aggregation Control Protocol
LAN	Local Area Network
LC	Local Connector
LDAP	Lightweight Directory Access Protocol
LEC	Local Exchange Carrier
LED	Light Emitting Diode
LLC	Logical Link Control
MAC	Media Access Control / Medium Access Control
Mbps	Megabits per second

ACRONYM	STANDS FOR
Mbps	Megabytes per second
MDF	Main Distribution Frame
MDI	Media Dependent Interface
MDIX	Media Dependent Interface Crossover
MIB	Management Information Base
MMF	Multimode Fiber
MPLS	MultiProtocol Label Switching
MS-CHAP	Microsoft Challenge Handshake Authentication Protocol
MT-RJ	Mechanical Transfer-Registered Jack
MX	Mail Exchanger
NAC	Network Access Control
NaaS	Network as a Service
NAS	Network Attached Storage
NAT	Network Address Translation
NCP	Network Control Protocol
NetBEUI	Network Basic Input / Output Extended User Interface
NetBIOS	Network Basic Input / Output System
NFS	Network File Service
NIC	Network Interface Card
NIPS	Network Intrusion Prevention System
nm	Nanometer
NNTP	Network News Transport Protocol
NTP	Network Time Protocol
NWLINK	Microsoft IPX/SPX Protocol
OCx	Optical Carrier
OS	Operating Systems
OSI	Open Systems Interconnect
OSPF	Open Shortest Path First
OTDR	Optical Time Domain Reflectometer
PAP	Password Authentication Protocol
PAT	Port Address Translation
PC	Personal Computer
PGP	Pretty Good Privacy
PKI	Public Key Infrastructure

ACRONYM	STANDS FOR
PoE	Power over Ethernet
POP3	Post Office Protocol version 3
POTS	Plain Old Telephone System
PPP	Point-to-Point Protocol
PPPoE	Point-to-Point Protocol over Ethernet
PPTP	Point-to-Point Tunneling Protocol
PRI	Primary Rate Interface
PSTN	Public Switched Telephone Network
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
RARP	Reverse Address Resolution Protocol
RAS	Remote Access Service
RDP	Remote Desktop Protocol
RFI	Radio Frequency Interface
RG	Radio Guide
RIP	Routing Internet Protocol
RJ	Registered Jack
RSA	Rivest, Shamir, Adelman
RSH	Remote Shell
RTP	Real Time Protocol
RTSP	Real Time Streaming Protocol
RTT	Round Trip Time or Real Transfer Time
SA	Security Association
SC	Standard Connector / Subscriber Connector
SCP	Secure Copy Protocol
SDSL	Symmetrical Digital Subscriber Line
SFTP	Secure File Transfer Protocol
SFP	Small Form-Factor Pluggable
SIP	Session Initiation Protocol
SLIP	Serial Line Internet Protocol
SMF	Single Mode Fiber
SMTP	Simple Mail Transfer Protocol
SNAT	Static Network Address Translation
SNMP	Simple Network Management Protocol
SNTP	Simple Network Time Protocol

ACRONYM	STANDS FOR
SOA	Start of Authority
SOHO	Small Office / Home Office
SONET	Synchronous Optical Network
SPS	Standby Power Supply
SSH	Secure Shell
SSID	Service Set Identifier
SSL	Secure Sockets Layer
ST	Straight Tip or Snap Twist
STP	Shielded Twisted Pair
T1	T-Carrier Level 1
TA	Terminal Adaptor
TACACS+	Terminal Access Control Access Control System+
TCP	Transmission Control Protocol
TCP / IP	Transmission Control Protocol / Internet Protocol
TDM	Time Division Multiplexing
TDR	Time Domain Reflectometer
Telco	Telephone Company
TFTP	Trivial File Transfer Protocol
TKIP	Temporal Key Integrity Protocol
TLS	Transport Layer Security
TTL	Time to Live
UDP	User Datagram Protocol
UNC	Universal Naming Convention
UPS	Uninterruptible Power Supply
URL	Uniform Resource Locator
USB	Universal Serial Bus
UTP	Unshielded Twisted Pair
VDSL	Variable Digital Subscriber Line
VLAN	Virtual Local Area Network
VNC	Virtual Network Connection
VoIP	Voice over IP
VPN	Virtual Private Network
VTP	Virtual Trunk Protocol
WAN	Wide Area Network
WAP	Wireless Application Protocol / Wireless Access Point

ACRONYM	STANDS FOR
WEP	Wired Equivalent Privacy
WINS	Window Internet Name Service
WPA	Wi-Fi Protected Access
www	World Wide Web
X.25	CCITT Packet Switching Protocol
XML	eXtensible Markup Language
XDSL	Extended Digital Subscriber Line
Zeroconf	Zero Configuration

Index

Symbols

- 3G technology, 141
- 4G technology, 141
- 6P4C RJ-14 connectors, 123
- 8P8C connectors, 122
- 10Base2 specification, 160
- 10Base5, 160
- 10Base5 specification, 160
- 10Base-FL specification, 160
- 10Base-T networks, 114
- 10Base-T specification, 160
- 10GBase-ER specification, 164
- 10GBase-EW specification, 164
- 10GBase-LR specification, 164
- 10GBase-LW specification, 164
- 10GBase-SR specification, 164
- 10GBase-SW specification, 164
- 10GBase-T specification, 164
- 100Base-FX specification, 161
- 100Base-T4 specification, 161
- 100Base-T specification, 161
- 100Base-TX specification, 161
- 110 blocks, 123
- 802.3 Ethernet standards, 78
- 802.11ac standard (wireless networks), 76
- 802.11a standards
 - 802.11b standards, compatibility between, 80
- 802.11a standard (wireless networks)
 - Orthogonal Frequency-Division Multiplexing (OFDM) and, 80
- 802.11b/g standards (wireless networks), 76
- 802.11b standards
 - 802.11a standards, compatibility between, 80
- 802.11g networks
 - OFDM modulation, 77
- 802.11n standard
 - 2x2 rated devices, speed of, 79
- 802.11n standard (wireless networks), 76
- 802.11 protocols, 240
- 802.11 standards, 78
- 1000Base-CX specification, 162
- 1000Base-LX specification, 162
- 1000Base-SX specification, 162
- 1000Base-T Gigabit Ethernet networks, 114

1000Base-T specification, 162

–ERR response code, 41

+OK response code, 41

A

AAA, 267–269

AAAA records, 46

AAA authentication

RADIUS and, 267–269

TACACS+ and, 269

access control

access control lists for, 245–246

remote access for, 256–261

tunneling/encryption for, 246–255

access control list (ACL)

MAC addresses and, 69

access control lists, 245–246

accessibility

to documentation, 227

access points. *See APs*

adding, to decrease latency, 90

antennae of, 74

basic service set (BSS) offered by, 75

incorrect placement of,

troubleshooting, 92

location of, choosing, 74

security concerns with, 74

signal bouncing and, 91

SSIDs and connecting devices to, 80

access security, 244–262

exam need to know, 245

acknowledgment (ACK) frame, 168

ACLS. *See access control lists*

Active Directory

NTP and, 41

Active Directory Domain Services (AD

DS), 46, 265

Active Directory Domain Services network

DNS and, 102

active optical network (AON), 140

Address Resolution Protocol. *See ARP*

Address Resolution Protocol (ARP), 7

APIPA and, 19

ad hoc wireless networks, 75

ADSL2, 137

ADSL2+, 137

ADSL Lite, 137
 ADSL Termination Unit-Remote (ATU-R), 137
 Advanced Encryption System (AES), 241
 AH. *See IP Authentication Header*
 air handling, 117
 Alien Crosstalk (AXT), 158
 alligator clips, 189
 American National Standards Institute (ANSI), 131
 analyzers
 sniffers vs., 222
 ANSI/TIA-568-C standard, 158, 171, 187
 ANSI/TIA-568-C standards, 196
 antennae
 beam, 74
 directional, 74
 gain, 74
 installing, 73, 74–75
 patch, 74
 replacement, 75
 anti-replay
 IPsec and, 251
 anycasts, 19
 APIPA. *See Automatic Private IP Addressing*
 appliances
 wireless networks and, 75, 87
 appliances and methods, security, 294–298
 exam need to know, 294
 honeypots/honeynets, categorization of, 298
 IDSes vs. IPSes, 294–296
 vulnerability scanners, categorization of, 296–298
 appliances, network. *See network appliances*
 Application layer (OSI model)
 visualization of layers, 2
 Application layer (TCP/IP model)
 visualization of layers, 3
 APs, 242, 243
 A records, 46
 ARP, 207–208
 networking protocols, function of, 39
 TCP/IP suite, 35
 Arp.exe, 207
 ARP messages
 duplicate IP addresses, troubleshooting, 101–102
 asset management, purpose of, 228
 asymmetrical services, 142
 Asymmetric Digital Subscriber Line (ADSL), 136
 Asynchronous Transfer Mode (ATM), 133–134
 attacks. *See threats and attacks*
 attenuation, 112, 158
 Augmented Category 6 (CAT6A) cables, 115
 authentication
 defined, 262
 PPP, 257
 PTP tunnels and, 248
 SSL and, 254
 Authentication, Authorization, and Accounting. *See AAA*
 Authentication phase (PPP), 258–259
 automatic allocation of IP addresses, 82
 automatic private IP addressing, 19–20
 Automatic Private IP Addressing (APIPA), 84
 autonomous system (AS), 30
 autonomous system numbers, 30
 awareness
 attack mitigation via, 279

B

backbone, 173
 backbone cabling, 172
 backups, 233
 bandwidth
 EIGRP, 25
 performance optimization and, 235
 Quality of Service and, 231
 throughput vs., 200
 baselines, purpose of, 228–229
 Basic Rate Interface (BRI), 135
 basic service set (BSS), 75
 basic switches, 10
 B channels, 135
 beam antenna, 74
 Berkeley Software Distribution (BSD), 288
 BNC connectors, 124
 bonding, 169
 Bootstrap Protocol (BOOTP)
 DHCP vs., 37
 Border Gateway Protocol (BGP), 30
 bridges
 collision domains and, 29
 OSI model, 12–13
 OSI model and, 6
 troubleshooting bridging loops, 95
 broadband over power lines (BPL), 119
 broadband, powerline networking
 vs., 119–120
 broadband routers, 105

broadcast domain,, 19
 broadcast domains, 29–30
 broadcasts
 IPv4 addresses and, 19
 broadcast traffic
 RIP version 1 and, 25
 routers and, 70
 routers and reduction of, 10
 switches and, 70
 broadcast transmissions
 multicast vs., 44
 buffer overflow, 278–279
 building entrance, 172
 bus topologies, 153–155
 bus topology, 153–154, 153–155
 Coaxial Ethernet networks, 12
 switches in, 10
 butt set, 189–190

C

cable, 113
 copper. *See* copper media
 fiber optic, 112–113
 cable certifier, 187–188
 cable certifiers, 158
 loopback plug and, 195
 cable management, purpose of, 227
 cables
 internal runs, troubleshooting, 95
 link lights, unreliability of, 95
 maximum lengths of, 104
 OSI model, 12
 OSI model and, 5
 planning for, in SOHO networks, 104–105
 troubleshooting, 95–96
 wire scheme and, 224
 cable television (CATV) networks, 138–139
 cable tester, 186–187
 cable testing devices, 187
 Cache Array Routing Protocol. *See* CARP
 caching
 by NAT routers, 183
 caching engines, 233–234
 caching-only servers, 46
 CARP, 234
 carrier sense, 165
 Carrier Sense Multiple Access with Collision Detection (CSMA/CD), 165
 CAT3 cable, 118
 CAT5, 161
 CAT5 cable, 118
 CAT5e cable, 118
 CAT6a cable, 164
 CAT6 cable, 118
 Category 3 (CAT3) cables, 114
 Category 5 (CAT5) cables, 114
 Category 5e (CAT5e) cables, 114
 Category 6 (CAT6) cables, 114
 ceilings, dropped, 117
 cells, 133
 cellular technologies (for Internet access), 141–142
 centralized authentication, 273
 Challenge Handshake Authentication Protocol (CHAP), 257
 change management, purpose of, 230
 channel service unit/data service unit (CSU/DSU), 130
 channels for wireless networks
 automatic configuration of, in devices, 77
 configuring, 76–78
 DSSS modulation, 76
 OFDM modulation, 77
 overcrowding, troubleshooting, 90
 troubleshooting problems with, 90
 CHAP authentication, 271
 CIDR. *See* classless inter-domain routing
 CIDR (classless inter-domain routing)
 IP addressing and, 16
 circuit-switching, 142–143
 circuit-switching networks, 143
 Cisco Systems, 26, 249
 port mirroring and, 72
 Switched Port Analyzer (SPAN), 72
 VTP (proprietary system), 71
 Citrix Corporation, 260
 cladding, 112
 Class A (IP addresses), 14
 Class B (IP addresses), 14
 Class C (IP addresses), 14
 Class D (IP addresses), 15
 Class E (IP addresses), 15
 classes
 IP addressing, 14–16
 obsolete IP addressing, 13
 classful IP addressing, 25
 Classless Inter-Domain Routing (CIDR), 13
 clear-to-send (CTS) messages, 168
 clients
 Kerberos encryption and, 265
 client/server networking, 155

client/server topology, 155
client-to-site connection, 248
clusters, 182, 232
CNAME records, 47
coaxial cable, 115–126, 118, 143
Coaxial Ethernet
 bus topology with, 12
collision detection phase, 166–167
collision domains, 169
collisions, 166
colon
 double, 17
colon-hexadecimal format, 17
committed burst information rate
 (CBIR), 132
committed information rate (CIR), 132
computers
 asset management and, 228
configuration issues (wireless networks)
 power saving mechanisms and, 89
 troubleshooting, 88–89
configuration management
 documentation, 223–230
 asset management, purpose of, 228
 baselines, purpose of, 228–229
 cable management, purpose of, 227
 change management, purpose of, 230
 exam need to know, 223
 network maps, 225–226
 purpose of documentation, 227
 wire schemes, 223–225
Configure Ack message, 258
Configure Request message, 258
congestion delays
 OSPF, 26
connectivity
 using hardware tools to
 troubleshoot, 186–198
 using software tools for
 troubleshooting, 198–215
connectivity problems
 exam need to know, 157
 troubleshooting, 157–159
connectivity software
 software tools, 200
connectors
 copper, 122–125
 exam need to know, 121
 fiber, 121–122
 standard, 121–125
Content Advisor feature (Internet Explorer), 185

content filters, 184–185
content switches, 182
continuity test, 157, 190
convergence, 27
converters, media. *See* media converters
copper cable testing
 optic cable testing vs., 196
copper connectors, 122–125
copper media, 113–116
 fiber media vs., 113–116
crimper, 188–189
crimping, 115
crossover cable, 146
crossover cables, 115
crosstalk, 116, 119, 157, 187
CSMA/CA, 167–169
CSMA/CD, 162
CSU/DSU, 176
curing oven, 122
customer premises equipment (CPE), 175

D

DARPA commands, 40
data centers
 central power supplies for, 98
datagrams, 8
 defined, 37
data integrity
 SSL and, 254
data-link layer
 bridges in, 12
 frames in, 8
 MAC addresses in, 6
data-link layer device, 5
Data-Link layer (OSI model)
 visualization of layers, 2
DB-9 connectors, 125
DB9 connectors, 125
DB loss, 158
D channels, 135
DE9 connectors, 125
default gateway, 22
default ports, identifying TCP and
 UDP, 31–33
delay
 EIGRP, 25
delay skew, 158
demarc, 175
demarc distribution, 175–176
demilitarized zone. *See* DMZ

denial of service (DoS) attacks, 276
 Denial of Service (DoS) attacks
 NAT and, 66
 dense wave division multiplexing
 (DWDM), 132
 Department of Defense (DoD) model.
 See TCP/IP model
 Description (protocol analyzer Frame
 Summary pane), 194
 desktops
 virtual, 54–55
 desktop virtualization, 54
 Destination Address field
 MAC addresses and, 6
 Destination (protocol analyzer Frame
 Summary pane), 194
 device placement
 wireless security and, 243–244
 devices
 limitations on, in SOHO networks, 106
 DHCP, 81–86
 automatic allocation of IP addresses, 82
 Automatic Private IP Addressing
 (APIPA), 84
 DHCP Message Type options, 85
 duplicate IP addresses and, 101
 dynamic allocation of IP addresses, 82
 exam need to know, 81–82
 Lease Duration settings, 83
 leases, 83–84
 leasing IPs across subnets with, 84
 manual allocation of IP addresses, 82
 Microsoft DHCP Server service, 83
 networking protocols, function of, 37
 options, 85–86
 renewing leases, 84
 reservations, 83
 scopes, 83
 subnets and, 83
 subnets, behavior when changing, 84
 zero-configuration networking, 84
 DHCPDISCOVER messages, 83
 DHCP Message Type option, 85
 DHCPOFFER messages, 83
 DHCPREQUEST unicast messages, 84
 diagnostics
 installing/configuring, 62
 differential backups, 233
 Differentiated services (DiffServ), 71, 231
 Digital Signal 0s (DS0s), 129
 Digital Signal 1 (DS1), 129
 digital signatures, 263
 Digital Subscriber Line Access Multiplexer
 (DSLAM), 137
 Digital Subscriber Line (DSL), 136–138
 dig utility, 203–204
 Dijkstra algorithm, 26
 Directional antennae, 74
 Direct-Sequence Spread Spectrum
 (DSSS), 126
 diskless workstation
 TFTP and UDP in startup of, 38
 distance
 DSL service and, 137
 Ethernet LANs, 170–171
 distance vector protocols, 24
 distance vector routing
 link state vs., 24
 distance vector routing protocols
 hop count used by, 22
 distance, WAN technologies classified
 by, 144
 distributed DoS attacks, 276
 DIX Ethernet II, 160
 DMZ
 firewalls and configuration of, 293–294
 DMZs, 246
 DNS
 Active Directory Domain Services
 network and, 102
 concepts and components,
 summarizing, 45–48
 dynamic, 47–48
 networking protocols, function of, 38
 DNS namespace, 46
 DNS records, 46–47
 DNS server
 round robin DNS and, 182
 DNS servers, 45–46
 documentation
 troubleshooting, 53
 documentation, configuration management.
 See configuration management
 documentation
 Domain Name Service (DNS)
 incorrect addresses,
 troubleshooting, 102
 Domain Name System. See DNS
 Domain Name System (DNS), 32
 dotted decimal notation, 16
 double colon, 17
 dropped ceilings, 117

- DS0, 130
- DS1, 130
- DS2, 130
- DS3, 130
- DS4, 130
- DSSS modulation, 76
- duplex mode (of ports), 96
- dynamic allocation of IP addresses, 82
- dynamic DNS, 47–48
- Dynamic Host Configuration Protocol.
 - See* DHCP
- Dynamic Host Configuration Protocol (DHCP), 32, 205. *See* DHCP
- dynamic IP addressing
 - static vs., 82
- Dynamic NAT, 291
- dynamic routing
 - convergence and, 27
 - static vs., 22
- dynamic routing protocols, 22
 - convergence issues with, troubleshooting, 99
- Dynamic Updates, 48

E

- E-1, 129
- E-3, 129
- EAP authentication, 272
- EAP-Transport Layer Security (EAP-TLS), 272
- E-carrier service, 129
- Echo Reply messages, 201
- Echo Request messages, 200
- EGP
 - IGP vs., 30
- EIGRP, 25–26
- electromagnetic interference (EMI)
 - fiber media and, 112
- Electronic Industries Alliance (EIA), 125
- EMI
 - coaxial cable and, 115
 - STP cable and, 115
- encapsulation, 247
- encryption. *See also* tunneling
 - IMAP4 and, 42
 - IPsec and, 250
 - secret key, 263
 - SSL and, 254
- encryption devices
 - OSI model, 11–12
 - OSI model and, 5
- encryption protocols, 240–242
 - WEP, 240–241
 - WPA, 241–242
- encryption type problems
 - troubleshooting, 91
- entrance, building, 172
- environmental monitors, 197–198
- ephemeral port numbers, 33
- Equal-Level Far-End Crosstalk (EL-FEXT), 119, 158
- equipment rooms, 172
- erbium-doped fiber amplifiers (EDFAs), 132
- error detection
 - data-link layer frames and, 8
- ESP. *See* IP Encapsulating Security Payload
- Ethernet, 160–161
 - IP vs., 7
 - modern, 167
 - speeds, 118
- Ethernet devices
 - MAC addresses used by, 6
- Ethernet frames
 - trunk ports/VLANs and, 67
- Ethernet LAN
 - Gateway values and, 63
 - port speed and duplex mode as functions of, 96
 - routing on, 21
- Ethernet network interface adapters, 12
- EUI-64
 - OSI model, 7
- EUI-64 relates
 - OSI model and, 5
- Europe, 129
- Event Viewer console, 220, 221
- evil twins, 275
- Evolved HSPA, 141
- exam need to know
 - access security, 245
 - appliances and methods, security, 294
 - configuration management
 - documentation, 223
 - connectivity problems, 157
 - connector types, 121
 - DHCP, 81–82
 - DNS concepts, 45
 - firewalls, 281
 - hardware tools, 186
 - IP addressing, 13–14
 - LANs, 159–160
 - network appliances, 181–182

- networking protocols, function of, 33–34
- Network Installation and Configuration domain, 62
- network media, 111–112
- network monitoring, 216
- OSI model, 2, 5–6
- performance optimization, 231
- routers, troubleshooting, 93–94
- routing and switching, 20–21
- software tools, 199
- SOHO networks, building, 103
- switches, troubleshooting, 93–94
- TCP and UDP default ports, identifying, 31–32
- TCP/IP model, 2
- threats and attacks, 274
- topologies, 145
- troubleshooting, 186, 199
- user authentication, 262–263
- virtual network components, identifying, 54
- WAN technology types/properties, categorization of network media by, 129
- wireless networks, installing and configuring, 73
- wireless networks, troubleshooting, 87
- wireless security, 239
- wireless standards for network media, 126
- wiring distribution, components of, 172
- exclusive filtering, 286
- extended TACACS (XTACACS), 269
- extended unique identifier-64. *See* EUI-64
- Extensible Authentication Protocol (EAP), 270
- exterior gateway protocols, 30. *See* EGP

F

- fallback speeds, 128
- Far-End Crosstalk (FEXT), 158
- Fast Ethernet, 161–162
- fast link pulse (FLP), 170
- fault, open, 191
- fault tolerance, 232–233
- F connectors, 124
- female connectors, 123
- fiber connectors, 121–122
- fiber media, 112–113
- fiber optic cable, 112–113, 119, 143
- fiber optic cables
 - troubleshooting, 196–197
- Fiber-to-the-curb (FTTC), 139
- Fiber-to-the-desk (FTTD), 139
- Fiber-to-the-home (FTTH), 139
- Fiber-to-the-node (FTTN), 139
- Fiber-to-the-premises (FTTP), 139
- File Transfer Protocol. *See* FTP
- File Transfer Protocol (FTP), 32
- filtering
 - access control for, 245
 - by NAT routers, 183
- filters
 - content, 184–185
- fire, 117
- firewalls, 281–294
 - content filters vs., 184
 - defined, 282
 - DMZ configuration and, 293–294
 - DoS attacks and, 279
 - exam need to know, 281
 - hardware vs. software, 282
 - ICMP Echo Requests and, 98
 - installing, 282
 - NAT/PAT configuration and, 290–293
 - NAT vs., for internet security, 66
 - Ping and, 201
 - port security configuration and, 282–283
 - rules, configuration of, 286–290
 - social engineering and, 279
 - stateful inspection vs. packet filtering and, 283–285
- floors, raised, 117
- FOIRL specification, 160
- formatting
 - IPv4 vs. IPv6, 14
- forwarders, 46
- frame aggregation, 127
- frame check sequence (FCS) field (data-link layer), 8
- Frame Details pane (protocol analyzer), 194
- Frame Number (protocol analyzer Frame Summary pane), 193
- frame relay, 132–133
- frame relay assembler/disassembler (FRAD), 133
- frames
 - OSI model, 8
 - OSI model and, 5
 - packets vs., 8
- Frame Summary pane (protocol analyzer), 193

- FreeRADIUS, 269
- freqencies of wireless networks
 - 802.11ac standard, 76
 - 802.11b/g standards, 76
 - 802.11n standard, 76
 - large office buildings and, 76
 - OFDM modulation, 77
 - selecting, 76
- Frequency-Hopping Spread Spectrum (FHSS), 126
- FTP
 - DARPA commands and, 40
 - networking protocols, function of, 37
 - TCP/IP suite, 35
- FTP bounce attacks, 277
- FTP client applications
 - need for, 37
- FTP protocol
 - FTP bounce and, 277
- Fully-Qualified Domain Names (FQNDs), 200

G

- gain (signal strength), 74
- gateway
 - default, 22
- Gateway column (IPv4 routing tables), 63
- gateways
 - troubleshooting, 101
 - values and Ethernet LANs, 63
- general logs
 - network monitoring using, 221
- Generic Routing Encapsulation. *See GRE*
- GetBulkRequest, 218
- GetNextRequest, 218
- Gigabit Ethernet, 162–163, 170
- Gigabit Interface Converter (GBIC) modular standard, 100
- gigabits per second (Gbps), 200
- GRE, 248

H

- hard disk backups, 233
- hard disk failure, 232
- hardware
 - asset management and, 228
 - butt set, 189–190
 - cable certifier, 187–188
 - cable tester, 186–187
 - crimper, 188–189
 - environmental monitors, 197–198
- loopback plug, 195–196
- multimeter, 197
- network diagrams, location in, 225
- OTDR, 196–197
- protocol analyzer, 192–195
- punch down tool, 191–192
- TDR, 196
- toner probe, 190–191
- troubleshooting, 51
- hardware addresses. *See MAC addresses*
- hardware firewalls, 282
- hardware tools
 - connectivity troubleshooting with, 186–198
 - exam need to know, 186
- hash message authentication code (HMAC), 251
- Hex Details pane (protocol analyzer), 194
- high availability, 232
- High-bit-rate Digital Subscriber Line (HDSL), 136
- High Speed Packet Access (HSPA), 141
- history logs
 - network monitoring using, 221
- honeynets, 298
- honeypots, 298
- hop count, 22
 - EIGRP, 25
 - OSPF, 26
- horizontal cabling, 172
- host identifier
 - boundary between network identifier and, 18
 - IPv4 addresses, 14
 - value assignment and, 15
- HSPA+, 141
- HTTP
 - HTTPS port vs., 39
 - networking protocols, function of, 38–39
 - TCP/IP suite, 35
- HTTPS
 - networking protocols, function of, 39
 - TLS and, 44
- hubs
 - OSI model, 11
 - OSI model and, 5
 - switches vs., 9, 167, 193
- hybrid fiber coaxial (HFC) networks, 138
- hybrid topology, 156
- Hypertext Transfer Protocol. *See HTTP*
- Hypertext Transfer Protocol (HTTP), 32

Hypertext Transfer Protocol Secure.
See [HTTPS](#)

Hypertext Transfer Protocol Secure (HTTPS), 32

Hyper-V, 54

hypervisor, 55

I

ICMP

- DoS attacks and, 276
- networking protocols, function of, 43
- TCP/IP suite, 35

IDF (intermediate distribution frames), 172–173

IDS, 294–296

IEEE 802.1q standard, 28

IEEE 802.1X, 241

IEEE 802.1X standard, 270

IEEE 802.3 standard, 160

IEEE 802.11a standard, 126

IEEE 802.11b standard, 127

IEEE 802.11g standard, 80, 127

IEEE 802.11n standard, 79, 80, 127–128

IEEE 802.11 standard, 29, 126

ifconfig, 204–206

IGMP

- networking protocols, function of, 44
- TCP/IP suite, 35

IGP

- EGP vs., 30

IMAP

- TCP/IP suite, 35

IMAP4

- networking protocols, function of, 42

incident response

- attack mitigation via, 280

inclusive filtering, 286

incoming signals

- hubs and, 11

incremental backups, 233

Independent Computing Architecture (ICA), 260

indirect route, 21

infrastructure networks, 75

installation

- questions to ask about, 224

installation/configuration, 61–110

- Dynamic Host Configuration Protocol (DHCP), 81–86
- exam need to know, 62

interfaces on routers/switches, 68

internal diagnostics on switches/routers, 70

managed/unmanaged routers and switches, 68

Network Address Translation (NAT), 64–65

Power over Ethernet (PoE), 69

Quality of Service (QoS), 71–72

routers/switches, 61–72

routing tables, 62–64

small office/home office (SOHO)

- networks, 102–107

traffic filtering, 69–70

troubleshooting routers/switches, 93–102

troubleshooting wireless networks, 86–93

VLAN Trunking Protocol (VTP), 70

wireless networks, 73–81

Institute of Electrical and Electronics Engineers (IEEE), 17, 78, 240

Integrated Services Digital Network (ISDN), 135–136

Integrated services (IntServ), 71, 231

integrated SNMP agent

- vs. managed switch/router, 68

integrity

- IPsec and, 251

Interface column (routing tables), 63

- router's use of, 64

Interface List (IPv6 routing tables), 64

interference on wireless networks

- appliances and, 75, 87
- between wireless channels, 76–78
- preventing, 75
- troubleshooting, 87–88
- walls and, 75, 88

interior gateway protocols. *See* [IGPs](#)

interior gateway routing

- protocols used for, 30

Interior Gateway Routing Protocol (IGRP), 26

intermediate distribution frames (IDFs), 173

internal diagnostics of routers/switches, 70

International Organization for Standardization (ISO), 58

International Telecommunications Union (ITU), 131, 141

Internet

- as backup medium, 233

Internet access, 134–142

Internet Assigned Numbers Authority (IANA), 31, 290

Internet Authentication Server (IAS), 269

Internet Connection Sharing (ICS), 292

Internet Control Message Protocol. *See* ICMP

Internet Control Message Protocol (ICMP), 200

 packets and, 9

Internet Control Message Protocol (ICMP)

 Echo Request messages, 98

Internet Digital Subscriber Line (IDSL), 137

Internet Engineering Task Force (IETF), 58, 217

 IPsec and, 250

Internet Explorer

 Content Advisor feature of, 185

Internet Group Management Protocol.

See IGMP

Internet Group Management Protocol (IGMP), 19

Internet Key Exchange (IKE), 253

Internet layer (TCP/IP model)

 visualization of layers, 3

Internet Mail Access Protocol (IMAP), 32

Internet Message Access Protocol. *See* IMAP

Internet Protocol addresses. *See* IP addresses

Internet Protocol Control Protocol (IPCP), 259

Internet Security Association and Key Management Protocol (ISAKMP), 253

Internet service providers (ISPs)

 addresses registered with, 15

Internet transmissions

 IP addresses and, 6

 MAC addresses and, 6

internetwork

 link state routing protocols on, 24

intrusion detection system. *See* IDS

intrusion prevention systems. *See* IPS

inventory scanners, 228

IP

 datagrams of, 37

 TCP/IP suite, 35

 versions of, 7

IP addresses

 automatic allocation, 82

 broadcasts and, 19

 duplicate, troubleshooting, 101–102

 dynamic allocation of, 82

 Internet transmissions and, 6

 load balancing and, 182

 MAC addresses, reconciling, 39

 manual allocation, 82

 multicasts and, 19

 NAT and, 290

 NAT, using to protect internal, 64–65

 OSI model, 6–7

 OSI model and, 5

 packet filtering and, 284

 Ping and, 200

 registered vs. unregistered and the Internet, 64

 static vs. dynamic, 82

 subnetting, 18

 unicasts and, 19

 VLAN assignment issues and, 97

IP address filtering, 245, 246

IP addressing, 13–20

 automatic private, 19–20

 CIDR and, 16

 classes, 14–16

 DHCP methods, 37

 exam need to know, 13–14

 MAC address formatting, 17–18

 obsolete classes, 13

 version 4 vs. version 6, 16–17

IP address space, conserving with DHCP, 82

IP Authentication Header (AH), 251

Ipconfig, 204–206

Ipconfig.exe, 17

IPCP (the Internet Protocol Control Protocol), 257

IP datagrams

 IPsec and, 250

IP Encapsulating Security Payload (ESP), 251–252

ipfw, 288

IP protocol

 IPsec and, 250

IPS, 294–296

IPsec, 250

 key exchange protocols and, 252–253

 transport/tunnel modes and, 253–254

IPsec protocols, 251–254

IP spoofing, 284

IPv4

 addressing in IPv6 vs., 16–17

 APIPA and, 20

 formatting in, 14

 IPv6 vs., 23

IPv6

 addressing in IPv4 vs., 16–17

 APIPA and, 20

 formatting in, 14

 IPv4 vs., 23

IPv6 addresses
 EUI-64 in, 7

IPv6 networks
 Interface List (routing table), 64
 MTU mismatches and, 98
 subnet masks and, 64

iterative queries, 46

J

J-1, 129
J-3, 129
Japan, 129
J-carrier service, 129

K

Kerberos, 265–267
Key Distribution Center (KDC), 265
key exchange protocols, 252–253
keys, encryption, 240
keystone connectors, 191, 192
kilobits per second (Kbps), 200

L

L2F, 249
L2TP, 249
LAN/MAN Standards Committee, IEEE
 802, 78
LAN protocol
 VPNs and, 247
LAN protocols
 MAC addresses and, 6
LANs
 hubs vs. switches and efficiency of, 11
LANs (local area networks), 159–171
 comparing types of, 160–165
 Ethernet, 160–161
 exam need to know, 159–160
 Fast Ethernet, 161–162
 Gigabit Ethernet, 162–163
 properties of, 165–171
 10 Gigabit Ethernet, 163–165
last mile fiber, 139–140
latency, 23
latency sensitivity, 234
latency (wireless networks)
 troubleshooting, 90–91
Layer 2 Forwarding. *See* L2F
layer 4 switches, 182
layer 7 switches, 182

LC connectors, 122
leased lines, 129–130
length, categorization of cables by, 118–119
level-2 memory cache, 234
link aggregation
 trunking vs., 68
Link Control Protocol (LCP), 257
Link Dead phase (PPP), 258
Link Establishment phase (PPP), 258
Link layer (TCP/IP model)
 visualization of layers, 3
link lights, unreliability of in troubleshooting
 cables, 95
Link Open phase (PPP), 259
link quality monitoring
 PPP, 259
link state routing
 distance vector vs., 24
link state routing protocols, 99
Link Termination phase (PPP), 259
Linux
 NAT and, 292
Linux systems
 dig with, 203
 syslogs with, 219
 traceroute on, 202
load
 EIGRP, 25
load balancer, 182–183
load balancing, 182
load-balancing switches, 182
local area networking protocols. *See* LAN
 protocols
logging
 by NAT routers, 183
Long Term Evolution (LTE), 141
loopback plug, 195–196

M

MAC address
 NIC, 12

MAC addresses
 access control list (ACL) and, 69
 Ethernet frames and, 8
 format of, 14
 formatting, 17–18
 IP addresses, reconciling, 39
 OSI model, 6
 OSI model and, 5
 packet filtering and, 284
MAC address filtering, 242–243, 245, 246

- macros
 - as malware, 278
- magnetic tape, 233
- main distribution frame (MDF), 174
- male connectors, 123
- malware, 278
- management information base (MIB), 43, 217
- man in the middle (MITM) attacks, 277
- manual allocation of IP addresses, 82
- maps, network, 225–226
- masquerading, 291. *See also* PAT
- Massachusetts Institute of Technology, 265
- maximum transfer unit (MTU)
 - EIGRP, 25
 - troubleshooting, 97–98
- maximum transmission unit (MTU), 23
- MDF (main distribution frame), 173–175
- media access control addresses. *See* MAC addresses
- media access control (MAC), 117, 242
- media converters, 117–118
- media, network. *See* network media
- median testers, 187
- media testers, 158
- megabits per second (Mbps), 200
- mesh shielding, 115
- mesh topology, 152–153
- Message Digest 5 Challenge (MD5-Challenge), 272
- Message Digest 5 (MD5), 271
- methods and appliances, security, 294–298
- Metric column (routing tables), 63
 - route performance and, 63
 - router's use of, 64
- metrics
 - routing, 22–23
- metropolitan area network (MAN), 141
- Microsoft, 155, 260, 271, 272
- Microsoft DHCP Server service, 83
- Microsoft Management Console (MMC), 288
- Microsoft Windows Server, 267
- mirroring
 - port, 29
- mitigation techniques, attack, 279–280
- modems, 134
- monitoring, network. *See* network monitoring
- monitoring resources
 - analyzing traffic with, 215
- monitoring tools
 - traffic analysis with, 215–223
- MPLS, 145–146
- MS-CHAP authentication, 271–272
- MS-CHAP v1, 271
- MS-CHAP v2, 271
- MT-RJ connectors, 122
- MTU black holes, 98
- MTU mismatches
 - IPv6 networks and, 98
 - MTU black holes, 98
 - Path MTU Discovery (ICMP), 98
 - Ping utility and, 98
 - troubleshooting, 97–98
- multicasts
 - IPv4 addresses and, 19
- multicast transmissions, 44
 - broadcast vs., 44
- multifactor authentication, 273
- multilayer switches, 182
 - OSI model, 10–11
 - OSI model and, 5
- multimeter, 197
- multimode fiber optic cable, 113, 119
- multiple access phase, 166
- Multiple-Input Multiple-Output (MIMO), 79, 127, 128
- Multiple VLAN Registration Protocol (MVRP), 71
- multiprotocol label switching (MPLS), 145–146
- Multitrate Symmetric Digital Subscriber Line (MSDSL), 136
- multistation access unit (MAU), 149
- mutual authentication
 - IPSec and, 251
 - MITM attacks and, 277
- MX records, 46
- MX resource records
 - SMTP servers and, 42

N

- NAC. *See* network access control
- namespace, DNS, 46
- NAT, 64–65
 - behavior of, 65
 - firewalls and configuration of, 290–293
 - firewalls vs., for internet security, 66
 - installing/configuring, 62
 - OSI model and, 65
 - PAT and, 65–66
 - proxy servers vs., 65

- using to protect internal IP addresses, 64
- NAT routers, 183
- Nbtstat.exe, 208
- Near-End Crosstalk (NEXT), 119, 158
- Nessus, 297
- NetBIOS Over TCP/IP (NetBT), 208
- Netfilter module, 285
- Netmask column (IPv4 routing tables), 63
 - subnet masks and, 63
- Netscape Communications, 254
- netstat, 210–213
- Netstat.exe, 210, 296
- network access control (NAC), 269–270
- network address translation (NAT), 290
 - Network Address Translation (NAT). *See* NAT
 - network address translation router. *See* NAT
- network administrator
 - change management and, 230
- network administrators
 - documentation and, 227
 - performance optimization and, 235
- network allocation vector (NAV), 168
- network appliances, 181–185
 - content filter, 184–185
 - exam need to know, 181–182
 - load balancer, 182–183
 - proxy server, 183–184
 - VPN concentrator, 185
- Network as a Service (NaaS), 57
- Network-based IDS (NIDS), 295
- Network Basic Input/Output System (Net-BIOS), 201
- Network Concepts domain, 1–60
 - answers to chapter questions, 57–60
 - common networking protocols, function of, 33–44
 - default ports, identifying TCP and UDP, 31–33
 - DNS concepts and components, summarizing, 45–48
 - IP addressing, purpose/properties of, 13–20
 - OSI and TCP/IP models compared, 1–4
 - OSI model layers, classifications related to, 4–13
 - troubleshooting methodology, 48–53
 - virtual network components, identifying, 53–57
- Network Control Protocol (NCP), 259
 - PPP and, 257
- Network Destination column (IPv4 routing tables), 63
- network devices
 - assigning values of given address class to, 15
- network diagrams
 - network maps vs., 226
- network elements
 - defining function of, 4
- Network+ exam objectives
 - TCP/IP section, revisions to, 4
- network identifier
 - boundary between host identifier and, 18
 - IPv4 addresses, 14
- networking protocols, function of, 33–44
 - ARP, 39
 - DHCP, 37
 - DNS, 38
 - exam need to know, 33–34
 - FTP, 37
 - HTTP, 38–39
 - HTTPS, 39
 - ICMP, 43
 - IGMP, 44
 - IMAP4, 42
 - NTP, 41
 - POP3, 41–42
 - RTP, 40
 - SIP, 40
 - SMTP, 42
 - SNMP2/3, 43
 - SSH, 41
 - TCP, 36
 - TCP/IP suite, 35–36
 - TELNET, 40
 - TFTP, 38
 - TLS, 44
 - UDP, 37
- network interface adapters. *See* NICs
- network interface cards. *See* NICs
- network interface device (NID), 175
- network interfaces
 - routers, 10
- network interface unit (NIU), 175
- network layer
 - PPP and, 257
 - switches and, 9

- Network layer (OSI model)
 - visualization of layers, 2
- network-layer protocol
 - PPP and, 259
- network management, 181–238
 - appliances, purpose/features of, 181–185
 - configuration management
 - documentation, 223–230
 - connectivity, using hardware tools to troubleshoot, 186–198
 - connectivity, using software tools to troubleshoot, 198–215
 - monitoring resources for analyzing traffic, 215–223
 - performance optimization, 230–235
- network management console, 43, 217
 - integrated SNMP agents and, 68
- Network Management domain, 181
- network maps, 225–226
- network media, 111–144
 - broadband vs. powerline networking
 - categorization by, 119–120
 - categorization of, 111–120
 - converters, media, 117–118
 - copper media, 113–116
 - distance/speed limitations,
 - categorization by, 118–119
 - exam need to know, 111–112
 - fiber media, 112–113
 - plenum vs. non-plenum cables, 117
 - standard connector types by, 121–125
- WAN technology types/properties,
 - categorization by, 128–144
- wireless standards for, 126–128

- Network Media and Topologies domain, 111
- network monitoring
- exam need to know, 216
- general logs, using, 221
- history logs, using, 221
- monitoring tools for, 215–223
- network sniffer, using, 222
- SNMP, using, 216–218
- SNMPv2, using, 218
- SNMPv3, using, 218
- syslog, using, 218–219
- system logs, using, 220–221
- traffic analysis, using, 221–222
- Network News Transfer Protocol (NNTP)
- SSL and, 254
- Network Security domain, 239
- network sniffer
- network monitoring using, 222
- network topologies. *See topologies*
- network traffic
- prioritizing, 71–72
- next hop, 23–24
- NICs
- OSI model, 12
- OSI model and, 5
- Nmap, 296
- nominal velocity of propagation (NVP), 196
- non-plenum cables, 117
- nonrepudiation, 263
- nslookup, 206–207
- NTP
- networking protocols, function of, 41
- numbering system, cable, 227

O

- Oakley, 253
- OFDM modulation, 77
- one-armed routers, 10
- open circuits, 190
- open fault, 191
- Open Shortest Path First (OSPF), 26
- Open Systems Interconnection (OSI)
 - reference model, 2–3. *See OSI model*
- optical carrier (OC) levels, 131
- optical-electrical-optical (OEO) regenerators, 132
- optical line terminal (OLT), 140
- optical network terminals (ONTs), 140
- optical time domain reflectometer. *See OTDR*
- organizationally unique identifier (OUI), 17, 243
- Orthogonal Frequency-Division Multiplexing (OFDM), 80, 126
- OSI model
 - bridges in, 12–13
 - cables in, 12
 - classifications related to, 4–13
 - encryption devices in, 11–12
 - EUI-64 in, 7
 - exam need to know, 2, 5–6
 - frames in, 8
 - hubs in, 11
 - IP addresses in, 6–7
 - layer numbers, use of, 3
 - MAC addresses in, 6
 - mnemonic for layer names of, 2
 - multilayer switches in, 10–11
 - NAT and, 65
 - NICs in, 12

- packets in, 9
- protocol analyzer and, 199
- routers in, 10
- switches in, 9–10
- TCP/IP model layer names, compared to, 3
- TCP/IP model vs., 1–4
- TCP/IP protocol stack vs., 3–4
- traditional numbering schemes, 3
- OSPF, 26
 - RIP vs., 26
- OTDR, 196–197

P

- packet filtering. *See* traffic filtering
 - stateful inspection vs., 283–285
- packets, 142
 - bridges and, 12
 - frames vs., 8
 - OSI model, 9
 - OSI model and, 5
 - TCP/IP network, 7
- packet sniffer, 192
- packet sniffing, 279
- packet switching, 132–134
- packet-switching, 142–143
- packet-switching networks, 142
- PAP
 - CHAP and, 271
- partial backups, 233
- passive optical network (PON), 140
- Password Authentication Protocol (PAP), 257
- passwords
 - multifactor authentication and, 273
- PAT, 65–66
 - firewalls and configuration of, 290–293
 - installing/configuring, 62, 65–66
 - NAT and, 65
- patch antenna, 74
- patch cables, 116, 124, 188, 192
- patches
 - general logs and, 221
- patch management
 - attack mitigation via, 280
- Path MTU Discovery, 23
- Path MTU Discovery (ICMP), 98
- PBX
 - virtual, 56
- peer-to-peer topology, 155–156
- performance
 - packet filtering and, 290
- Performance Monitor tool, 228
- performance optimization
 - caching engines and, 233–234
 - CARP and, 234
 - exam need to know, 231
 - fault tolerance and, 232–233
 - high availability and, 232
 - methods of, 231–234
 - methods/rationales for, 230–235
 - Quality of Service and, 231
 - reasons for, 234–235
 - traffic shaping and, 232
- peripheral networks, 246
- permanent virtual circuit (PVC), 132
- phishing, 277
- physical layer
 - TCP/IP suite and, 36
- Physical layer (OSI model)
 - visualization of layers, 2
- Ping, 200–202
 - DoS attacks using, 276
- Ping (diagnostic tool)
 - running from router interface, 70
- Ping utility (ICMP), 43
- PKI
 - MITM attacks and, 277
- PKI authentication, 263–264
- Plain Old Telephone Service (POTS), 134
- planning phase, 227
- plenum cables, 117
- PoE
 - limitations on, 69
- point-to-multipoint topology, 147–150
- Point-to-Point Protocol (PPP), 57
 - frames and, 8
 - VPNs and, 247
- point-to-point topology, 146–147
- policies and procedures
 - attack mitigation via, 280
- POP 3
 - TCP/IP suite, 35
- POP3
 - IMAP4 vs., 42
 - networking protocols, function of, 41–42
- port address translation (PAT), 291. *See* PAT
- port mirroring, 29
 - installing/configuring, 62, 72
 - Switched Port Analyzer (SPAN), 72

port number filtering, 245

ports

- configuration issues, troubleshooting, 96
- firewalls and security of, 282–283
- HTTP vs. HTTPS, 39
- identifying TCP and UDP default, 31–33
- well-known, 31

port scanners, 296

Post Office Protocol 3 (POP3)

- firewalls and, 283

power

- provided through PoE, 69

power failures

- troubleshooting, 98

powerline networking, broadband vs., 119–120

Power over Ethernet (PoE), 69

- installing/configuring, 62

Power Sum EL-FEXT (PS-ELFEXT), 158

Power Sum Equal-Level Far-End Crosstalk (PS-ELFEXT), 119

Power Sum Near-End Crosstalk (PS-NEXT), 119

Power Sum NEXT (PS-NEXT), 158

PPP

- Authentication phase, 258–259
- L2TP and, 249
- Link Dead phase in, 258
- Link Establishment phase, 258
- Link Open phase, 259
- link quality monitoring, 259
- Link Termination phase, 259
- network-layer protocol configuration, 259

PPPoE, 259–260. *See Point-to-Point Protocol over Ethernet*

PPP (Point-to-Point protocol). *See PPP*

PPPTP, 248–249

- L2TP and, 249

prefabricated patch cables, 189

presentation layer

- encryption devices and, 11
- Presentation layer (OSI model)
- visualization of layers, 2

Primary Rate Interface (PRI), 135

private branch exchange (PBX), 56

private key, 263

Process Name (protocol analyzer Frame Summary pane), 194

promiscuous mode, 193

propagation delay, 158

Protected EAP (PEAP), 272

protocol analyzer, 192–195, 199

protocol analyzers, 50, 279

- port scanners vs., 297

protocol data unit (PDU), 218

protocol data units (PDUs)

- TCP/UDP generation of, 36

protocol identifiers (packet filtering), 283

Protocol Name (protocol analyzer Frame Summary pane), 194

proxy servers, 183–184

- NAT routers vs., 65

PSTN, 56, 141

PTR records, 47

public key, 263

public key infrastructure (PKI), 263

Public Switched Telephone Network (PSTN), 134–135

punch down tool, 191–192

punching down, 115, 123, 191, 192

Q

QoS. *See Quality of Service*

Differentiated services (DiffServ), 71

- installing/configuring, 62

Integrated services (IntServ), 71

quality of service (QoS)

- ATM, 133

Quality of Service (QoS), 231

- installing/configuring, 71–72

R

RADIUS, 241

RADIUS (Remote Authentication Dial In User Service), 267–269

raised floors, 117

RAS. *See remote access server*

RAS (remote access server), 256

Rate-Adaptive Digital Subscriber Line (RADSL), 137

ratings, cable installation, 123

Real-Time Transport Protocol. *See RTP*

receiver-SMTP, 42

records

- DNS, 46–47

recursive queries, 46

Redundant Array of Independent Disks (RAID), 232

referrals, 47

- reliability
 - EIGRP, 25
- remote access, 256–261
 - PPP and, 256–259
 - PPPoE and, 259–260
 - RAS and, 256
 - remote terminal emulation and, 260
 - Secure Shell and, 260–261
- Remote Desktop, 200
- Remote Desktop Connection, 260
- Remote Desktop Protocol (RDP), 32, 260
- Remote Desktop Services, 260
- remote systems
 - TELNET vs. FTP access to, 40
 - remote terminal emulation and, 260
- repeater
 - media converters vs., 117
- repeaters, 11
- replacement antennae, 75
- requests
 - HTTP, 38
- request-to-send (RTS) messages, 168
- resistance, 197
- resource records (RRs), 46
- responses
 - HTTP, 38
- Return Loss (RL), 119, 158
- reverse name lookup, 45
- RFC 1122, "Requirements for Internet Hosts
Communication Layers", 4
- RFC 1510, 265
- RG-8 cable, 160
- RG-8/U cable, 118
- RG-58A/U cable, 118
- RG-58 coaxial cable, 160
- ring topology, 148
- RIP
 - metric values used by, 23
 - OSPF vs., 26
- RIPs, 24–25
- RJ-11 connector
 - butt set with, 189
- RJ-11 connectors, 123
- RJ-45 connector
 - butt set with, 189
- RJ-45 connectors, 122, 123, 124
 - attaching, using a crimper, 188
 - types of, 191
- rogue access points, 275
- rooms and enclosures, telecommunications, 172
- round robin DNS, 182, 183
- route cost, 23
 - OSPF, 26
- Route.exe, 213–215
- routers
 - diagnostics. installing/configuring, 70
 - DNS server addresses, 102
 - exam need to know, 93–94
 - filtering and, 246
 - firewalls vs., 282, 284
 - gateways, troubleshooting, 101
 - installing/configuring, 61–72
 - integrated web server, 68
 - interface configuration, 62, 69
 - internal diagnostics, 70
 - looping traffic and, 94
 - managed vs. unmanaged, 62
 - modules, troubleshooting problems
 - with, 100
 - NAT and, 290
 - one-armed, 10
 - OSI model, 10
 - OSI model and, 5
 - power failures, troubleshooting, 98
 - routing table use of, 64
 - stub, 10
 - traffic filtering, configuring on, 69–70
 - troubleshooting, 93–102
- routes
 - bad or missing, troubleshooting, 99
 - Metric column (routing table) and, 63
- routing
 - broadcast domains and, 29–30
 - convergence and, 27
 - EIGRP and, 25–26
 - exam need to know, 20–21
 - IGP vs. EGP, 30
 - link state vs. distance vector, 24
 - metrics for, 22–23
 - next hop and, 23–24
 - OSPF and, 26
 - static vs. dynamic, 22
- Routing and Remote Access Service
(RRAS), 286
- Routing Information Protocol. *See* RIP
- routing tables, 21–22
 - entries on a Windows Server 2008 R2 workstation, 63
 - errors in, troubleshooting, 99
 - formats of, 62
 - installing/configuring, 62–64

routing tables *Continued*
 Interface List (IPv6), 64
 IPv4 column functions, 63
 IPv6 format of, 64
 multiple entries for single network address, 64
 sending packets to networks not in the table, 64
 traffic filtering and, 69
 RS-232 standard, 125
 RTP
 networking protocols, function of, 40
 RTP Control Protocol (RTCP), 40

S

satellite-based services, 139
 scanners, 187
 scanners, inventory, 228
 scanning
 by NAT routers, 184
 scissors, 191
 screening, 115
 SC (Subscriber Connector) connectors, 122
 SDH, 142
 SDH (Synchronous Digital Hierarchy), 130–132
 secret key encryption, 263
 Secure Shell. *See* SSH
 Secure Shell (SSH), 32, 260–261
 Secure Sockets Layer (SSL)
 HTTPS utilization of, 39
 load balancing and, 182
 TLS and, 44
 security, 239–300. *See also* firewalls
 access, 244–262
 and accessibility to documentation, 227
 access point locations and, 74
 appliances and methods, 294–298
 encryption type problems, troubleshooting on wireless networks, 91
 incompatibility between protocols (wireless networks), 89
 netstat and, 212
 protocol analyzers and, 195
 SNMP, 43
 SNMPv2 and, 218
 SOHO network planning and, 103
 SSID broadcasts and, 79–80
 threats and attacks, common, 274–280
 user authentication and, 262–274
 VPN concentrators and, 185
 wireless security measures, 239–244

security association (SA), 252
 security protocols
 incompatibility between, 89
 segments, 8
 sender-SMTP, 42
 Sequence Number values (TCP headers), 285
 sequences, 8
 server
 syslog, 218
 server farms, 182
 servers
 caching-only, 46
 DNS, 45–46
 failure of, 232
 Kerberos encryption and, 265
 proxy, 183–184
 virtual, 55
 service-dependent filtering, 245
 service set identifier (SSID). *See* SSID
 Session Initiation Protocol. *See* SIP
 Session layer (OSI model)
 visualization of layers, 2
 shared secret, 240
 sheath
 plenum-grade, 117
 shielded twisted pair (STP) cable, 115
 shielding, 115
 short circuits, 190
 signal quality error (SQE), 166
 signal strength
 wireless security and, 244
 signal strength problems, 88
 bounce problems, 91
 sign-on authentication, 273–274
 Simple Mail Transfer Protocol. *See* SMTP
 Simple Mail Transfer Protocol (SMTP), 32, 219
 SSL and, 254
 Simple Network Management Protocol.
 See SNMP
 Simple Network Monitoring Protocol.
 See SNMP2/3
 Simple Password Exponential Key Exchange (SPEKE), 253
 single key encryption, 264
 singlemode fiber optic cable, 113, 119
 single sign-on, 273
 SIP
 networking protocols, function of, 40
 site survey
 wireless networks and, 74
 site-to-site connection, 248

- Small Form-Factor Pluggable (SFP) modular standard, 100
- small office/home office (SOHO) network.
 - See* SOHO networks
- smart cards, 273
 - EAP authentication and, 272
- smartjack, 176
- SMTP
 - firewalls and, 283
 - networking protocols, function of, 42
 - POP3 and, 41
 - TCP/IP suite, 35
- smurf attacks, 276–279
- sniffer, network
 - network monitoring using, 222
- SNMP
 - network monitoring using, 216–218
 - TCP/IP suite, 35
- SNMPv2/3
 - networking protocols, function of, 43
- SNMP-based management agents on routers/switches, 70
- SNMPv2
 - network monitoring using, 218
- SNMPv3
 - network monitoring using, 218
- social engineering, 277
- software
 - connectivity, 200
 - network diagram, 226
- software firewalls, 282
- software tools
 - ARP, 207–208
 - connectivity software, 200
 - connectivity troubleshooting with, 198–215
 - dig, 203–204
 - exam need to know, 199
 - Ipconfig/ifconfig, 204–206
 - Nbtstat.exe, 208
 - netstat, 210–213
 - nslookup, 206–207
 - Ping, 200–202
 - protocol analyzer, 199
 - Route.exe, 213–215
 - throughput tester, 199–200
 - Traceroute, 202–203
- SOHO networks
 - broadband routers and, 105
 - cable length, planning for, 104–105
 - device compatibility requirements and, 106–107
- devices on, planning for, 105
- environmental limitations on, 105–106
- equipment limitations, planning for, 106
- exam need to know, 103
- installing/configuring, 102–107
- operating systems and, 107
- requirements for, listing, 103–104
- virtual private network (VPN) connections and, 103
- SONET, 142
- SONET (Synchronous Optical Network), 130–132
- Source Address field
 - MAC addresses and, 6
- Source (protocol analyzer Frame Summary pane), 194
- spanning tree protocols, 27
- Spanning Tree Protocol (STP), 27, 59, 94
- Spatial Division Multiplexing (SDM), 127
- speed
 - Ethernet LAN, 169–170
 - WAN technologies classified by, 143
- speed limitations, categorization of cables by, 118–119
- split pairs, 157, 187
- split wire, 157
- spyware, 278
- SSH. *See* Secure Shell
 - networking protocols, function of, 41
- SSID, 79–80
 - mismatch problems, troubleshooting, 92
 - text string, anatomy of, 92
- SSL, 254–255
- SSL VPN, 250
- standards for wireless networks, 78–79
 - 802.11 standards, 78
 - compatibility between, 80–81
 - IEEE 802.11g standard, 80
 - IEEE 802.11n standard, 79, 80
 - incompatibility issues, troubleshooting, 89–90
 - Institute of Electrical and Electronics Engineers (IEEE), 78
- star topology, 150–152
 - hubs in, 11
 - switches in, 10
- stateful inspection
 - packet filtering vs., 283–285
- stateful packet inspection, 284
- stateless address autoconfiguration, 20

- static IP addressing
 - dynamic vs., 82
- Static NAT (SNAT), 291
- static routing
 - dynamic vs., 22
 - metric values in, 23
- STP cable. *See* shielded twisted pair cable
- ST (Straight Tip) connectors, 121
- stub routers, 10
- subnet masks, 16, 18
 - incorrect, troubleshooting, 100–101
 - in RIP, 25
 - IPv6 and, 64
 - Netmask column and, 63
- subnets
 - behavior of DHCP when changing, 84
 - DHCP scopes and, 83
 - rationale for, 18
- subnetting, 14
 - IP addresses, 18
- suffix (classless addresses), 16
- swapping, 232
- Switched Port Analyzer (SPAN), 72
- switches
 - auto-negotiation of port configurations, 96
 - basic, 10
 - cable length maximums and, 104
 - configuration interfaces in SOHO switches, 106
 - diagnostics. *installing/configuring*, 70
 - duplex mode settings on ports, 96
 - exam need to know, 93–94
 - hubs vs., 11, 167, 193
 - incorrect placement of, troubleshooting, 92–93
 - installing/configuring, 61–72
 - integrated web server, 68
 - interface configuration, 62, 68–69
 - internal diagnostics, 70
 - load balancing and, 182
 - managed, 68
 - managed vs. unmanaged, 62
 - manual port configuration, requirements for, 96
 - modules, troubleshooting problems with, 100
 - multilayer
 - OSI model, 9–10
 - OSI model and, 5
 - port configuration issues, troubleshooting, 96
- port mirroring on, 72
- power failures, troubleshooting, 98
- redundancy of, 94
- SOHO networks and, 105
- switching loops, troubleshooting, 94–95
- traffic filtering, configuring on, 69–70
- troubleshooting, 93–102
 - trunk ports, 67–68
 - unmanaged, 68
 - virtual, 55–56
- VLAN assignment problems, troubleshooting, 97
- VLANS, installing/configuring on, 67–68
- VTP, installing/configuring, 70–71

switching

- exam need to know, 20–21
- port mirroring and, 29
- spanning tree protocols and, 27
- VLANs and, 28–29

switching loops

- Spanning Tree Protocol (STP) and, 94
- troubleshooting, 94–95

symmetrical services, 142

Symmetric Digital Subscriber Line (SDSL), 136

Synchronous Digital Hierarchy (SDH), 131

SYN floods, 284

syslog

- network monitoring using, 218–219

syslog server, 218

system logs

- network monitoring using, 220–221

T

- T-1**, 129, 135
- T-3**, 129
- T568A** pinout, 192
- T678B** pinout, 192
- TACACS+**, 269
- T-carrier service**, 129
- TCP**
 - default ports, identifying, 31–33
 - networking protocols, function of, 36
 - TCP/IP suite, 36
- TCP/IP**
 - configuration parameters, with DHCP, 85
 - IP addressing, 13–20
 - Network+ exam objectives, revisions to, 4
 - OSI model vs., 1–4

- protocol stack vs. OSI model layers, 3–4
- subnet mask issues,
- troubleshooting, 100
- TCP/IP model**
 - exam need to know, 2
 - IP addresses and, 6
 - OSI model layer names, compared to, 3
 - RFC 1122, "Requirements for Internet Hosts Communication Layers", 4
- TCP/IP networks**
 - packets on, 7
 - router tables on, 22
- TCP/IP suite**
 - networking protocols, function of, 35–36
- TDR**, 196
- teardrop attacks, 284
- Telecommunications Industry Association/**
 - Electronic Industries Alliance (TIA/EIA), 114
- telecommunications rooms and enclosures, 172
- telephone connections, leased lines for, 129
- telephone system
 - PSTN, 134
- Telnet**, 260
- TELNET**, 32
 - FTP vs., 37
 - networking protocols, function of, 40
- temperature monitor, 198
- Temporal Key Integrity Protocol (TKIP)**, 241
- Tenable Network Security**, 297
- 10 Gigabit Ethernet, 163–165
- Terminal Access Controller Access-Control System (TACACS)**, 269
- TFTP**
 - networking protocols, function of, 38
- Thick Ethernet**, 144, 153, 160
- Thin Ethernet**, 153, 160
- threats and attacks, 274–280
 - buffer overflow, 278–279
 - common attacks, 276–279
 - denial of service (DoS), 276
 - distributed DoS attacks, 276
 - exam need to know, 274
 - FTP bounce attacks, 277
 - malware, 278
 - man in the middle (MITM) attacks, 277
 - mitigation techniques, 279–280
 - packet sniffing, 279
 - smurf attacks, 276–279
 - social engineering, 277
- wireless networks, common threats to, 274–275
- three-way handshake**, 36
- throughput**
 - bandwidth vs., 200
- throughput tester**, 199–200
- Ticket-Granting Service (TGS)**, 265
- ticket-granting ticket (TGT)**, 265
- Time Date (protocol analyzer Frame Summary pane)**, 194
- time division multiplexing**, 129
- time domain reflectometry**. See **TDR**
- Time Offset (protocol analyzer Frame Summary pane)**, 194
- Time to Live (TTL) field**, 201, 202
- TLS**, 254–255
 - networking protocols, function of, 44
- token passing**, 148
- Token Ring**, 115
- Token Ring protocol**, 148
- tone generator**
 - wiremap tester vs., 187
- tone generator and locator**, 190
- toner probe**, 190–191
- topologies**, 144–156
 - bus, 153–155
 - client/server, 155
 - exam need to know, 145
 - hybrid, 156
 - mesh, 152–153
 - MPLS, 145–146
 - peer-to-peer, 155–156
 - point-to-multipoint, 147–150
 - point-to-point, 146–147
 - star, 150–152
- Traceroute**, 202–203
- Traceroute (diagnostic tool)**
 - running from router interface, 70
- Tracert.exe**, 202
- traffic analysis**
 - network monitoring using, 221–222
- traffic filtering**
 - configuring, 62, 69–70
 - unicast packets, 69
- traffic shaping**, 232
- training**
 - attack mitigation via, 279
- Transmission Control Protocol**. See **TCP**
- Transmission Control Protocol (TCP)**, 58
- transmission media**
- WAN technologies classified by**, 143

- transmission speed
 - OSPF, 26
- transmission speeds
 - reduction of due to bad cables, troubleshooting, 95
- Transport layer (OSI model)
 - visualization of layers, 2
- Transport Layer Security. *See* TLS
- Transport Layer Security (TLS)
 - HTTPS utilization of, 39
 - SSL VPN and, 250
- Transport layer (TCP/IP model)
 - visualization of layers, 3
- transport mode, 253–254
- traps, 217
- Trivial File Transfer Protocol. *See* TFTP
- Trojan Horse attacks, netstat and detection of, 212
- trojans, 278
- troubleshooting
 - connectivity problems, 157–159
 - connectivity, using hardware tools to troubleshoot, 186–198
 - connectivity, using software tools for troubleshooting, 198–215
 - determine cause, testing theory to, 50–51
 - documentation and, 53
 - escalation of, 52
 - exam need to know, 48, 199
 - identifying the problem and, 49–50
 - methodology for, 48–53
 - plan of action, establishing a, 51
 - preventative measures, implementing, 52
 - routers/switches, 93–102
 - solution, implementing the, 52
 - theory of probable cause, establishing a, 50
 - wireless networks, 86–93
 - wire schemes for, 224
- trunking
 - between VLAN switches, 67–68
 - failures due to VLAN assignment problems, 97
 - link aggregation vs., 68
 - Multiple VLAN Registration Protocol (MVRP), 71
 - VLAN Trunking Protocol (VTP), 70
- trunk ports (VLANs), 67–68
- tunneling, 246–255
 - defined, 248
 - IPsec and, 250–254
 - key exchange protocols and, 252–253
 - L2TP and, 249
 - OSI model encapsulation and, 255
 - PPTP and, 248–249
 - SSL/TLS and, 254–255
 - SSL VPN and, 250
 - VPN and, 246–250
 - tunnel mode, 253–254
 - twisted pair, 143
 - twisted pair cable
 - unshielded, 113–115
 - twisted pair cables
 - connectors for, 122
 - two-factor authentication, 273
 - Type of Service (TOS) field, 231

U

- UDP
 - default ports, identifying, 31–33
 - diskless workstation startup with, 38
 - networking protocols, function of, 37
 - RADIUS and, 267
 - TCP and, 36
 - TCP/IP suite, 36
 - TFTP use of, 38
- unicast packets, 69
- unicasts
 - IPv4 addresses and, 19
- Uniform Resource Locator (URL), 254
- UNIX
 - NAT and, 292
- UNIX systems
 - dig with, 203
 - syslogs with, 219
 - traceroute on, 202
- unshielded twisted pair (UTP) cable, 113–115
- unshielded twisted pair (UTP) cables
 - ratings of, 187
- Update messages, 48
- updates
 - general logs and, 221
- user authentication, 262–274
 - AAA, 267–269
 - CHAP, 271
 - EAP, 272
 - exam need to know, 262–263
 - Kerberos, 265–267
 - MS-CHAP, 271–272

multifactor authentication, 273
 network access control and, 269–270
 PKI, 263–264
 sign-on authentication, 273–274
 two-factor authentication, 273

User Datagram Protocol. *See* UDP

User Datagram Protocol (UDP), 58, 217
 L2TP and, 249

UTP cable, 118. *See also* unshielded twisted pair cable
 straight-through wiring with, 191

UTP cables
 troubleshooting, 158

V

vertical cross connect, 174

Very High-rate Digital Subscriber Line (VDSL), 137

virtual desktops, 54–55

virtual LANs (VLANs). *See also* VLANs
 switches and creation of, 10

virtual machine (VM) environment, 55

virtual network components,
 identifying, 53–57
 exam need to know, 54
 Network as a Service (NaaS), 57
 onsite vs. offsite locations, 57

virtual PBX, 56

virtual private networks. *See* VPNs

virtual private networks (VPNs), 185

virtual private network (VPN) connection
 SOHO networks and, 103

virtual servers, 55

virtual switches, 55–56

virtual WANs
 VPNs and, 247

viruses, 278

VLANs
 assignment problems, troubleshooting, 97
 installing/configuring trunks for, 62, 67–68
 switching and, 28–29
 trunking failures due to assignment problems, 97

VLAN Trunking Protocol (VTP)
 installing/configuring, 62, 70–71

Voice over IP, 235

Voice over IP (VoIP)
 SIP use by, 40
 wireless networks and, 90

Voice over IP (VOIP)
 RTP use of, 40

voice telephone systems
 UTP used by, 116

VoIP
 PBX and, 56

voltage event recorder, 198

Volt-Ohm Meter (VOM), 197

VPN, 246–250

VPN concentrators, 185

vulnerability scanners, categorization of, 296–298

W

walls, physical
 bounce problems with wireless networks, 91
 materials of, and wireless networks, 75, 88

WAN technology
 VPNs and, 246

WAN technology types/properties, categorization of network media by, 128–144
 distance, by, 144
 exam need to know, 129
 Internet access and, 134–142
 leased lines, 129–130
 packet switching, 132–134
 packet- vs. circuit-switching, 142–143
 SONET/SDH, 130–132
 speed, by, 143
 transmission media, by, 143

war chalking, 275

war driving, 275

wavelength division multiplexing (WDM), 132

web browsers
 FTP client functioning, 37

webpages
 HTTP transactions for displaying, 39

web servers
 public IP addresses for, 15

websites
 HTTPS support on, 39

web switches, 182

well-known ports, 31

WEP, 240–241
 cracking, 275
 IEEE 802.1X standard and, 270

wide area networks (WANs)
 frames and, 8

- Wi-Fi Alliance, 241
- Wi-Fi Protected Access. *See* WPA
- Windows Firewall, 286
- Windows Server
 - NAT and, 292
- Windows Terminal Services, 260
- Windows XP Mode, 54
- WinFrame, 260
- Wired Equivalent Privacy. *See* WEP
- wireless media, 143
- wireless networks
 - 802.11ac standard, 76
 - 802.11b/g standards, 76
 - 802.11n standard, 76
 - access points, locating, 74
 - access points, locations of, 73
 - antennae, 73
 - appliances and, 75
 - channel configuration for large, 77
 - channels, 73
 - channels for, configuring, 76–78
 - channels, troubleshooting problems with, 90
 - common threats to, 274–275
 - compatibility between standards, 80–81
 - configuration issues, troubleshooting, 88
 - device compatibility problems, troubleshooting, 89–90
 - encryption type problems, troubleshooting, 91
 - exam need to know, 73, 87
 - frequencies, selecting, 73
 - IEEE 802.11n standard, 79
 - incompatibility between security protocols (wireless networks), 89
 - installing/configuring, 73–81
 - Institute of Electrical and Electronics Engineers (IEEE), 78
 - interference, 73
 - interference, preventing, 75
 - interference problems, troubleshooting, 87–88
 - large office buildings and, 76
 - latency, troubleshooting, 90–91
 - Multiple-Input Multiple-Output (MIMO), 79
 - power saving mechanisms and, 89
 - signal bounce problems, troubleshooting, 91
 - signal strength problems, 88
 - SSID broadcasts, disabling, 73, 79–80
 - SSID mismatch problems, troubleshooting, 92
 - standards, 73, 78–79
 - standards, compatible, 73
 - switches, troubleshooting incorrect placement of, 92–93
 - troubleshooting, 86–93
 - Voice over IP (VoIP) and, 90
 - wireless security, 239–244
 - device placement and, 243–244
 - encryption protocols and, 240–242
 - exam need to know, 239
 - MAC address filtering and, 242–243
 - signal strength and, 244
 - wireless standards for network media, 126–128
 - 802.11a/b/g/n standards, comparing, 126–128
 - exam need to know, 126
 - wiremap tester, 157, 186
 - wiremap testers
 - loopback plug and, 195
 - wire schemes
 - purpose of, 223–225
 - wiring distribution, components of, 171–177
 - CSU/DSU, 176
 - demarc, 175
 - demarc distribution, 175–176
 - exam need to know, 172
 - IDF, 172–173
 - MDF, 173–175
 - smartjack, 176
 - work area, 172
 - Worldwide Interoperability for Microwave Access (WiMAX), 141
 - worms, 278
 - WPA, 241–242
 - cracking, 275
 - WPA2, 241
 - cracking, 275
 - WPA2-AES encryption protocol algorithm, 91
 - WPA2-Enterprise, 241
 - WPA2-TKIP encryption protocol algorithm, 91

Z

- Zenmap, 296
- zero-configuration networking, 84
- zombies, 278

About the author

Craig Zacker is the author or co-author of dozens of books on operating systems, networking topics, and PC hardware, including *Windows Small Business Server 2011 Administrator's Pocket Consultant* and *MCITP Self-Paced Training Kit for Exam 70-686: Windows 7 Desktop Administrator*, both for Microsoft Learning. He has also been an English professor, a network administrator, a webmaster, a corporate trainer, a darkroom technician, a library clerk, a student, and a newspaper boy. He lives in the Susquehanna Valley with his wife and a neurotic cat.

What do you think of this book?

We want to hear from you!

To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Tell us how well this book meets your needs—what works effectively, and what we can do better. Your feedback will help us continually improve our books and learning resources for you.

Thank you in advance for your input!

Microsoft®
Press