

EXAM ✓ **CRAM**

Cisco CCENT

ICND1 100-101

Second Edition



PEARSON IT
CERTIFICATION

MICHAEL VALENTINE
KEITH BARKER

FREE SAMPLE CHAPTER



SHARE WITH OTHERS

EXAM ✓ CRAM

Cisco CCENT

ICND1 100-101

Second Edition

**Michael Valentine
Keith Barker**

Cisco CCENT ICND1 100-101 Exam Cram

Copyright © 2014 by Pearson

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-7897-5150-8

ISBN-10: 0-7897-5150-x

Library of Congress Publication Control Number: 2013955306

First Printing: December 2013

Trademarks

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Pearson cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the CD or programs accompanying it.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact international@pearsoned.com.

Publisher

Paul Boger

Associate Publisher

David Dusthimer

Executive Editor

Brett Bartow

Senior Development Editor

Christopher
Cleveland

Managing Editor

Sandra Schroeder

Project Editor

Seth Kerney

Copy Editor

Keith Cline

Indexer

Ken Johnson

Proofreader

Jess DeGabriele

Technical Editors

Brian D’Andrea

Andrew Whitaker

Publishing Coordinator

Vanessa Evans

Multimedia Developer

Lisa Matthews

Book Designer

Mark Shirar

Page Layout

Bumpy Design

Contents at a Glance

	Introduction	1
	Self-Assessment	17
CHAPTER 1	Networking Fundamentals	25
CHAPTER 2	Network Models	39
CHAPTER 3	Concepts in IP Addressing	69
CHAPTER 4	Working with Cisco Equipment	133
CHAPTER 5	Securing Your Cisco Devices	169
CHAPTER 6	Basic Switch Operations and Configuration	189
CHAPTER 7	Basic Routing	225
CHAPTER 8	Visualizing Data Flow Between Two Hosts on a Network	271
CHAPTER 9	IP Access Lists	289
CHAPTER 10	IP Services	319
CHAPTER 11	Troubleshooting	347
	Practice Test 1	377
	Answers to Practice Test 1	413
	Practice Test 2	423
	Answers to Practice Test 2	459
	Glossary	469
	Index	487

Table of Contents

Introduction	1
Taking a Certification Exam	2
Tracking Your Certification Status	3
How to Prepare for an Exam	4
What This Book Will Not Do	5
What This Book Is Designed to Do	5
About This Book	6
CCENT Official Exam Topics	6
Chapter Formats	12
Pearson IT Certification Practice Test Engine and Questions on the CD-ROM	14
Install the Software from the CD	14
Activate and Download the Practice Exam	15
Activating Other Exams	16
Contacting the Authors	16
Self-Assessment	17
CCENT in the Real World	17
The Ideal CCENT Candidate	18
Put Yourself to the Test	19
CHAPTER 1:	
Networking Fundamentals	25
Essential Terms and Components	26
Topologies	27
Ethernet	28
WAN Technologies	31
Review Questions	34
Answers to Review Questions	36
What Next?	37
CHAPTER 2:	
Network Models	39
Cisco Hierarchical Design	40
Access Layer	40

Distribution Layer	41
Core Layer	41
Advantages of the Three-Layer Model	42
OSI Model	44
Layer 7: The Application Layer	45
Layer 6: The Presentation Layer	47
Layer 5: The Session Layer	47
Layer 4: The Transport Layer	47
Layer 3: The Network Layer	54
Layer 2: The Data Link Layer	56
Layer 1: The Physical Layer	57
TCP/IP Model	61
Review Questions	64
Answers to Review Questions	67
CHAPTER 3:	
Concepts in IP Addressing	69
Binary	70
Hexadecimal	74
Conversion Between Binary, Hex, and Decimal	75
Decimal to Hexadecimal Conversions	75
Hexadecimal to Decimal Conversions	75
IP Address Components	77
Address Class	77
Public and Private IP Addresses	78
Default Subnet Mask	78
The Network Field	79
The Host Field	80
Nondefault Masks	80
The Subnet Field	81
Subnetting	82
Address Class and Default Mask	83
The Increment	84
Number of Hosts	86
The Broadcast ID	88
Number of Subnets	89

Working with Subnetting Questions	91
The Subnetting Chart	93
VLSM	96
Route Summarization	100
Summarization Guidelines	102
Determining Summary Addresses	104
IPv6	110
IPv6 Address Allocation	112
IPv6 Address Notation	113
Types of IPv6 Addresses	114
IPv6 Address Configuration	115
IPv6 Router Configuration	116
IPv6 Features	117
The IPv6 Header	118
IPv6 Transition Strategies	119
Review Questions	122
Answers to Review Questions	128
What Next?	131
CHAPTER 4:	
Working with Cisco Equipment	133
Introducing Your Cisco Gear	134
External Connections	135
Other Connections	139
Connecting and Configuring Cisco Devices	141
Device Memory Locations	142
IOS Startup Process	142
Setup Mode	144
Configuration Register	144
Password Recovery	145
Command-Line Modes	146
Command Shortcuts	148
Context-Sensitive Help	149
Managing IOS Licensing	150
Remote Management	152
Telnet	152
Secure Shell	154

Cisco Discovery Protocol	155
Internet Control Message Protocol	156
Network Time Protocol	158
Simple Network Management Protocol	158
Syslog	159
NetFlow	159
Basic Switch Configuration	159
Basic Router Configuration	160
Review Questions	164
Answers to Review Questions	166
What Next?	167
CHAPTER 5:	
Securing Your Cisco Devices	169
Securing Routers and Switches	170
Configuring Password Security	170
Configuring SSH Access to Your Router	174
External Authentication Methods	176
Switch Port Security	178
Security Best Practices for Switches	178
Review Questions	184
Answers to Review Questions	186
What Next?	187
CHAPTER 6:	
Basic Switch Operations and Configuration	189
Bridging and Switching	190
Functions of Bridges and Switches	191
The Differences Between Switches and Bridges	195
STP	200
VLAN Concepts and Applications	204
The Definition of a VLAN	206
Benefits of VLANs	206
Implementing VLANs	208
VLAN Membership	210
Trunking	212
802.1Q	214
Configuring Switches for Trunking	215

Review Questions	219
Answers to Review Questions	222
What Next?	224
CHAPTER 7:	
Basic Routing	225
Routing Fundamentals	226
The Next Hop Concept	227
Connected Networks	228
Static Routes	229
Default Routes	230
Dynamic Routes	232
Understanding the Router Lookup Process and the IP Route Table	235
Inter-VLAN Routing	242
Router on a Stick	243
L3 Switching	245
OSPF	248
Characteristics	249
Configuring Single-Area OSPFv2	254
Configuring Single-Area OSPFv3	258
Verifying and Troubleshooting OSPF	260
Verification of OSPFv3	261
Review Questions	265
Answers to Review Questions	269
CHAPTER 8:	
Visualizing Data Flow Between Two Hosts on a Network	271
Host-to-Host Direct Communications	272
Host to Host Through a Switch	275
Host to Host Through a Router	278
Review Questions	283
Answers to Review Questions	286
What Next?	288
CHAPTER 9:	
IP Access Lists	289
IP Access Lists as Packet Filters	290
Types of ACLs	293

Configuring and Implementing	296
Configuring Standard ACLs	297
Configuring Extended ACLs	300
Filtering Telnet and SSH Access	302
Troubleshooting and Verifying ACL Configurations	308
Review Questions	311
Answers to Review Questions	316
What Next?	318
CHAPTER 10:	
IP Services	319
Network Address Translation	320
NAT Terminology	321
Applications, Advantages, and Disadvantages of NAT	323
Static NAT	324
Dynamic NAT	326
PAT	328
Verification and Troubleshooting NAT and PAT	330
Network Time Protocol	334
Configuring an IOS Router as an NTP Client	335
Verification of NTP	336
DHCP	338
Configuring an IOS Router as a DHCP Client	338
Configuring an IOS Router as a DHCP Server	339
Verification of DHCP Operation	340
Review Questions	343
Answers to Review Questions	345
CHAPTER 11:	
Troubleshooting	347
Troubleshoot Layer 1 Problems	348
Problems with the Wiring	348
Duplex	349
Finding the Errors	349
Verify and Troubleshoot VLANs and Port Membership	353
Troubleshooting VLANs	354

Verify Trunk Configuration and Operation	357
Is the Link Up and Configured?.	357
VLANs Must Exist for the Trunk to Support Them.	357
Administrative Versus Operational	358
Troubleshoot Problems with IP Addressing and Host Configurations . .	361
Verify Hosts Have Correct IP Addresses	362
How to Remotely Change a Cisco IP Address	362
Making Sure That We Have a Default Gateway on Hosts.	362
Troubleshooting NAT	363
Troubleshoot ACL Issues	365
Review Questions	368
Answers to Review Questions.	375
Practice Exam 1	377
Answer Key to Practice Exam 1	413
Answers at a Glance to Practice Exam 1	413
Answers with Explanations.	413
Practice Exam 2	423
Answer Key to Practice Exam #2	459
Answers at a Glance to Practice Exam 2	459
Answers with Explanations.	460
Glossary	469
Index	487

About the Authors

Michael Valentine has been in the IT field for 16 years, focusing on network design and implementation. He is a Cisco Certified Systems Instructor (#31461) and specializes in Cisco Unified Communications instruction and in CCNA and CCNP courses. His accessible, humorous, and effective teaching style has demystified Cisco for hundreds of students since he began teaching in 2002. Michael has a Bachelor of Arts degree from the University of British Columbia and currently holds CCNA, CCDA, CCNP, and CCNP-Voice certifications. In addition to the popular *Exam Cram 2: CCNA* and *CCENT* books, Michael co-authored the *Official Certification Guide for CCNA-Voice* and has contributed to and served as technical editor for the Cisco Press titles *CCNP ONT Official Exam Certification Guide* and *CCNA Flashcards*, among others. Michael has also developed courseware and lab guides for Official Cisco Curriculum courses and custom classes for individuals and corporations.

Keith Barker, CCIE No. 6783 R/S & Security, is a 27-year veteran of the networking industry. He currently works at CBT Nuggets. His past experience includes EDS, Blue Cross, Paramount Pictures, and KnowledgeNET, and he has delivered CCIE-level training for several years. As part of the original set of Cisco VIPs for the Cisco Learning Network, he continues to give back to the community in many ways. He is CISSP, HP, PaloAlto, Brocade, and Juniper certified, loves to teach, and keeps many of his video tutorials at <http://www.youtube.com/keith6783>. You can reach him on Facebook at “Keith Barker Networking.”

About the Technical Editors

Brian D'Andrea started his career working as a bench technician for a large computer manufacturer. He then progressed to a consultant position for various financial and medical institutions across Pennsylvania, New Jersey, and Delaware. He is now a long-time instructor and courseware developer of Cisco courses that include CCNA Routing & Switching, CCDA, CCNA Security, CCNP Routing & Switching, and CCDP. He has been privileged to be part of several Cisco Press published materials. He enjoys sharing his knowledge and 17 years of experience in the information technology field.

Andrew Whitaker (CCNA:Security, CCNP, CCSP, CCVP, CCDP, CCDA, CCENT, CISSP, CEH, CEPT, CPT, LPT, MCT, CEI, CICP, CHFI, ECSA, MCTS, MCSE, CNE, EMCPA, CTP, A+, Network+, Security+, Convergence+, Linux+, CERE, WAPT, CSSA, LPI-1) is a nationally recognized expert on cybersecurity, an author of best-selling networking and security books, and an award-winning technical trainer. He was also contributing author to previous editions of the *CCNA* and *CCENT Exam Crams*. His work has gained media coverage by NBC, *The Wall Street Journal*, *The Philadelphia Inquirer*, *San Francisco Gate*, *Business Week* magazine, and others. He is a frequent conference speaker and has given talks at GFIRST8, DefCon, TakeDownCon, ChicagoCon, BSides, and SecurePhilly. As an instructor, he is the recipient of both the EC-Council Instructor of Excellence award and the EC-Council Instructor of the Year award.

Dedication

Michael Valentine:

This one is for Andy DeMaria, a good boss and a stalwart friend.

Keith Barker:

*To Dan and Johnna Charbonneau and the CBT Nuggets family:
If anyone had the opportunity to be surrounded by such awesome people,
even for a few hours, they would never be the same again. Thank you!*

Acknowledgments

With this book in its latest iteration, I have thanked several people several times over. But a few of those people deserve to be thanked every day.

Brett Bartow, Chris Cleveland, Vanessa Evans, and all the unseen warriors at Pearson: You put up with a lot and rarely lose your patience with your authors and their egos. I'm impressed and grateful for your long-suffering help.

The crew at Skyline—Toby, Ed, Marshall and Suad—have taught me, coached me, saved me, and/or kicked my butt as required and been such constant, dependable, professional colleagues, and friends I can't imagine working without. They make me better at what I do.

Andy DeMaria (despite being a fan of a disreputable New York hockey team) is possibly the best boss I have ever had. Through thick and thin, a serious and debilitating injury, border guard zeal, and general gong show snafus, he has been a deep well of patience, logic, perceptive observation, and effective action. If either of us ever has 10 minutes to spare, I owe him (at last count) two fresh salmon and three cases of good Canadian beer. Seriously, Andy: Thanks.

My little girl has grown up seeing me glued to my chair, grinding away on chapters, slides, lab guides, and lectures. She has learned a whole new sign language for when I am not able to talk because I am lecturing, brought me dozens of cups of coffee and meals, and generally been a great kid and assistant. Pretty soon she will be able to type better than I can (not that that is any great challenge) and she can take over while I dictate. Thank you for being a great kid and a big help.

My wonderful wife: You and me, together. Thank you for being the kind of girl who enjoys driving around New Zealand in a hippie van named Vince.

And last of all: You, the reader, the student, the learner, the exam taker. Thanks for all your work. I wish you success.

We Want to Hear from You!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our way.

We welcome your comments. You can email or write us directly to let us know what you did or didn't like about this book—as well as what we can do to make our books better.

Please note that we cannot help you with technical problems related to the topic of this book.

When you write, please be sure to include this book's title and authors as well as your name, email address, and phone number. We will carefully review your comments and share them with the author and editors who worked on the book.

Email: feedback@pearsonitcertification.com.

Mail: Pearson IT Certification
ATTN: Reader Feedback
800 East 96th Street
Indianapolis, IN 46240 USA

Reader Services

Visit our website and register this book at <http://www.pearsonitcertification.com/register> for convenient access to any updates, downloads, or errata that might be available for this book.

Introduction

Welcome to *Cisco CCENT ICND1 100-101 Exam Cram*! Whether this is your first or your fifteenth Exam Cram series book, you'll find information here that will help ensure your success as you pursue knowledge, experience, and certification. This introduction explains Cisco's certification programs in general and talks about how the Exam Cram series can help you prepare for the Cisco CCENT ICND1 100-101 exam. The materials in this book have been prepared with a very clear focus on testable concepts, configurations, and skills. As much extraneous material as possible, beyond what is needed for background comprehension, has been eliminated so that the book is a distillation of the necessary knowledge to take—and pass—the CCENT ICND1 100-101 exam. The two sample tests with answer keys at the end of the book should give you a reasonably accurate assessment of your knowledge. We have also included challenge labs to give you the critical hands-on practice you will need to master the simulator questions on the CCENT ICND1 100-101 exam. Read the book, understand the material, practice the labs, and you'll stand a very good chance of passing the test.

Exam Cram books help you understand and appreciate the subjects and materials you need to pass Cisco certification exams. Exam Cram books are aimed strictly at test preparation and review. They do not teach you everything you need to know about a topic. Instead, we present and dissect the topics and key points we've found that you're likely to encounter on a test. We've worked to bring together as much accurate information as possible about the latest CCENT ICND1 100-101 exam.

Nevertheless, to completely prepare yourself for any Cisco test, we recommend that you begin by taking the Self-Assessment that is included in this book, immediately following this introduction. The Self-Assessment tool will help you evaluate your knowledge base against the requirements for a CCENT under both ideal and real circumstances.

Based on what you learn from the Self-Assessment, you might decide to begin your studies with some classroom training, some practice with the Cisco IOS, or some background reading. Or, you might decide to pick up and read one of the many study guides available from Cisco

or third-party vendors on certain topics. We also recommend that you supplement your study program with visits to <http://www.examcram2.com> to receive additional practice questions, get advice, and track the CCENT and CCNA programs.

We also strongly recommend that you practice configuring the Cisco devices that you'll be tested on because nothing beats hands-on experience and familiarity when it comes to understanding the questions you're likely to encounter on a certification test. Book learning is essential, but without a doubt, hands-on experience is the best teacher of all.

Taking a Certification Exam

After you've prepared for your exam, you need to register with a testing center. The cost of the CCENT ICND1 100-101 exam is \$125 and is administered by VUE in the United States and Canada.

You can sign up for a test or get the phone numbers for local testing centers through the Web at <http://www.vue.com>.

To sign up for a test, you must possess a valid credit card or contact VUE for mailing instructions to send a check (in the United States). Only when payment is verified or your check has cleared can you actually register for the test.

To schedule an exam, you need to call the number or visit either of the web pages at least one day in advance. To cancel or reschedule an exam, you must call before 7 p.m. Pacific Standard Time the day before the scheduled test time (or you might be charged, even if you don't show up to take the test). When you want to schedule a test, you should have the following information ready:

- ▶ Your name, organization, and mailing address
- ▶ Your Cisco test ID
- ▶ The name and number of the exam you want to take
- ▶ A method of payment (As mentioned previously, a credit card is the most convenient method, but alternative means can be arranged in advance, if necessary.)

After you sign up for a test, you are told when and where the test is scheduled. You should try to arrive at least 15 minutes early. You must supply two forms of identification—one of which must be a photo ID—and sign a nondisclosure agreement to be admitted into the testing room.

All Cisco exams are completely closed book. In fact, you are not permitted to take anything with you into the testing area, but you are given a blank sheet of paper and a pen (or in some cases, an erasable plastic sheet and an erasable pen). We suggest that you immediately write down on that sheet of paper all the information you've memorized for the test. In Exam Cram books, this information appears on a tear-out sheet inside the front cover of each book. You are given some time to compose yourself, record this information, and take a sample orientation exam before you begin the real thing. We suggest that you take the orientation test before taking your first exam, but because all the certification exams are more or less identical in layout, behavior, and controls, you probably don't need to do this more than once.

When you complete a Cisco certification exam, the software tells you immediately whether you've passed or failed. If you need to retake an exam, you have to schedule a new test with VUE and pay another \$100 or \$125.

Note

If you fail a Cisco test, you must wait 5 full days before you can take it again. For example, if you failed on Tuesday, you would have to wait until Monday to take it again.

Tracking Your Certification Status

As soon as you pass the CCENT ICND1 100-101 exam, you are a CCENT. Cisco generates transcripts that indicate which exams you have passed. You can view a copy of your transcript at any time by going to Cisco.com and going to the certifications tracking tool. This tool enables you to print a copy of your current transcript and confirm your certification status.

After you pass the necessary exam, you are certified. Official certification is normally granted after 3 to 6 weeks, so you shouldn't expect to get your credentials overnight. The package for official certification includes the following

- ▶ A certificate that is suitable for framing, along with a wallet card.
- ▶ A license to use the applicable logo, which means that you can use the logo in advertisements, promotions, and documents, as well as on letterhead, business cards, and so on. Along with the license comes information on how to legally and appropriately use the logos.

Many people believe that the benefits of Cisco certification are among the most powerful in the industry. We're starting to see more job listings that request or require applicants to have CCNA, CCDA, CCNP, and other certifications, and many individuals who complete Cisco certification programs can qualify for increases in pay and/or responsibility. As an official recognition of hard work and broad knowledge, one of the Cisco credentials is a badge of honor in many IT organizations.

How to Prepare for an Exam

Preparing for the CCENT ICND1 100-101 exam requires that you obtain and study materials designed to provide comprehensive information about the product and its capabilities that will appear on the specific exam for which you are preparing. The following list of materials can help you study and prepare:

- ▶ The official Cisco study guides by Cisco Press.
- ▶ Practicing with real equipment or simulators.
- ▶ The CCNA Prep Center on Cisco's website, which features articles, sample questions, games, and discussions to focus and clarify your studies.
- ▶ The exam-preparation advice, practice tests, questions of the day, and discussion groups on the <http://www.examcram.com> e-learning and certification destination website.
- ▶ The *Cisco CCENT ICND1 100-101 Exam Cram*—This book gives you information about the material you need to know to pass the tests. Seriously, this is a great book.
- ▶ Classroom training. Cisco training partners and third-party training companies (such as the Training Camp) offer classroom training for CCENT and CCNA related skills and topics. These companies aim to help you prepare to pass the CCENT and CCNA exams. Although such training can be expensive, most of the individuals lucky enough to partake find this training to be very worthwhile.
- ▶ Other publications. There's no shortage of materials available about CCENT and CCNA. Try to remember not to drown yourself in reading material—at some point, you are just ready to test and should go for it.

This set of required and recommended materials represents a good collection of sources and resources about the CCENT exam and related topics. We hope that you'll find that this book belongs in this company.

What This Book Will Not Do

This book will not teach you everything you need to know about networking with Cisco devices, or even about a given topic. Nor is this book an introduction to computer technology. If you're new to networking and looking for an initial preparation guide, check out <http://www.quepublishing.com>, where you will find a whole section dedicated to Cisco certifications and networking in general. This book will review what you need to know before you take the test, with the fundamental purpose dedicated to reviewing the information needed on the Cisco CCENT exam.

This book uses a variety of teaching and memorization techniques to analyze the exam-related topics and to provide you with ways to input, index, and retrieve everything you need to know to pass the test. Once again, it is not a comprehensive treatise on Cisco networking.

What This Book Is Designed to Do

This book is designed to be read as a pointer to the areas of knowledge you will be tested on. In other words, you might want to read the book one time, just to get an insight into how comprehensive your knowledge of networking with Cisco is. The book is also designed to be read shortly before you go for the actual test and to give you a distillation of the entire field of CCENT knowledge in as few pages as possible. We think you can use this book to get a sense of the underlying context of any topic in the chapters—or to skim read for Exam Alerts, bulleted points, summaries, and topic headings.

We've drawn on material from the Cisco listing of knowledge requirements, from other preparation guides, and from the exams themselves. We've also drawn from a battery of third-party test-preparation tools and technical websites, as well as from our own experience with Cisco equipment and the exam. Our aim is to walk you through the knowledge you will need—looking over your shoulder, so to speak—and point out those things that are important for the exam (Exam Alerts, practice questions, and so on).

The CCENT exam makes a basic assumption that you already have a strong background of experience with the general networking and its terminology. However, because the CCENT is an introductory-level test, we've tried to demystify the jargon, acronyms, terms, and concepts.

About This Book

If you're preparing for the CCENT exam for the first time, we've structured the topics in this book to build upon one another. Therefore, the topics covered in later chapters might refer to previous discussions in earlier chapters.

CCENT Official Exam Topics

The following is the list of official exam topics for CCENT. Each CCENT chapter references a selection of the topics in this list.

Operation of IP Data Networks

- ▶ Recognize the purpose and functions of various network devices such as routers, switches, bridges, and hubs.
- ▶ Select the components required to meet a given network specification.
- ▶ Identify common applications and their impact on the network
- ▶ Describe the purpose and basic operation of the protocols in the OSI and TCP/IP models.
- ▶ Predict the data flow between two hosts across a network.
- ▶ Identify the appropriate media, cables, ports, and connectors to connect Cisco network devices to other network devices and hosts in a LAN

LAN Switching Technologies

- ▶ Determine the technology and media access control method for Ethernet networks.
- ▶ Identify basic switching concepts and the operation of Cisco switches:
 - ▶ Collision domains
 - ▶ Broadcast domains
 - ▶ Types of switching
 - ▶ CAM table
- ▶ Configure and verify initial switch configuration, including remote access management:
 - ▶ Cisco IOS commands to perform basic switch setup
- ▶ Verify network status and switch operation using basic utilities such as ping, Telnet, and SSH.

- ▶ Describe how VLANs create logically separate networks and the need for routing between them:
 - ▶ Explain network segmentation and basic traffic management concepts
- ▶ Configure and verify VLANs.
- ▶ Configure and verify trunking on Cisco switches:
 - ▶ DTP
 - ▶ Autonegotiation

IP Addressing (IPv4 / IPv6)

- ▶ Describe the operation and necessity of using private and public IP addresses for IPv4 addressing.
- ▶ Identify the appropriate IPv6 addressing scheme to satisfy addressing requirements in a LAN/WAN environment.
- ▶ Identify the appropriate IPv4 addressing scheme using VLSM and summarization to satisfy addressing requirements in a LAN/WAN environment.
- ▶ Describe the technological requirements for running IPv6 in conjunction with IPv4 such as dual stack.
- ▶ Describe IPv6 addresses:
 - ▶ Global unicast
 - ▶ Multicast
 - ▶ Link local
 - ▶ Unique local
 - ▶ eui 64
 - ▶ Autoconfiguration

IP Routing Technologies

- ▶ Describe basic routing concepts:
 - ▶ CEF
 - ▶ Packet forwarding
 - ▶ Router lookup process

- ▶ Configure and verify utilizing the CLI to set basic router configuration:
 - ▶ Cisco IOS commands to perform basic router setup
- ▶ Configure and verify operation status of an Ethernet interface.
- ▶ Verify router configuration and network connectivity:
 - ▶ Cisco IOS commands to review basic router information and network connectivity
- ▶ Configure and verify routing configuration for a static or default route given specific routing requirements.
- ▶ Differentiate methods of routing and routing protocols:
 - ▶ Static vs. dynamic
 - ▶ Link state vs. distance vector
 - ▶ Next hop
 - ▶ IP routing table
 - ▶ Passive interfaces
- ▶ Configure and verify OSPF (single area):
 - ▶ Benefit of single area
 - ▶ Configure OSPFv2
 - ▶ Configure OSPFv3
 - ▶ Router ID
 - ▶ Passive interface
- ▶ Configure and verify interVLAN routing (router on a stick):
 - ▶ Subinterfaces
 - ▶ Upstream routing
 - ▶ Encapsulation
- ▶ Configure SVI interfaces

IP Services

- ▶ Configure and verify DHCP (IOS router):
 - ▶ Configuring router interfaces to use DHCP

- ▶ DHCP options
- ▶ Excluded addresses
- ▶ Lease time
- ▶ Describe the types, features, and applications of ACLs:
 - ▶ Standard
 - ▶ Sequence numbers
 - ▶ Editing
 - ▶ Extended
 - ▶ Named
 - ▶ Numbered
 - ▶ Log option
- ▶ Configure and verify ACLs in a network environment:
 - ▶ Named
 - ▶ Numbered
 - ▶ Log option
- ▶ Identify the basic operation of NAT:
 - ▶ Purpose
 - ▶ Pool
 - ▶ Static
 - ▶ 1 to 1
 - ▶ Overloading
 - ▶ Source addressing
 - ▶ One-way NAT
- ▶ Configure and verify NAT for given network requirements.
- ▶ Configure and verify NTP as a client.

Network Device Security

- ▶ Configure and verify network device security features such as:
 - ▶ Device password security

- ▶ Enable secret vs. enable
- ▶ Transport
- ▶ Disable telnet
- ▶ SSH
- ▶ VTYS
- ▶ Physical security
- ▶ Service password
- ▶ Describe external authentication methods
- ▶ Configure and verify switch port security features such as:
 - ▶ Sticky MAC
 - ▶ MAC address limitation
 - ▶ Static/dynamic
 - ▶ Violation modes
 - ▶ Err disable
 - ▶ Shutdown
 - ▶ Protect restrict
 - ▶ Shutdown unused ports
 - ▶ Err disable recovery
 - ▶ Assign unused ports to an unused VLAN
 - ▶ Setting native VLAN to other than VLAN 1
- ▶ Configure and verify ACLs to filter network traffic.
- ▶ Configure and verify ACLs to limit Telnet and SSH access to the router.

Troubleshooting

- ▶ Troubleshoot and correct common problems associated with IP addressing and host configurations.
- ▶ Troubleshoot and resolve VLAN problems:
 - ▶ Identify that VLANs are configured

- ▶ Port membership correct
- ▶ IP address configured
- ▶ Troubleshoot and resolve trunking problems on Cisco switches:
 - ▶ Correct trunk states
 - ▶ Correct encapsulation configured
 - ▶ Correct VLANs allowed
- ▶ Troubleshoot and resolve ACL issues:
 - ▶ Statistics
 - ▶ Permitted networks
 - ▶ Direction
 - ▶ Interface
- ▶ Troubleshoot and resolve Layer 1 problems:
 - ▶ Framing
 - ▶ CRC
 - ▶ Runts
 - ▶ Giants
 - ▶ Dropped packets
 - ▶ Late collision
 - ▶ Input/output errors

We suggest that you read this book from front to back. You won't be wasting your time because nothing we've written is a guess about an unknown exam. We've had to explain certain underlying information on such a regular basis that those explanations are included here.

After you've read the book, you can brush up on a certain area by using the Index or the Table of Contents to go straight to the topics and questions you want to reexamine. We've tried to use the headings and subheadings to provide outline information about each given topic. After you've been certified, we think you'll find this book useful as a tightly focused reference and an essential foundation of CCENT knowledge.

Chapter Formats

Each Exam Cram chapter follows a regular structure, along with graphical cues about especially important or useful material. The structure of a typical chapter is as follows:

- ▶ **Cram Savers and Cram Quizzes:** Responding to our readers' feedback, we have added more quizzes and answers. At the beginning of each major topic, you will find a Cram Saver. If these questions are no problem for you, you might be able to back off a little on how hard you cram that section. If you can't get them right, start cramming. Similarly, at the end of each major section, you will find a new Cram Quiz, which serves as an immediate check on your understanding and retention of that topic.
- ▶ **Topical coverage:** Each chapter covers topics drawn from the official exam topics as they relate to the chapter's subject.
- ▶ **Alerts:** Throughout the topical coverage section, we highlight material most likely to appear on the exam by using a special Exam Alert layout that looks like this:

ExamAlert

This is what an Exam Alert looks like. An Exam Alert stresses concepts, terms, software, or activities that will most likely appear in one or more certification exam question. For that reason, we think any information found offset in Exam Alert format is worthy of unusual attentiveness on your part.

Even if material isn't flagged as an Exam Alert, all the content in this book is associated in some way with test-related material. What appears in the chapter content is critical knowledge.

- ▶ **Notes:** This book is an overall examination of basic Cisco networking. Therefore, we dip into many aspects of Cisco networking. Where a body of knowledge is deeper than the scope of the book, we use notes to indicate areas of concern or specialty training, or refer you to other resources.

Note

Cramming for an exam will get you through a test, but it won't make you a competent IT professional. Although you can memorize just the facts you need to become certified, your daily work in the field will rapidly put you in water over your head if you don't know the underlying principles of networking with Cisco gear.

- ▶ **Tips:** We provide tips that will help you to build a better foundation of knowledge or to focus your attention on an important concept that will reappear later in the book. Tips provide a helpful way to remind you of the context surrounding a particular area of a topic under discussion.
- ▶ **Exam Prep Questions:** This section presents a short list of test questions related to the specific chapter topic. Each question has a following explanation of both correct and incorrect answers. The practice questions highlight the areas we found to be most important on the exam.

The bulk of the book follows this chapter structure, but there are a few other elements that we would like to point out:

- ▶ **Practice tests:** The practice tests are very close approximations of the types of questions you are likely to see on the current CCENT ICND1 100-101 exam. The answer key for each practice test will help you determine which areas you have mastered and which areas you need to study further.
- ▶ **Answer keys:** These provide the answers to the sample tests, complete with explanations of both the correct responses and the incorrect responses.
- ▶ **Glossary**—This is an extensive glossary of important terms used in this book.
- ▶ **Cram Sheet**—This appears as a tear-out sheet inside the front cover of this Exam Cram book. It is a valuable tool that represents a collection of the most difficult-to-remember facts and numbers we think you should memorize before taking the test. Remember, you can dump this information out of your head onto a piece of paper as soon as you enter the testing room. These are usually facts that we've found require brute-force memorization. You only need to remember this information long enough to write it down when you walk into the test room. Be advised that you will be asked to surrender all personal belongings before you enter the exam room itself.

You might want to look at the Cram Sheet in your car (no, not while you are driving!) or in the lobby of the testing center just before you walk into the testing center. The Cram Sheet is divided under headings so that you can review the appropriate parts just before each test.

- ▶ **CD-ROM:** The CD-ROM contains the Pearson IT Certification Practice Test engine, which provides multiple test modes that you can use for exam preparation. The practice tests are designed to appropriately balance the questions over each technical area (domain) covered by the exam.

Pearson IT Certification Practice Test Engine and Questions on the CD-ROM

The CD-ROM in the back of the book includes the Pearson IT Certification Practice Test engine—software that displays and grades a set of exam-realistic multiple-choice questions. Using the Pearson IT Certification Practice Test engine, you can either study by going through the questions in Study Mode or take a simulated exam that mimics real exam conditions. You can also serve up questions in a Flash Card Mode, which will display just the question and no answers, challenging you to state the answer in your own words before checking the actual answers to verify your work.

The installation process requires two major steps: installing the software and then activating the exam. The CD in the back of this book has a recent copy of the Pearson IT Certification Practice Test engine. The practice exam—the database of exam questions—is not on the CD.

Note

The cardboard CD case in the back of this book includes the CD and a piece of paper. The paper lists the activation code for the practice exam associated with this book. Do not lose the activation code. On the opposite side of the paper from the activation code is a unique, one-time use coupon code for the purchase of the Premium Edition eBook and Practice Test.

Install the Software from the CD

The Pearson IT Certification Practice Test is a Windows-only desktop application. You can run it on a Mac using a Windows virtual machine, but it was built specifically for the PC platform. The minimum system requirements are as follows:

- ▶ Windows XP (SP3), Windows Vista (SP2), Windows 7, or Windows 8
- ▶ Microsoft .NET Framework 4.0 Client
- ▶ Pentium class 1GHz processor (or equivalent)
- ▶ 512MB RAM
- ▶ 650MB disc space plus 50MB for each downloaded practice exam

The software installation process is pretty routine as compared with other software installation processes. If you have already installed the Pearson IT Certification Practice Test software from another Pearson product, there is no need for you to reinstall the software. Simply launch the software on your desktop and proceed to activate the practice exam from this book by using the activation code included in the CD sleeve.

The following steps outline the installation process:

1. Insert the CD into your PC.
2. The media interface that automatically runs allows you to access and use all CD-based features, including the exam engine and sample content from other Cisco self-study products. From the main menu, click the Install the Exam Engine option.
3. Respond to windows prompts as with any typical software installation process.

The installation process gives you the option to activate your exam with the activation code supplied on the paper in the CD sleeve. This process requires that you establish a Pearson website login. You will need this login to activate the exam, so please do register when prompted. If you already have a Pearson website login, there is no need to register again. Just use your existing login.

Activate and Download the Practice Exam

Once the exam engine is installed, you should then activate the exam associated with this book (if you did not do so during the installation process) as follows:

1. Start the Pearson IT Certification Practice Test software from the Windows Start menu or from your desktop shortcut icon.
2. To activate and download the exam associated with this book, from the My Products or Tools tab, click the Activate Exam button.
3. At the next screen, enter the activation key from paper inside the cardboard CD holder in the back of the book. Once entered, click the Activate button.
4. The activation process will download the practice exam. Click Next, and then click Finish.

Once the activation process is completed, the My Products tab should list your new exam. If you do not see the exam, make sure that you have selected the My

Products tab on the menu. At this point, the software and practice exam are ready to use. Simply select the exam and click the Open Exam button.

To update a particular exam you have already activated and downloaded, select the Tools tab and click the Update Products button. Updating your exams will ensure you have the latest changes and updates to the exam data.

If you want to check for updates to the Pearson IT Certification Practice Test engine software, simply select the Tools tab and click the Update Application button. This will ensure that you are running the latest version of the software engine.

Activating Other Exams

The exam software installation process, and the registration process, only has to happen once. Then, for each new exam, only a few steps are required. For instance, if you buy another new Pearson IT Certification Cert Guide, extract the activation code from the CD sleeve in the back of that book; you don't even need the CD at this point. From there, all you have to do is start the exam engine (if not still up and running), and perform steps 2 through 4 from the previous list.

Contacting the Authors

We've tried to create a real-world tool that you can use to prepare for and pass the CCNA certification exams. We're interested in any feedback you would care to share about the book, especially if you have ideas about how we can improve it for future test takers. We'll consider everything you say carefully and will respond to all reasonable suggestions and comments. You can reach us via email at examcram@mikevalentine.ca and keithb@hurry.ws.

Let us know if you found this book to be helpful in your preparation efforts. We'd also like to know how you felt about your chances of passing the exam before you read the book and then after you read the book. Of course, we'd love to hear that you passed the exam—and even if you just want to share your triumph, we'd be happy to hear from you.

Thanks for choosing us as your personal trainers, and enjoy the book. We would wish you luck on the exam, but we know that if you read through all the chapters and work with the product, you won't need luck—you'll pass the test on the strength of real knowledge!

Self-Assessment

This section helps you to determine your readiness for the Cisco Certified Entry Networking Technician certification exam. You will be invited to assess your own skills, motivations, education, and experience and see how you compare against the thousands of CCENT and CCNA candidates we have met.

Tip

You can also pre-assess your CCENT readiness by using the accompanying CD.

CCENT in the Real World

The Cisco Certified Entry Networking Technician remains one of the most popular certifications in the IT industry. Although Cisco does not publish certification statistics, it is safe to say that thousands of new CCENTs are minted each year from all over the world. In the face of a backlash against so-called paper-only certification holders, Cisco has worked hard to maintain the credibility of its certifications by making them difficult to achieve, in addition to ensuring that the exams test not only their own products and services but also general networking knowledge. In the past few years, Cisco has added more complex and involved questions and interfaces to computer-based tests to test the applied knowledge of candidates, and we can expect this trend to continue. A Cisco certification is still the gold standard for networking professionals.

A candidate who has passed the CCENT has demonstrated three significant capabilities:

- ▶ **A mastery of technical knowledge:** The successful CCENT candidate knows the technical material and has an elevated level of retention and accuracy. The CCENT exam has a pass mark of 849 out of 1,000. Very little room for technical error exists. Successful candidates know their stuff.

- ▶ **A demonstrated ability to apply the technical knowledge:** The addition of simulator questions has greatly reduced the possibility that a candidate can simply memorize all the information and pass the exam. A CCENT is supposed to be able to apply basic router and switch configurations; the simulator questions help prove that the candidate can do so.
- ▶ **The ability to perform under pressure:** The CCENT exam requires that you proceed at a fairly rapid pace, spending about 1 minute per question on average. Many candidates find that they have little time left when they finish, and indeed many run out of time altogether—and some fail as a result. Add to this the stress of being in an exam environment, the potential of having an employer’s performance expectations, personal expectations, and possibly financial or career implications pressuring you as well, and the exam turns into a stress test. All of this is intimidating—and unfortunate for the unprepared—so be prepared.

Imagine yourself as an employer looking for a junior networking professional. You want someone who knows their stuff, who can reliably do the actual work of setting up and configuring equipment, and who can do all that under the pressures of time, screaming bosses and customers, and critical deadlines. Enter the successful CCENT candidate.

The Ideal CCENT Candidate

Other than a photographic memory, typing speed that would make Mavis Beacon jealous, and nerves of steel, what makes for the “ideal” CCENT candidate? A combination of skills and experience is the short answer. The successful candidates we have seen—and we have seen thousands from classes that we have taught—had a good mix of the following traits:

- ▶ **Motivation:** Why are you taking the CCENT? Here are some of the most common answers to this question that we have seen:

Because I want to further my career and get a promotion.

To expand my knowledge; I’m interested in it.

My job is changing, and the company needs me to get the certification.

I am unemployed and/or starting a new career.

The company needs more Cisco-certified people to gain a certain partner status as a reseller.

We’re just burning the training budget for this year.

I've heard that the computer industry is a good field and that a CCENT guarantees you \$85,000 a year.

So what motivates you? Who is paying for the training and exams? What are the implications if you fail? Successful candidates are highly motivated. If you don't care, your chances of passing drop tremendously.

- ▶ **An interest in learning and an ability to learn:** Passing a CCENT exam requires taking on board a great deal of new information, much of it obscure and without a referential pattern to make it easier to recall. Candidates who have acquired the skills to do this—and rest assured, these are skills that can be learned—will do better than those who have trouble retaining information. Candidates who simply enjoy learning will find it easier and will do better as a result.

If you have trouble retaining and recalling information quickly and accurately, you will find CCENT certification a difficult thing to achieve. This book is not aimed at teaching you these skills; other books are. In the absence of the ability to learn and retain quickly, patience and persistence are a good substitute. If it takes you a year to pass, you have still passed.

- ▶ **A decent background in IP networking:** *Decent* is intentionally vague. We have seen candidates with little experience succeed and candidates with extensive experience fail. Experience is not a guarantee, but it absolutely helps. Many CCENT questions test the basics of networking; many others assume that you know the basics and incorporate the requirement of that knowledge into a more advanced question—the old “question-within-a-question” trick. As a guideline, if you have been involved with business-class networks for about a year, you will probably have absorbed enough knowledge to give you an advantage when it comes to the basics. After a certain point, experience can be a weakness: In the immortal words of Han Solo, “Don't get cocky.” If you think that CCENT will be easy because you have 10 years of experience, you are in for a rude awakening.

Put Yourself to the Test

Now is the time to take a close look at your education, experience, motivation, and abilities. It's worth being honest with yourself; being aware of your weaknesses is as important as being aware of your strengths. Maybe you know someone who can help you with an objective assessment—a friend, a teacher, or an HR person perhaps. Above all, realize that the following questions and comments simply summarize our experience with CCENT candidates. That

experience is pretty solid; we have taught CCENT to more than a thousand people. By the same token, though, there is no magic formula; every person is a different story. Your best plan is to be as prepared as you can be in all respects. Now, time to look inward.

Educational Background

Although in theory anyone can attempt the CCENT exam, in reality some are better prepared than others. Educational background forms a big part of this preparation. These questions will help to identify education and training that will be of benefit:

1. Have you ever taken any computer science courses at a college level?

Most college-level IT courses include an element of networking theory. Also, if you are taking this kind of course, you are probably already interested in this topic and will find it easier to master the basics and pick up the advanced stuff. If you have never taken an IT course at this level, you have a steeper learning curve and might be at a disadvantage.

2. Did you attend college and major in a computer-related field?

If so, you should have most of the basics covered—unless you studied programming; in which case, you might not have covered much in the way of networking. Some colleges actually offer the CCENT as part of the curriculum. Doing a college major in IT is not a prerequisite by any means, but it might be helpful.

3. Have you ever held an IT certification?

If you have been certified before, you have some idea of what is coming in terms of the depth of knowledge required and the examination process; it also implies at least some involvement in computers and networking.

4. Which certification(s) have you held?

A previous CCNA or CCENT will definitely be an asset—but not a guarantee. At the time of this writing, the CCENT recently changed dramatically. Previous certification in general networking (perhaps a Net+), or an MCSE, will cover the basics, but not the Cisco-specific information. On the flip side, a certification in Visual Basic or Oracle might not be very helpful for CCENT.

5. Do you currently hold any IT certifications?

Current information is more relevant—especially in the IT world. Some certifications are more relevant than others, of course, as noted previously.

6. Which certification(s) do you currently hold?

You might hold other Cisco specialization certifications, or current certifications from Microsoft, CompTIA, or Novell. Again, anything that has tested your networking knowledge will be an asset.

7. Have you ever taken any IT training courses in networking?

Many people take training courses but do not certify. Any exposure and knowledge gained from these courses will be useful.

8. How much self-study have you done?

Although it is difficult to do pure self-study and pass the new CCENT, the more you study, the better the chances are that you will retain information. In our experience, it is always more productive to get some training—whether online, with a mentor/tutor, or from a training company—but a significant amount of self-study is always required regardless. The fact that you are holding this book is a very good sign. Read all of it!

9. How long have you been studying for your CCENT?

This is a tricky equation. The longer you study, the more you are likely to know—but the more you are likely to forget, as well.

10. Is there a formal or informal training plan for you at your workplace?

Work experience is a great way to gain the knowledge and skills you need for the exam. A training plan can be a good motivator because you might have someone coaching and encouraging you and also because there may be a reward—perhaps a promotion or raise—for completing the program.

Hands-On Experience

It is the rare individual who really understands networks but has never built, broken, and then rebuilt one. For the CCENT exam, a certain amount of hands-on experience is a must. The simulator questions require you to actually type in router or switch configurations. Ask yourself the following:

- 1.** Does your job allow you to work with Cisco routers and switches on a regular basis?
- 2.** Is there a lab where you can practice? Perhaps at home with borrowed or purchased gear?
- 3.** How long have you been working with Cisco equipment?
- 4.** Are you completely fluent in subnetting?

At a minimum, you should get a simulator that includes lab exercises for you to practice key skills. If you have access to a lab and equipment you can play with, as you become more advanced, you can build more complicated and realistic test networks.

The major skill areas you need hands-on experience in are as follows:

- ▶ Basic configuration: IP addresses, passwords
- ▶ Subnetting
- ▶ Dynamic routing protocol configuration
- ▶ NAT/PAT (Network/Port Address Translation)
- ▶ Basic WAN protocols and configuration
- ▶ Switching, VLANs, VLAN Trunking Protocol, trunking
- ▶ IP access lists
- ▶ Troubleshooting

As you think about those areas, picture yourself in front of a Cisco business-class router or switch and assess your level of confidence in being able to quickly and correctly configure it. You should feel no intimidation or uncertainty in being able to tackle these kinds of configurations. Subnetting is emphasized and is one of the main areas where people have difficulty. You must be totally, unequivocally confident with subnetting or you will face a serious challenge on your exam. Even more than subnetting, troubleshooting has become a significant emerging theme in CCENT exams, and we have included a chapter devoted entirely to that topic.

Testing Your Exam Readiness

The CCENT exam will demand a high degree of technical accuracy, applied skill, and the ability to perform quickly under pressure. You can give yourself experience in this environment by practicing on an exam simulator until you are comfortable. You must become technically accurate to about 90% to 95%, have no difficulty with the simulator tasks, and be able to complete the exam in the appropriate time frame. This can be achieved by repetition, but be careful that you do not simply memorize all the questions in the test pool!

Assessing Your Readiness for the CCENT Exam

There are three “pillars” of success on the CCENT exam: technical excellence, applied skills, and the ability to perform under pressure. Technical excellence

is achieved with study, training, and self-testing. Applied skills are learned through practice labs and exams, work experience, and hands-on training and experience. The ability to perform under pressure is gained from situational training such as exam simulators and challenge labs, perhaps with a trainer or mentor. The goal is to increase your confidence level so that you feel as if you own the material and want to be challenged to a duel by the exam.

With a combination of educational and work experience, CCENT-specific training, self-study, and hands-on practice, you will put yourself in the best position to approach the exam with a high degree of confidence—and pass. Good luck; study hard.

This page intentionally left blank

CHAPTER THREE

Concepts in IP Addressing

This chapter covers the following official ICND1 100-101 exam topics:

- ▶ Describe the operation and necessity of using private and public IP addresses for IPv4 addressing.
- ▶ Identify the appropriate IPv4 addressing scheme using VLSM and summarization to satisfy addressing requirements in a LAN/WAN environment.
- ▶ Describe IPv6 addresses.
- ▶ Identify the appropriate IPv6 addressing scheme to satisfy addressing requirements in a LAN/WAN environment.
- ▶ Describe the technological requirements for running IPv6 in conjunction with IPv4 such as dual stack.

The exam requires a perfect fluency in subnetting. Success requires speed and accuracy in answering the many questions you will see on this topic. The key to this level of fluency is practice—you must work at your skills until they become second nature.

The following sections discuss binary and hexadecimal numbering systems as compared with the more familiar decimal system. An understanding of binary, in particular, is crucial to success on the test as it is fundamental to computer systems in general, and to topics such as subnetting, access lists, routing, and route summarization. This chapter also looks at the relationship between the IP address and subnet mask in more detail, as well as how it can be manipulated for more efficient network functionality using route summarization. Finally, IP Version 6 is introduced and some of its features are explained.

Binary

CramSaver

1. What are the eight binary values found in a single octet of an IP address?
 - A. 256 128 64 32 16 8 4 2
 - B. 254 62 30 14 6 4 0
 - C. 128 64 32 16 8 4 2 1
 - D. 0 2 4 6 8 10 12 14
2. What is the decimal value of binary 00001110 ?
 - A. 13
 - B. 14
 - C. 1,110
 - D. 16
 - E. 15
3. What is the binary value of decimal 256?
 - A. 11111111
 - B. 1111111111
 - C. 100000000
 - D. 10000000

Answers

1. Answer C is correct.
2. Answer B is correct.
3. Answer C is correct.

Binary is the language of digital electronic communication. Binary is another name for Base 2 numbering. Our usual numbering system is Base 10, in which a single character or column can represent one of 10 values: 0, 1, 2, 3, 4, 5, 6, 7, 8, or 9. The first column indicates how many ones there are in a given value. To represent a value greater than 9, we need another column, which represents how many “tens” there are; if the value we want to represent is greater than 99, we use another column for the “hundreds,” and so on. You might notice that each additional column is ten times greater than the preceding one: ones, tens,

hundreds, thousands, and so forth—all “powers of 10”: 10⁰, 10¹, 10², 10³, and so on. Base 10 is easy because most of us have 10 fingers and have known how to count from an early age.

In binary, or Base 2, a single character or column can represent one of only two values: 0 or 1. The next column represents how many “twos” there are; the next column how many “fours,” and so on. You’ll notice here that the value of each additional column is two times greater than the previous—all “powers of 2”: 2⁰, 2¹, 2², 2³, and so on. This is not a coincidence.

Given that a Base 2 or binary column can have only two possible values (0 or 1), this makes it easy to represent a binary value as an electrical value: either off (0) or on (1). Computers use binary because it is easily represented as electrical signals in memory or digital values on storage media. The whole system works because computers are quick at computing arithmetic, and as you’ll learn, pretty much all computer operations are really just fast binary math.

Let’s take a look at some Base 10 (or decimal) to binary conversions. Take the decimal number 176. Those three digits tell us that we have one 100, plus seven 10s, plus six 1s. Table 3.1 illustrates how decimal numbers represent this distribution of values.

TABLE 3.1 **Decimal Values**

100,000s	10,000s	1000s	100s	10s	1s
0	0	0	1	7	6

Notice that we have some 0s in the high-value columns; we can drop those from the beginning if we want to. You will not have to analyze decimal numbers in this way on the exam; we are simply demonstrating how Base 10 works so it can be compared to Base 2 and Base 16 in the same way.

In binary, the columns have different values—the powers of 2. Table 3.2 lists the values of the lowest 8 bits in binary.

TABLE 3.2 **Binary Values**

128	64	32	16	8	4	2	1

Note

The biggest values in a binary string (the 1s at the left) are often called the *high-order* bits because they have the highest value. Similarly, the lowest-value bits at the right are referred to as the *low-order* bits.

Tip

You must know the value of each binary bit position! If you have difficulty memorizing them, try starting at 1 and keep doubling as you go to the left.

To represent the decimal number 176 in binary, we need to figure out which columns (or bit positions) are “on” and which are “off.” Now, because this is arithmetic, there are a few different ways to do this.

Start with the decimal number you want to convert:

176

Next, look at the values of each binary bit position and decide if you can subtract the highest column value and end up with a value of 0 or more. Ask yourself: “Can I subtract 128 from 176?” In this case, $176 - 128 = 48$.

Yes, you can subtract 128 from 176 and get a positive value, 48. Because we “used” the 128 column, we put a 1 in that column, as shown in Table 3.3.

TABLE 3.3 Building a Binary String, Part 1

128	64	32	16	8	4	2	1
1							

Now, we try to subtract the next highest column value from the remainder. We get $176 - 128 = 48$. We take the 48 and subtract 64 from it.

Notice that you can’t do this without getting a negative number; this is not allowed, so we can’t use the 64 column. Therefore, we put a 0 in that column, as shown in Table 3.4.

TABLE 3.4 Building a Binary String, Part 2

128	64	32	16	8	4	2	1
1	0						

Move along and do the math for the rest of the columns: $48 - 32 = 16$. We then subtract 16 from 16 and get 0.

Note that when you get to 0, you are finished—you need to only fill the remaining bit positions with 0s to complete the 8-bit string. So, we used only the 128 column, the 32 column, and the 16 column. Table 3.5 is what we end up with.

TABLE 3.5 **Completed Binary Conversion**

128	64	32	16	8	4	2	1
1	0	1	1	0	0	0	0

176 decimal = 10110000 binary.

If you add up $128 + 32 + 16$, you get 176. That is how you convert from binary to decimal: Simply add up the column values where there is a 1.

ExamAlert

You will see several questions on converting from decimal to binary and back, so prepare accordingly.

Hexadecimal

The CCENT exam will ask you a few questions on the conversion of binary to hexadecimal and back, so you need to understand how it works. An understanding of hex is also a useful skill for other areas of networking and computer science.

Binary is Base 2; decimal is Base 10; hexadecimal is Base 16. Each column in hex can represent 16 possible values, from 0 through 15. To represent a value of 10 through 15 with a single character, hex uses the letters A through F. It is important to understand that the values of 0 through 15 are the possible values of a 4-bit binary number, as shown in Table 3.6.

TABLE 3.6 **Decimal, Binary, and Hex Values Compared**

Decimal	Binary	Hex
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F

ExamAlert

You should be able to reproduce Table 3.6 as a quick reference for the exam.

Conversion Between Binary, Hex, and Decimal

The following sections provide an introduction to converting between binary, hex, and decimal. Again, there is more than one mathematical approach to finding the correct answer, but the method shown is simple and reliable.

Decimal to Hexadecimal Conversions

The easiest way to get from decimal to hexadecimal and back is to go through binary. Take the example we used earlier in which we converted 176 decimal to binary:

$$176 = 10110000$$

Given that a single hex character represents four binary bits, all we need to do is to break the 8-bit string 10110000 into two 4-bit strings like this:

$$1011 \quad 0000$$

Now, simply match the 4-bit strings to their hex equivalent:

$$1011 = B$$

$$0000 = 0$$

The answer is simply $10110000 = 0xB0$.

The $0x$ in front of the answer is an expression that means “the following is in hex.” This is needed because if the hex value was 27, we could not distinguish it from 27 decimal.

Hexadecimal to Decimal Conversions

The reverse of the procedure is easier than it seems, too. Given a hex value of $0xC4$, all we need to do is to first convert to binary, and then to decimal.

To convert to binary, take the two hex characters and find their binary value:

$$C = 1100$$

$$0100 = 4$$

Now, make the two 4-bit strings into one 8-bit string:

$$11000100$$

Finally, add the bit values of the columns where you have a 1:

$$128 + 64 + 4 = 196$$

ExamAlert

It is critical to polish your skills in binary. You must be confident and quick in conversions, and the better your understanding of binary, the easier subnetting and other advanced IP topics will be for you. Practice, practice, practice!

Cram Quiz

1. Write the following binary IP in dotted-decimal format:
11000000.10101000.00000001.11111110.
2. Write the following subnet mask in binary format: 255.255.255.240.
3. What is 127.0.0.1 in binary?

Cram Quiz Answers

1. 192.168.1.254
 2. 11111111.11111111.11111111.11110000
 3. 01111111.00000000.00000000.00000001
-

IP Address Components

CramSaver

1. What class of IP address is 191.168.1.0?
2. What is the range (in decimal) of Class B addresses?
3. What is the range of private Class A addresses?
4. Is the address 172.16.1.0/24 subnetted?

Answers

1. Class B.
2. 128.0.0 .0 to 191.255.255.255
3. 10.0.0.0 to 10.255.255.255
4. Yes. The default mask is /16; /24 is longer so the address is subnetted.

CCNA candidates need to be fluent in their understanding of IP addressing concepts. The following sections detail how IP addresses are organized and analyzed, with a view to answering subnetting questions.

Address Class

Early in the development of IP, RFC 791 designated five classes of IP address: A, B, C, D, and E. These classes were identified based on the pattern of high-order bits (the high-value bits at the beginning of the first octet). The result is that certain ranges of networks are grouped into classes in a pattern based on the binary values of those high-order bits, as detailed in Table 3.7.

TABLE 3.7 **Address Class and Range**

Class	High-Order Bits	1st Octet Range
A	0	1–127
B	10	128–191
C	110	192–223
D	1110	224–239
E	11110	240–255

You might notice that 127 is missing. This is because at some point the address 127.0.0.1 was reserved for the loopback (sometimes called *localhost*) IP—this is the IP of the TCP/IP protocol itself on every host machine.

ExamAlert

You absolutely must be able to identify the class of an address just by looking at what number is in the first octet. This is critical to answering subnetting questions.

Public and Private IP Addresses

As the popularity of TCP/IP increased, many organizations wanted to use it in their own networks, without paying to be connected to the Internet. The IETF published a Recommended Best Practice (RFC 1918) that defined several “private” IP networks that could be used by individuals, corporations, or other organizations without needing to pay to lease them from an Internet service provider (ISP). The tradeoff was that these networks were not routable on the Internet because these addresses were filtered from the route tables on Internet routers and so were unreachable from the Internet. They work exactly the same as any other IP address in every other respect; in fact, if they weren’t filtered they would work on the Internet too.

These address ranges will probably be familiar to you already, because every LAN you have ever joined likely used one of them. To access the Internet from a LAN using one of these private, non-routable networks, we have to go through a Network Address Translation (NAT) router. NAT is explained in Chapter 10, “IP Services.” Table 3.8 lists the private IP Address ranges. You should commit these to memory and be able to recognize them at a glance.

TABLE 3.8 Private IP Address Ranges

Class	Range
A	10.0.0.0 to 10.255.255.255
B	172.16.0.0 to 172.31.255.255
C	192.168.0.0 to 192.168.255.255

Default Subnet Mask

Each class of address is associated with a default subnet mask, as shown in Table 3.9. An address using its default mask defines a single IP broadcast

domain—all the hosts using that same network number and mask can receive each other's broadcasts and communicate via IP.

TABLE 3.9 **Address Class and Default Masks**

Class	Default Mask
A	255.0.0.0
B	255.255.0.0
C	255.255.255.0

One of the rules that Cisco devices follow is that a subnet mask must be a contiguous string of 1s followed by a contiguous string of 0s. There are no exceptions to this rule: A valid mask is always a string of 1s, followed by 0s to fill up the rest of the 32 bits. (There is no such rule in the real world, but we will stick to the Cisco rules here; it's a Cisco exam, after all.)

Therefore, the only possible valid values in any given octet of a subnet mask are 0, 128, 192, 224, 240, 248, 252, 254, and 255. Any other value is invalid.

ExamAlert

You should practice associating the correct default subnet mask with any given IP address; this is another critical skill in subnetting.

The Network Field

Every IP address is composed of a network component and a host component. The subnet mask has a single purpose: to identify which part of an IP address is the network component and which part is the host component. Look at a 32-bit IP address expressed in binary, with the subnet mask written right below it. Figure 3.1 shows an example.

IP Address and Mask: 192.168.0.96 255.255.255.0

Binary IP: 11000000.10101000.00000000.01100000

Binary Mask: 11111111.11111111.11111111.00000000

Network
Field

 Host
Field

FIGURE 3.1 IP address and mask in binary, showing network and host fields.

Anywhere you see a binary 1 in the subnet mask, it means “the matching bit in the IP address is part of the network component.” In this example, the network part of the address is 192.168.0.X, and the last octet (X) will be the host component.

Because there are 24 bits in a row in the mask, we can also use a shortcut for the mask notation of /24. These examples show how a dotted-decimal mask can be expressed in slash notation:

192.168.1.66 255.255.255.0 = 192.168.1.66 /24

172.16.0.12 255.255.0.0 = 172.16.0.12 /16

10.1.1.1 255.0.0.0 = 10.1.1.1 /8

This slash notation is sometimes called CIDR (classless interdomain routing) notation. For some reason, it’s a concept that confuses students, but honestly it’s the easiest concept of all: The slash notation is simply the number of 1s in a row in the subnet mask. The real reason to use CIDR notation is simply that it is easier to say and especially to type—and it appears interchangeably with dotted-decimal throughout the exam. CIDR notation also appears in the output of various IOS commands.

Every IP address has a host component and a network component, and the 1s in the mask tell us which bits in the address identify the network component.

The Host Field

If the 1s in the mask identify the network component of an address, the 0s at the end of the mask identify the host component. In the preceding example, the entire last octet is available for the host IP number.

The number of 0s at the end of the mask mathematically define how many hosts can be on any given network or subnet. The 1s in the mask always identify the network component, and the 0s at the end of the mask always identify the host component of any IP address.

Nondefault Masks

At this point, you should be able to recognize what class an address belongs to, and what its default mask is supposed to be. Here’s the big secret: If a mask is longer than it is supposed to be, that network has been subnetted. So, it is clearly another critical skill that you be able to spot those nondefault masks.

The Subnet Field

Because we have extended the subnet mask past the default boundary into the bits that were previously host bits, we identify the bits we “stole” from the host part as the subnet field. The subnet field is relevant because those bits mathematically define how many subnets we create. Figure 3.2 uses the same IP address from our previous example, but now we have applied a mask that is longer than the default. Note that this creates the subnet field.

IP Address and Mask: 192.168.0.96 255.255.255.192

Binary IP: 11000000.10101000.00000000.01100000

Binary Mask: 11111111.11111111.11111111.11000000




FIGURE 3.2 IP address and non-default mask in binary illustrating the subnet field.

Figure 3.2 identifies the two extra bits past the default boundary as the subnet field—they used to be in the host field, but we subnetted and stole them to become the subnet field.

Cram Quiz

1. If the mask assigned to a private Class C address is 24 bits, is the address subnetted?
2. Which of the following are private IP addresses that can be assigned to a host?
 - A. 12.17.1.45
 - B. 10.255.255.254
 - C. 172.15.255.248
 - D. 192.168.1.5
 - E. 239.0.0.1
3. Why can't Duncan assign the address of 17.21.12.1111 to his Internet web server?

Cram Quiz Answers

1. No. /24 is the default mask, so the address is not subnetted.
 2. Answers B and D are correct. The other addresses are not in the private unicast ranges.
 3. Because it is not a valid IP address: .1111 in any octet is not a valid IP.
-

Subnetting

CramSaver

1. How many hosts are on the network 172.16.41.0/27?
 - A. 65,534
 - B. 32
 - C. 254
 - D. 30
 - E. 27
 - F. 14

2. How many subnets are created by the address 192.168.1.0 255.255.255.248?
 - A. 1
 - B. 2
 - C. 4
 - D. 8
 - E. 16
 - F. 32
 - G. 64

3. What is the broadcast ID of the seventh subnet created using 172.16.0.0/28?
 - A. 172.16.111.0
 - B. 172.16.0.0
 - C. 172.16.0.7
 - D. 172.16.0.96
 - E. 172.16.0.110
 - F. 172.16.0.111
 - G. 172.16.0.112

Answers

1. Answer D is correct. Five 0s at the end of the mask; $(2^5) - 2 = 30$
2. Answer F is correct. 5 bits were stolen to extend the mask: $(2^5) = 32$
3. Answer F is correct. The seventh subnet ranges from the network ID of 172.16.0.96 to the broadcast ID of 172.16.0.111.

Subnetting is not as difficult as it initially seems. Because we are dealing with arithmetic, there is definitely more than one way to do this, but the method shown here has worked well. The following sections work through the process of subnetting. Then, we work on some shortcuts to show how you can subnet quickly because CCNA exam candidates often find that they are pressed for time on the exam.

Address Class and Default Mask

Subnetting happens when we extend the subnet mask past the default boundary for the address we are working with. So it's obvious that we first need to be sure of what the default mask is supposed to be for any given address. Previously, we looked at the RFC791 designations for IP address classes and the number ranges in the first octet that identify those classes. If you didn't pick up on this before, you should memorize those immediately.

When faced with a subnetting question, the first thing to do is decide what class the address belongs to. Here are some examples:

192.168.1.66

The first octet is between 192 and 223: Class C

Default mask for Class C: 255.255.255.0

188.21.21.3

The first octet is between 128 and 191: Class B

Default mask for Class B: 255.255.0.0

24.64.208.5

The first octet is between 1 and 126: Class A

Default mask for Class A: 255.0.0.0

It's important to grasp that if an address uses the correct default mask for its class, it is not subnetted. This means that regardless of how many hosts the 0s at the end of the mask create, all those hosts are on the same network, all in the same broadcast domain. This has some implications for classful networks (ones that use the default mask for the address). Take a Class A for example: A Class A network can have 16,777,214 hosts on it. Almost 17 million PCs on one network would never work—there would be so much traffic from broadcasts alone, never mind regular data traffic, that nothing could get through and the network would collapse under its own size. Even a Class B network has 65,534 possible host IPs. This is still too many. So, either we waste a lot of addresses by not using the whole classful A or B network, or we subnet to make the networks smaller.

This is actually one of the most common reasons we subnet: The default or classful networks are too big, causing issues such as excessive broadcast traffic and wasted IP address space. Subnetting creates multiple smaller subnetworks out of one larger classful network, which allows us to make IP networks the “right” size—big or small—for any given situation.

The Increment

By definition, the process of subnetting creates several smaller classless subnets out of one larger classful one. The size of these subnets, or how many IP addresses they contain, is called the increment. Because we are working with binary numbers, a pattern emerges in which the increment is always one of those powers of 2 again—another good reason to memorize those numbers.

The increment is really easy to figure out. It is simply the value of the last 1 in the subnet mask. Let’s look at some examples. Figure 3.3 shows an IP address and subnet mask in binary.

IP Address and Mask: 192.168.21.1 255.255.255.0

Binary IP: 11000000.10101000.00010101.00000001

Binary Mask: 11111111.11111111.11111111.00000000

FIGURE 3.3 IP address and mask in binary.

Note that this is a Class C address, and it uses the correct default mask—so it is not subnetted. This means that there is only one network, so there isn’t really an increment to worry about here. It’s sufficient at this point to recognize that an address that uses its default mask creates one network (no subnets), so there is no subnetted increment to consider.

Let’s take the same address and subnet it by extending the mask past the default boundary, as shown in Figure 3.4.

IP Address and Mask: 192.168.21.1 255.255.255.192

Binary IP: 11000000.10101000.00010101.00000001

Binary Mask: 11111111.11111111.11111111.11000000




FIGURE 3.4 IP address and subnetted mask.

The very last 1 in the subnet mask in the figure is in the bit position worth 64—so the increment in this case is 64, which means that the subnets we made are evenly spaced at 64 IP addresses apart.

Think about this for a second. We are doing the subnetting in the fourth octet—that is where the mask changes from 1s to 0s. (The octet where this happens is sometimes referred to as the *interesting* octet.) The lowest possible value in that fourth octet is 0. If the subnets are 64 IP addresses apart, this means that the first subnet starts at 0, the next one starts at 64, the third at 128, and the fourth at 192—all multiples of the increment. Note that if we add another 64 to that last 192, we get 256—and that is larger than 255, the largest value that is possible in one octet. So this means we only have room for four subnets. Figure 3.5 illustrates this pattern more clearly.

.0	.64	.128	.192
.1	.65	.129	.193
.	.	.	.
.	.	.	.
.	.	.	.
.62	.126	.190	.254
.63	.127	.191	.255

FIGURE 3.5 Subnets created with increment of 64.

The multiples of the increment—0, 64, 128, and 192—are the starting addresses of the subnets we created. The subnets are all 64 addresses long, so we have room to make four subnets before we run out of addresses in the fourth octet.

Figure 3.6 shows our IP and subnet mask—note that the value of the last “1” in the mask is 16—and the subnets created with that increment of 16.

192.168.21.0 255.255.255.240

IP: 11000000.10101000.00010101.00000000
Mask: 11111111.11111111.11111111.11110000

Subnets Created with Increment of 16:

.0	.16	.32	.48	.64	.80	.96	.112	.128	.144	.160	.176	.192	.208	.224	.240
.
.
.
.15	.31	.47	.63	.79	.95	.111	.127	.143	.159	.175	.191	.207	.223	.239	.255

FIGURE 3.6 IP address and subnet mask with increment of 16.

First of all, you should notice that we are subnetting again—the mask extends past the default boundary. The last 1 in the mask is in the bit position worth 16, so our increment is 16. The multiples of 16 are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. Again, we can't make another subnet because $240 + 16 = 256$. Be careful not to start doubling as we did with the binary values; here we are just adding the increment value each time. It's easy to get confused!

The increment is really the key to subnetting; if you can determine the increment, you can see how big your subnets are and how many you have created. Remember, the easy way to find the increment is to just determine the bit value of the last 1 in the mask.

Number of Hosts

The number of 0s at the end of the mask always defines the number of hosts on any network or subnet. There is a simple mathematical formula that defines how many IP addresses are available to be assigned to hosts.

Note

Hosts is another word for computers, router interfaces, printers, or any other network component that can be assigned an IP address.

Now, no one expects you to be a big fan of algebra, but you need to see and understand the formula.

The number of binary bits you have to use determines the maximum number of different values you can express using those bits. If you have 3 bits, you can make eight different values—0 through 7, or 000 through 111 in binary; 3 bits, and $2^3 = 8$ —this is not a coincidence. The binary values you learned earlier—1, 2, 4, 8, 16, 32, 64, and 128—are all powers of 2 and define the maximum number of different values you can create if the mask ends in that bit position. So it should come as no surprise that the formula for the number of hosts on any network or subnet is $2^H - 2$, where H is the number of 0s at the end of the mask.

But why do we subtract 2 in the formula? It's pretty straightforward: Every network or subnet has two reserved addresses that cannot be assigned to a host. The rule is that no host can have the IP address in which all the host bits are set to 0, and no host can have the IP address in which all the host bits are set to 1. These addresses are called the network ID and the broadcast ID, respectively. They are the first and last IP addresses in any network or subnet. We lose those two IP addresses from the group of values that could be assigned to hosts.

Think of a network or subnet as a street with houses on it. Each house has a unique address, and the street has a name. The network ID is like the street name, and all the houses are hosts on a subnet that is known by its network ID street name. If two hosts have identical network and subnet fields in their addresses, they are on the same network, and can ping each other and exchange data and all that good stuff. If the network and subnet fields are different, even by 1 bit, they are on different networks and can't communicate until we put a router between them. The routers act like street intersections; you must get to the right intersection (router) before you can get on to the street you want (but we'll save that for later).

In a network where there are no routers, devices running TCP/IP make a decision about whether a particular IP address is on the network by performing a logical AND operation. The AND is a Boolean function that works like this:

```
1 AND 1 = 1
0 AND 1 = 0
1 AND 0 = 0
0 AND 0 = 0
```

This operation applies to IP networking like this: A host does a logical AND between its own IP and its mask. This determines its network ID. The host can then do an AND between another IP address and its own mask to determine if that second address is on the same network or some other one.

Let's take the IP address and mask of an imaginary host and display them in binary, as shown in Figure 3.7. The AND operation takes each bit in the address and ANDs it with the corresponding bit in the mask below it; the result is the network ID of the host.

```
IP Address and Mask: 192.16.20.12 255.255.255.0

Binary IP:      11000000.00010000.00010100.00001100
Binary Mask:    11111111.11111111.11111111.00000000
AND Result:     11000000.00010000.00010100.00000000

NetID = 192.16.20.0
```

FIGURE 3.7 The AND operation determines the network ID.

Now the host knows its own network ID and can compare any other host's address to that to see if the other host has the same network ID. If the two network IDs are different, traffic has to be sent through a router to get to the other network—and if there is no router, the two hosts can't communicate.

ExamAlert

Being able to do the AND operation is a useful skill. A lot of test questions center around the network ID, and being able to find it quickly is a big help.

The Broadcast ID

The broadcast ID is the address that represents every host on that network or subnet. Sometimes called a directed broadcast, it is the common address of all hosts on that network ID. This should not be confused with a full IP broadcast to the address of 255.255.255.255, which hits every IP host that can hear it; the broadcast ID hits only hosts on a common subnet.

Let's take the previous example of an increment of 64 and expand on the detail, as shown in Figure 3.8.

Subnets Created with Increment of 64 – NetID and Broadcast ID shown:

.0 N	.64 N	.128 N	.192 N
.1	.65	.129	.193
.	.	.	.
.	.	.	.
.	.	.	.
.62	.126	.190	.254
.63 B	.127 B	.191 B	.255 B

FIGURE 3.8 Subnets from increment of 64 with network ID and broadcast ID shown.

Note that all the multiples of the increment—the numbers that mark the start of each subnet—have been identified by an *N* for network ID, and the last IP in every subnet is marked with a *B* for broadcast ID. This leaves us with 62 IPs left over in each subnet, and any of these (but only these) can be assigned to a host.

This leaves us with a range of IP addresses within every network or subnet that can be assigned to hosts. There is an unofficial convention that the gateway or router for a subnet is assigned the first or the last IP address available, but that is entirely arbitrary.

ExamAlert

You need to know exactly what the first and last IP addresses are in any subnet; a lot of questions ask for them, and it's fundamental to understanding what is happening when you subnet.

The first valid IP address is defined as

Network ID + 1

In Figure 3.8, the first valid host IPs in each subnet are .1, .65, .129, and .193.

The last valid host is defined as

Broadcast ID - 1

In Figure 3.8, the last valid host IPs in each subnet are .62, .126, .190, and .254.

Tip

Here are some handy tips to help you keep track of the network ID, first and last hosts, and broadcast ID:

Network ID:	Always even
First host:	Always odd
Last host:	Always even
Broadcast ID:	Always odd

See how the subnetted mask in the previous example has shortened the number of 0s at the end of the mask as compared to the default of 8? We now have only six 0s in the host part, so our formula would be

$$2^6 - 2 = 62$$

Here's something interesting: It doesn't matter what IP address you use with this mask; that mask will always give you 62 hosts on each subnet. You can pick a Class A address, say 22.1.1.0, and that mask would still make 62 hosts per subnet. The number of 0s at the end of the mask always drives how many hosts are on each subnet, regardless of the address.

So, what happened to all the other host IPs we started with? Remember that subnetting takes a classful A, B, or C network and splits it into several equal-sized pieces. It's just like cutting a pie into pieces; the original amount of pie is still there, but each piece is now separate and smaller.

Remember that the number of 0s at the end at the mask always defines how many hosts are on each subnet, regardless of the address in use.

Number of Subnets

Following on with the pie analogy, we know that we slice a classful network into pieces—but how many pieces? There is a simple mathematical relationship

to this as well, but it is slightly more complex because of an old rule that we sometimes have to deal with.

The basic formula for the number of subnets is similar to the hosts formula. It is simply 2^S , where S is the number of bits in the subnet field—that means the number of 1s in the mask past the default boundary for that address. If you look at Figure 3.9, you can see how this works.

The default boundary for that Class C address should be at the 24th bit, where the third octet ends and the fourth begins. The subnetted mask extends that by 2 bits into the fourth octet. So, we have stolen 2 bits, and our formula would look like this:

$$\# \text{ of subnets} = 2^S$$

$$S = 2$$

$$2^2 = 4$$

IP Address and Mask: 192.168.21.1 255.255.255.192

Binary IP: 11000000.10101000.00010101.00000001

Binary Mask: 11111111.11111111.11111111.11000000




FIGURE 3.9 Subnetted Class C with increment of 64.

We made four subnets, as you saw earlier. To figure out how many bits we stole, we first must know where the default boundary is so that we know where to start counting. This is where knowing the address classes and the correct default masks is critical; if you can't figure this out, you will not be able to answer most subnetting questions correctly, and that would be bad.

Now here's where things get tricky. A rule that some older systems use says that the first and last subnets created are invalid and unusable. The rule is known as the Subnet Zero Rule, and obviously if it is in effect, we lose two subnets from the total we create. These two subnets will be referred to from now on as the zero subnets. Newer systems do not use the Zero Subnets Rule, including newer Cisco devices. This is confusing and makes things more difficult—but difficult is not something Cisco shies away from on its certification exams. So if you want your CCNA, pay attention to the question and don't complain about how hard it is.

ExamAlert

Cisco tests might be difficult and tricky, but they are fair; they do not withhold information you need to answer the question. The test question will always tell you whether somehow the Zero Subnets Rule is in effect; yes, both types of questions are asked.

The Cisco IOS supports the use of the zero subnets. The command **ip subnet zero** turns on the ability to use them, so that might be how the question is telling you whether they are in effect. Once you pass your CCNA, you will not likely have to worry about the Zero Subnets Rule again, unless you lose your mind and decide to become a Cisco trainer.

Tip

After you determine whether the zero subnets are available, use the following to get the calculation for the number of subnets right:

Zero subnets not available?	Subtract two subnets: The formula is $2^s - 2$.
Zero subnets available?	Keep all subnets: The formula is 2^s .

Working with Subnetting Questions

ExamAlert

The approach you need to take to any subnetting question is simple. After you become fluent in subnetting, you can take some shortcuts; but to build a solid understanding, you need to be methodical.

Every subnetting question you ever see will revolve around one of three things:

- ▶ Number of hosts
- ▶ Number of subnets
- ▶ The increment

Your task will be to simply figure out what the question is asking for and solve it without getting confused or distracted.

Determining Host Requirements

There are only two scenarios when determining the host requirements: Either you are given a mask and asked how many hosts per subnet this creates or you are given a requirement for a certain number of hosts and asked to provide the

appropriate mask. Either way, the number of 0s at the end of the mask drives how many hosts per subnet there will be; the address to which that mask is applied is irrelevant. Your task is to put the correct number of 0s at the end of the mask such that $2^H - 2$ is greater than or equal to the desired number of hosts, or to determine what the value of $2^H - 2$ actually is. From there, you must choose the correct expression of the mask, either in dotted-decimal or CIDR notation.

Determining Subnet Requirements

The scenarios for determining subnet requirements are quite similar to the host questions. Either you are told how many subnets you need and asked to provide the appropriate mask or you are given a mask and asked how many subnets it creates. Note that in both cases (unlike hosts questions), you must know the IP address or at least the class of address you are working with.

Creating subnets happens by extending the default mask, so you must know where the mask should end by default—and for that you need to know the class of address. Once you know where to start, simply extend the mask by the correct number of subnet bits such that $2^S - 2$ (or possibly just 2^S) gives you the correct number of subnets.

ExamAlert

Remember that the Zero Subnets Rule might come into play here. Although the majority of questions say that the zero subnets are valid and therefore the formula should be 2^S , it's possible that a few questions may clearly state that zero subnets are not available. Read the question!

Determining Increment-Based Requirements

Increment questions are the most challenging and complex subnetting questions, often requiring you to do a lot of legwork before you can get to the answer.

Increment questions often give you two or more IP addresses and masks, and ask you things such as, “Why can't Host A ping Host B?” The answer could be that A and B are on different subnets. To determine this, you need to understand where those subnets begin and end, and that depends on the increment. Another popular question gives you several IP addresses and masks that are applied to PCs, servers, and routers. The system, as it is described, is not working, and you need to determine what device has been incorrectly configured—perhaps two IPs in different subnets, perhaps a host that is using a network ID or broadcast ID as its address.

The key is to first determine what the increment is or should be; then, carefully plot out the multiples of the increment—the network IDs of all the subnets. Then you can add the broadcast IDs, which are all one less than the next network ID. Now you have a framework into which you can literally draw the host IP ranges, without risk of “losing the picture” if you do this all in your head.

All of these skills take practice. Everyone goes through the same process in learning subnetting: For quite a while, you will have no idea what is going on—then suddenly, the light goes on and you “get it.” Rest assured that you will get it. It takes longer for some than others, and you do need practice or you will lose the skill.

The Subnetting Chart

You should now understand concepts and mechanics of subnetting. You can do it and get the right answer almost all of the time, but it takes you a while. This is good—congratulations! If you are not at that point yet, you should practice more before you look at this next section.

What follows is one of many variations of a subnetting chart. This is a good one because it is easy to use under pressure when your brain will behave unpredictably.

ExamAlert

You must be able to re-create this chart exactly and correctly before you start your exam. If you make a simple mistake in creating your chart, you could easily get all of your subnetting questions wrong, and that would probably cause you to fail.

The chart represents the last two octets of a subnet mask, and what effect a 1 or a 0 in the different bit positions will have. It lists the increment, CIDR notation, the mask in decimal, the number of hosts created, and the number of subnets formed from a Class B and C address. Use an acronym to help get the rows correct: “Internet Class May Have Been Canceled.” (I = increment, C = CIDR, M = mask, H = host, B = B hosts, C = C hosts). Figure 3.10 shows a completed version.

Subnetting Chart

	32768	16384	8192	4096	2048	1024	512	256											
Increment	128	64	32	16	8	4	2	1	-	128	64	32	16	8	4	2	1	-	-
CIDR	/17	/18	/19	/20	/21	/22	/23	/24	-	/25	/26	/27	/28	/29	/30	/31	/32	-	-
Mask	128	192	224	240	248	252	254	255	-	128	192	224	240	248	252	254	255	-	-
Hosts	32766	16382	8190	4094	2046	1022	510	254	-	126	62	30	14	6	2	0	0	-	-
B-Subnet	2	4	8	16	32	64	128	256	-	512	1024	2048	4096	8192	16384	~	~	-	-
C-Subnet	1	1	1	1	1	1	1	1	-	2	4	8	16	32	64	~	~	-	-

FIGURE 3.10 The subnetting chart.

The following are steps to re-create the chart:

1. The first row is simply the binary bit position values—the powers of 2. Start at the right with 1 and keep doubling the value as you go left: 1, 2, 4, 8, 16, 32, 64, 128. Repeat for the third octet.
2. The second row is the CIDR notation—the number of 1s in a row in the mask. Our chart starts at the 17th bit, so number the second row starting at 17, through 32.
3. The third row is the mask in binary. Add consecutive bit values together from left to right to get the valid mask values of 128, 192, 224, 240, 248, 252, 254, and 255. Or you can just memorize them.
4. The fourth row is the number of hosts created. Starting at the right side of the fourth octet, subtract 2 from the increment line (the first line) and enter that value. Do this for the whole fourth octet. When you get to the third octet (the left half of the chart), you will have to change your approach: The value will keep increasing in the same pattern, but subtracting 2 from the top row won't work anymore because the top row resets for the third octet. The simplest approach is to double the last value and add 2. For example, $(126 \times 2) + 2 = 254$, $(254 \times 2) + 2 = 510$, and so on.
5. The fifth row is the number of subnets created from a Class B address. Starting at the left side of the chart (the third octet), repeat the values from the first line, but in reverse order. Remember to start at 2!

Caution

Remember that the Zero Subnets Rule will change your answers and how you use your chart. If the zero subnets are not allowed, simply deduct 2 from the values in lines 5 and 6 of your chart in the appropriate octet.

6. The sixth row of the chart is the number of subnets created from a Class C address. Remember, with a Class C, we do not make any subnets (that is, we have only one network) in the third octet, so we have all 1s there. For the fourth octet, the numbers are the same as in row 5; just start them in the fourth octet instead. The same caution and tactic about the zero subnets applies.

Provided you have built it correctly, your chart is a huge help in answering subnetting questions quickly and accurately. All you need to do is determine what the question is asking for, and then look up that value on your chart. All of the

answers you need will be in the same column. Practice building and using the chart until it becomes something you can do without thinking. You will need your brain for other more complicated problems.

Cram Quiz

1. What is the ideal mask to use on point-to-point serial links?
2. Is 172.16.255.0/18 a valid host IP?
3. Julie's IP address is 192.168.1.21 255.255.255.240. Joost's IP is 192.168.1.14/28. Their computers are connected together using a crossover Ethernet cable. Why can't they ping each other?
 - A. The subnet masks are different
 - B. They can. This is another trick question.
 - C. Because they are in different subnets.
 - D. Because the router does not support subnetting.
 - E. Because it should be a straight-through cable.

Cram Quiz Answers

1. /30, or 255.255.255.252. That provides two valid host IPs, one for each end of the link.
 2. Yes. The subnet's valid host IPs range from 172.16.192.1 to 172.16.255.254 .
 3. Answer C is correct. Answer A is incorrect; the masks are the same, just written differently. Answer B is incorrect, sorry. Answers D and E are incorrect. First of all, the question specifically mentions that they are cabled back to back with a crossover; Second, a router that doesn't support subnetting probably came as a prize in a box of cereal. Real routers support subnetting.
-

VLSM

CramSaver

1. Why is VLSM important to modern IP networks?
 - A. Because networks that use the same mask cannot route.
 - B. Because every subnet must use a different mask to avoid conflicts.
 - C. Because it allows us to uniquely identify each subnet by its mask number.
 - D. Because it allows each subnet in a routed system to be correctly sized for the requirement.
2. What must a routing protocol be able to do to support VLSM?
 - A. Multicast
 - B. Automatically summarize networks to a common mask
 - C. Advertise the mask for each subnet in the routing update

Answers

1. Answer D is correct. Answers A and B are simply untrue. Answer C sounds good, but networks are uniquely identified by the network ID, not the mask.
2. Answer C is correct. Answer A is not a requirement. Answer B is effectively the opposite of VLSM

Variable-length subnet masking or VLSM can be defined as the capability to apply more than one subnet mask to a given class of addresses throughout a routed system. Although this is common practice in modern networks, there was a time when this was impossible because the routing protocols in use could not support it. Classful protocols such as RIPv1 do not include the subnet mask of advertised networks in their routing updates; therefore, they cannot possibly learn the existence of more than one mask length. Only classless routing protocols—EIGRP, OSPF, RIPv2, IS-IS, and BGP—include the subnet mask for the networks they advertise in their routing updates and thus publish a level of detail that makes VLSM possible.

The main push for VLSM came from the need to make networks the right size.

Subnetting logically creates the appropriately sized networks, but without the capability for routing protocols to advertise the existence (for example) of both a /26 and a /30 network within the same system. Prior to VLSM-capable routing protocols, the network in our example would have been confined to using

only /26 masks throughout the system. The use of VLSM has two main advantages that are closely linked:

- ▶ It makes network addressing more efficient.
- ▶ It provides the capability to perform route summarization (discussed in the next section).

ExamAlert

Know the definition of VLSM and its two main advantages.

An illustration of the need for VLSM is shown in Figure 3.11.




FIGURE 3.11 Inefficient addressing without VLSM.

The diagram shows several branch offices using subnetted Class C (/26) addresses that provide each branch with 62 possible host IPs. The branches are connected to the central office via point-to-point WAN links. The ideal mask to use for such a link is /30 because it provides only two hosts, one for each end of the link. The problem arises when the routing protocols are configured: Prior to VLSM, the /30 networks could not be used because the /26 networks existed in the same system and the classful routing protocols could only advertise one mask per class of address. All networks, including the little /30 links, had to use the same mask of /26. This wastes 60 IP addresses on each WAN link.

With the implementation of VLSM-capable routing protocols, we can deploy a /30 mask on the point-to-point links, and the routing protocols can advertise them as /30s along with the /26s in the branches because the subnet mask for each network is included in the routing updates. Figure 3.12 illustrates the preferred, optimized addressing scheme that takes advantage of VLSM.




FIGURE 3.12 Optimized addressing using VLSM.

Note that using VLSM has allowed us to make the point-to-point link networks the ideal size (two hosts on each) using /30 masks. This has allowed us to use a single subnetted Class C network for all the addressing requirements in this scenario—and as you’ll see, it makes a perfect opportunity to summarize these routes. This is what is meant by “more efficient addressing”—in other words, making networks the right size without depleting the limited address space or limiting future growth.

Cram Quiz

1. True or false: It is impossible to subnet a subnet.
2. How does VLSM make IP addressing more efficient?
 - A. By increasing the total number of IP addresses.
 - B. By decreasing the total number of IP addresses
 - C. By creating subnets
 - D. By allowing a routed system to include subnets of different mask lengths to suit requirements

Cram Quiz Answers

1. False. It is not only possible, you can in fact subnet a subnet of a subnet—and keep going as long as there are sufficient bits left in the mask.
 2. Answer D is correct. Answers A and B are incorrect: VLSM neither creates nor deletes IP addresses. (You might argue that subnetting reduces the available IPs, but VLSM itself is not necessarily subnetting.) Likewise, Answer C is not correct; VLSM does not create subnets. It is simply the application of different subnet masks within a routed system.
-

Route Summarization

CramSaver

1. What is the best summary for the following range of subnets?
172.20.32.0/24 to 172.20.47.0/24
2. Why is summarization so important to an efficient routed system?
 - A. It adds detail to the route tables of routers.
 - B. Summarization sends all subnets as classful networks, eliminating the overhead of transmitting the mask in routing updates.
 - C. Summarization reduces the size of route tables, prevents route table instability due to flapping routes, and reduces the size of routing updates.
 - D. Summarization enforces router authentication, preventing spurious updates from excessively loading the router.

Answers

1. 172.20.32.0/20 or 172.20.32.0 255.255.240.0
2. Answer C is correct. A is incorrect; that is the opposite effect of summarization. Answer B is incorrect; this describes classful routing protocols (in part). Answer D is completely untrue.

If subnetting is the process of lengthening the mask to create multiple smaller subnets from a single larger network, route summarization can be described as shortening the mask to include several smaller networks into one larger network address. As the network grows large, the number of individual networks listed in the IP route table becomes too big for routers to handle effectively. They get slower, drop packets, and even crash. This, of course, is an undesirable state of affairs. With more than 450,000 routes (at the time of this writing, anyway) known to major Internet routers, some way to reduce the number of entries is not only desirable, but also critical.

In the previous VLSM example, all the subnets for the branches and the WAN links were created from the 192.168.0.0 /24 Class C network. If we take that diagram and put it into context, we can see how route summarization can reduce the number of entries in the route table, as shown in Figure 3.13.

The Central Office router can either send a routing update with all the subnets it knows about listed individually, or it can send a single line in the update that

essentially says, “Send anything that starts with 192.168.0 to me.” Both methods work; the issue is one of scalability. No router will ever collapse under the load of advertising six subnets, but make it 6,000 subnets and it makes a huge difference in performance if you summarize as much as possible.

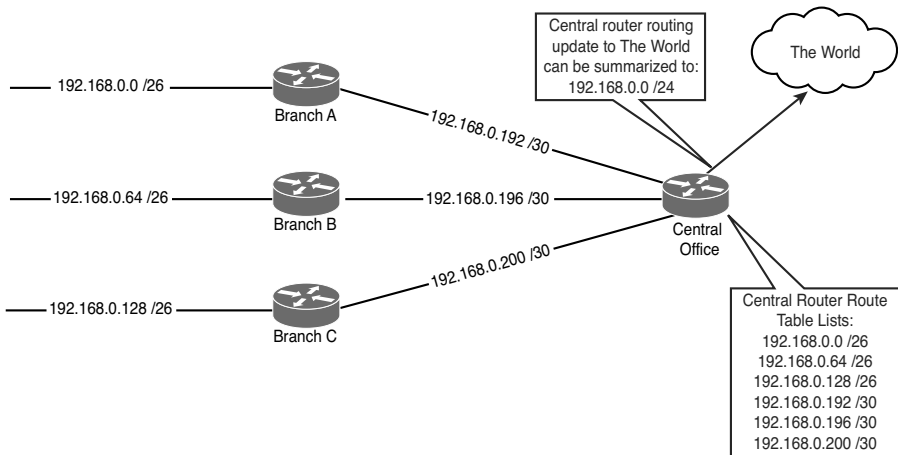


FIGURE 3.13 Simple route summarization example.

Route summarization takes a set of contiguous networks or subnets and groups them together using a shorter subnet mask. The advantages of summarization are that it reduces the number of entries in the route table, which reduces load on the router and network overhead, and hides instability in the system behind the summary, which remains valid even if summarized networks are unavailable.

Note

The word *contiguous* sometimes confuses people. It is not a typo of *continuous*; the word means “adjacent or adjoining.” For example, when we make subnets using a 16 increment, the first four network IDs are .0, .16, .32, and .48. Those four subnets are contiguous because they are adjacent to each other. If we take the last four subnets from that same increment (.192, .208, .224, and .240), they are contiguous with each other, but not with the first four—there are a bunch of subnets between the two sets.

ExamAlert

Know the definition and advantages of route summarization.

Summarization Guidelines

It is important to follow a few rules and guidelines when summarizing. Serious routing problems will happen otherwise—such as routers advertising networks inaccurately and possibly duplicating other routers' advertisements, suboptimal or even totally incorrect routing, and severe data loss.

The first rule is to design your networks with summarization in mind, even if you don't need it yet. This means that you will group contiguous subnets together behind the router that will summarize them—you do not want to have some subnets from a summarized group behind some other router. The summary is essentially saying, "I can reach the networks represented by this summary; send any traffic for them through me." If one (or more) of the networks behind the summarizing router is unavailable, traffic will be dropped—but not by the summarizing router, because the individual routes to the networks that were summarized are still valid, and have a longer match entry than the summary. The packet will get routed to the router that connects to the dead network, and dropped there. Advance planning, including making plenty of room for future growth, will give you a solid, scalable network design that readily lends itself to summarizing. Figure 3.14 shows a badly designed network that will be almost impossible to summarize because the subnets are discontinuous, with individual subnets scattered all over the system.

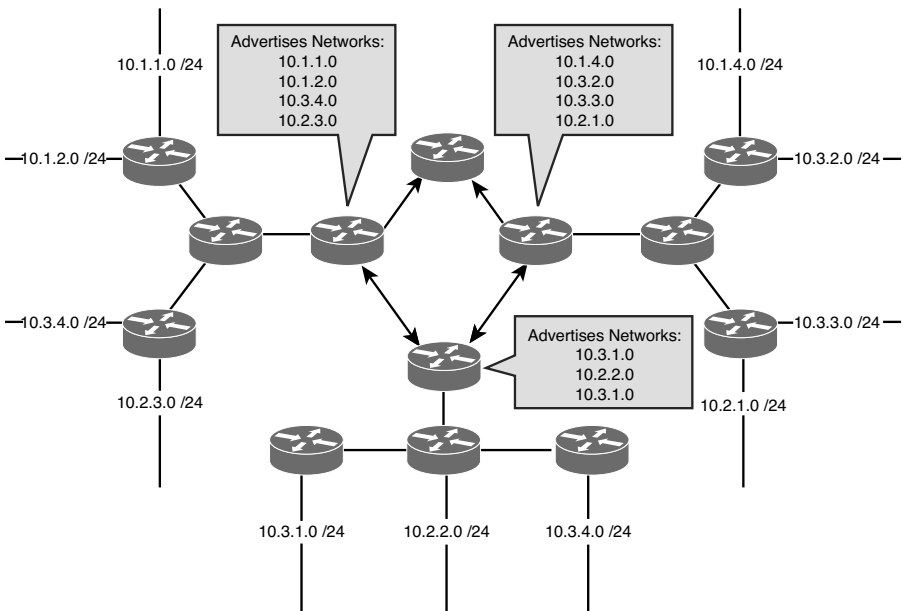


FIGURE 3.14 Poor planning prevents proper performance.

The second rule is to summarize into the core of your network. The core is where the bigger, faster, busier routers are—like the Central Office router in the previous example. These routers have the job of dealing with high volumes of traffic headed for all different areas of the network, so we do not want to burden them with big, highly detailed route tables. The further you get from the core, the more detail the routers need to get traffic to the correct destination network. It's much like using a map to drive to a friend's house; you don't need a great deal of detail when you are on the highway, but when you get into the residential areas, you need to know very precise information if you have a hope of finding the place.

Figure 3.15 illustrates the same network after your friendly neighborhood Cisco Certified Internetwork Expert has spent the afternoon re-addressing the network and configuring summarization. This network will scale beautifully and have minimal performance issues (at least because of route table and routing update overhead).

Following these rules will give you one of the additional benefits of summarization as well: hiding instability in the summarized networks. Let's say that one of the branches is having serious spanning-tree problems because Sparky the Junior Woodchuck was allowed to configure a Cisco switch. (This is actually a felony in some states.) That route could be "flapping"—up, down, up, down—as spanning-tree wreaks havoc with your network. The router will be doing its job, sending out updates every time the route flaps. If we were not summarizing, those flapping messages would propagate through the entire corporate system, putting a totally unnecessary and performance-robbing load on the routers. Once you summarize, the summary is stable: It can't flap because it is not a real network. It's just like a spokesperson at a press conference: "The rumors of a fire at the Springfield plant have had no impact on production whatsoever." Meanwhile, the Springfield plant could be a charred hulk. The summary is still valid, and traffic will still be sent to the router connected to the flapping network. This keeps people from asking any more questions about the Springfield fire. However, if someone were to send a shipment to Springfield, it would be hastily redirected to another site (or dropped). All we have done is hide the problem from the rest of the world so that we don't flood the Internet with rapid-fire routing updates.

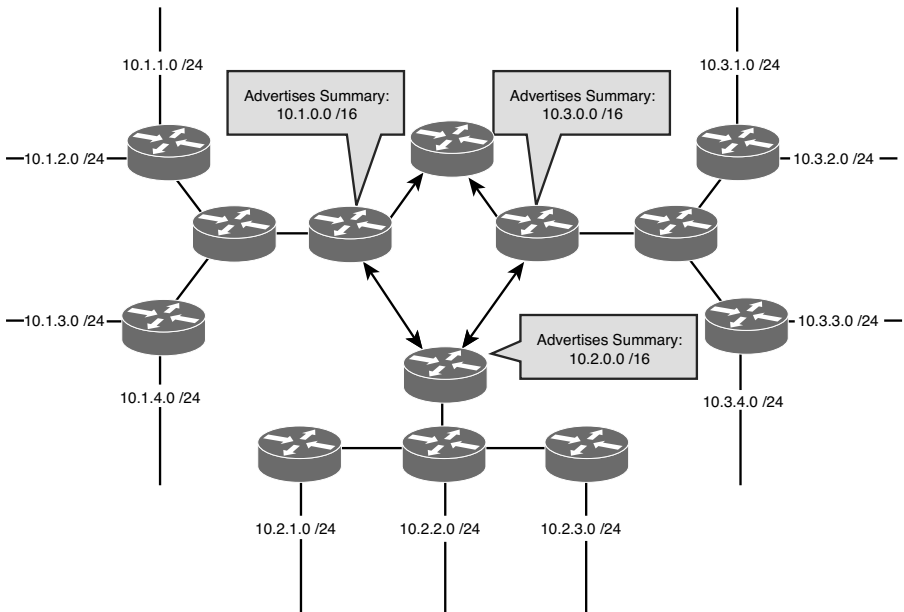


FIGURE 3.15 Proper planning prevents poor performance.

Determining Summary Addresses

When using routing protocols in classless configurations, creating summary addresses is a totally manual process. Classful routing protocols perform automatic summarization, but that is not as fancy as it sounds. They simply treat any subnet as the classful address from which it was created, which works if your networks are built with this in mind; however, in reality that is too simplistic and real networks need more customized summarization. The upshot of all this is that you need to understand how to determine the summary address given a set of networks to be summarized, and you also need to be able to figure out if a particular network is included in a given summary.

Remember that summarization is exactly the opposite of subnetting; in fact, another term for summarization is supernetting. (You might also see it called aggregation.) When we subnet, we lengthen the mask, doubling the number of networks each time we add an extra bit to the mask. Supernetting does the opposite: For each bit we retract or shorten the mask, we combine networks into groups that follow the binary increment numbers.

To illustrate this, let's look at the private Class B address space. These networks are listed as follows:

```

172.16.0.0 /16
172.17.0.0 /16
172.18.0.0 /16
172.19.0.0 /16
172.20.0.0 /16
172.21.0.0 /16
172.22.0.0 /16
172.23.0.0 /16
172.24.0.0 /16
172.25.0.0 /16
172.26.0.0 /16
172.27.0.0 /16
172.28.0.0 /16
172.29.0.0 /16
172.30.0.0 /16
172.31.0.0 /16

```

If you look carefully, you will notice that the range of networks is identified in the second octet. The octet where the range is happening is referred to as the *interesting octet*. This is your first clue where to begin your summarization.

The next step is to figure out what the binary values of the network's range are. The binary values for the interesting octet are shown in Figure 3.16.

```

16 = 0 0 0 1 0 0 0 0
17 = 0 0 0 1 0 0 0 1
18 = 0 0 0 1 0 0 1 0
19 = 0 0 0 1 0 0 1 1
20 = 0 0 0 1 0 1 0 0
21 = 0 0 0 1 0 1 0 1
22 = 0 0 0 1 0 1 1 0
23 = 0 0 0 1 0 1 1 1
24 = 0 0 0 1 1 0 0 0
25 = 0 0 0 1 1 0 0 1
26 = 0 0 0 1 1 0 1 0
27 = 0 0 0 1 1 0 1 1
28 = 0 0 0 1 1 1 0 0
29 = 0 0 0 1 1 1 0 1
30 = 0 0 0 1 1 1 1 0
31 = 0 0 0 1 1 1 1 1

```

FIGURE 3.16 Binary values for Class B private range second octet.

You should see a pattern in the binary values: The first 4 bits are all the same. The range is actually happening in the last 4 bits in the second octet; those 4 bits range from 0000 through 1111; the first 4 bits are common for all 16 networks in the range.

The next step is to identify those common bits. While you are learning how to do this, it's a good idea to write out the binary for the range and draw a line that represents the boundary between the common bits and the variable bits in the range. Remember, be absolutely sure that your boundary line is in the right place: For all the networks in the range, everything to the left of the line must be identical, and everything to the right will be the ranging values.

The next step is easy. We are about to summarize: All we need to do is to build a subnet mask that puts a 1 under all of the common bits in the range, and a 0 under everything else—1s to the left of the boundary, and 0s to the right, as shown in Figure 3.17.

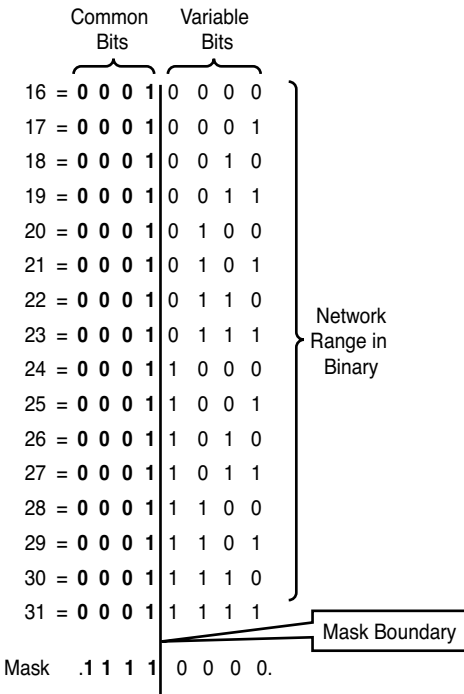


FIGURE 3.17 Identifying and masking the common bits in a summary.

The last step is to actually create the summary statement. A summary is always an IP address plus a mask; the IP is usually a network ID, and it should be the first network in the range. In our example, the first network ID is 172.16.0.0 so that is the IP we will use. For the mask, the first octet is the same in the whole range, and we have figured out that the first four bits in the second octet are always the same. Remembering that a mask is always a string of 1s followed by a string of 0s, this means that we should mask all 8 bits in the first octet and the first 4 in the second octet, so our mask looks like this:

```
11111111.11110000.00000000.00000000
```

That can also be expressed as

```
255.240.0.0 or /12
```

So, our summary statement becomes

```
172.16.0.0 255.240.0.0
```

or

```
172.16.0.0 /12
```

Reverse engineering this is the same process. You are given a summary statement and asked what networks it includes. The octet in which the mask changes from 1s to 0s is the interesting one, where the range will be defined. Jot down the address and mask in that octet in binary and see what possible values are in the range. Then check the networks to see if those are in the range. Figure 3.18 gives an example.

Given the Summary: 192.168.8.0 /21:

- 3rd octet is interesting

- 3rd octet of Mask in Binary:

.1 1 1 1 1 0 0 0.

- 3rd octet of IP in Binary:

.0 0 0 0 1 0 0 0.

Last 3 bits are the IP range:

.0 0 0 0 1 **0 0 0.** = .8

... through

.0 0 0 0 1 **1 1 1.** = .15

Mask Boundary

Therefore, the range of networks is 192.168.8.0 through 192.168.15.0

- Network 192.168.12.0 /24 would be in this range.

- Network 192.168.16.0 /24 would not be in this range.

- Network 192.168.0.0 /24 would not be in this range.

FIGURE 3.18 Summary address analysis.

Cram Quiz

1. You are given the following ranges of subnets:

- | | |
|-----------------|------------------|
| 192.168.1.0/29 | 192.168.1.128/29 |
| 192.168.1.8/29 | 192.168.1.136/29 |
| 192.168.1.16/29 | 192.168.1.144/29 |
| 192.168.1.24/29 | 192.168.1.152/29 |
| 192.168.1.32/29 | 192.168.1.160/29 |
| 192.168.1.40/29 | 192.168.1.168/29 |
| 192.168.1.48/29 | 192.168.1.176/29 |
| 192.168.1.56/29 | 192.168.1.184/29 |

Your task is to summarize these two ranges of subnets. Do *not* include any subnets not named in the ranges in your summary. (Hint: You may use more than one summary address).

2. Your boss complains that manual route summarization is difficult and complex, and wonders if maybe you should not bother with it. What are the most compelling arguments in favor of route summarization? Choose all that apply.
- A. Utilizes the full RAM and CPU performance capacity of the routers
 - B. Can suppress the effects of an unstable or “flapping” interface
 - C. Advertises complete and detailed route tables
 - D. Increases security by advertising “fake” networks
 - E. Reduces the size of the route tables
 - F. Reduces the load on RAM, CPU, and bandwidth of routers

Cram Quiz Answers

1. Two summary statements are required. Because these two network ranges are discontinuous, we cannot use a single statement without including the ranges between and after, which is both not allowed in the question and not generally a good idea in practice. The two summary statements are 192.168.1.0/26 and 192.168.1.128/26. The /26 in each statement can also be expressed as 255.255.255.192.
 2. Answers B, E, and F are correct. Answer A is incorrect because summarization actually reduces the load on routers, and maxing out your router is not a good idea to begin with. Answer C is incorrect; summarization sends out summary routes that represent the detailed routes. Answer D is incorrect, but tricky: Summarization does in fact send out fake routes, but this does nothing to increase security.
-

IPv6

CramSaver

1. Which of the following are valid types of IPv6 address? Choose all that apply.
 - A. Global unicast
 - B. Unique local
 - C. Link local
 - D. Multicast
 - E. Anycast
 - F. Broadcast
 - G. Directed broadcast
2. Which of the following are valid IPv6 addresses? Choose all that apply.
 - A. 2001:0db8:0000:0000:0000:ff00:0042:8329
 - B. 2001:db8:0:0:0:ff00:42:8329
 - C. 2001:db8::ff00:42:8329
 - D. 0000:0000:0000:0000:0000:0000:0000:0001
 - E. ::1
 - F. ::192:168:1:1
3. Which of the following is a valid command to apply an IPv6 address to a router interface?
 - A. **interface fastethernet 1/0 ip address 2001:AB00:00FF:1::/64 eui-64**
 - B. **interface fastethernet 1/0 ipv6 address 2001:AB00:00FF:1::/64 eui-64**
 - C. **line con 0 ipv6 address 2001:AB00:00FF:1::/64 eui-64**
 - D. **interface fastethernet 1/0 ipv6 address 2001:AB00:000FF:1/64 eui-64**

Answers

1. Answers A, B, C, D, and E are correct. IPv6 does not broadcast, so Answers F and G are wrong.
2. All answers are correct! Even F, which looks bogus, is sometimes used (probably to mke the admin feel a little better).
3. Answer B is correct. Answer A is wrong because it uses **ip** instead of **ipv6**. Answer C is wrong because the Console port cannot be given an IPv6 (or IP for that matter) address. Answer D is wrong because the IPv6 address is missing the :: to make it the correct 64-bit length.

Up to this point, when we talked about IP or an IP address, we were referring to IP Version 4. IPv4 was created to build a Defense Department network in the early 1970s. At the time, no one foresaw that the Internet as we know it today was going to happen. The designers of the TCP/IP suite of protocols did not plan for their little project to balloon into the largest network in the world and revolutionize the commercial, cultural, and communications behavior of the whole planet.

But it did, and a couple problems came to light rather quickly when the Internet started to really catch on. One really tricky one was that the address “space” was originally handed out without quite enough thought and planning as to who got what size chunks, and what routers would be responsible for those chunks. At the time it didn’t matter; there were plenty of addresses to go around. But as the routers started to get really large route tables, with all these networks being added, they had trouble dealing with it. Routers at the time were relatively small and slow, and when the route tables became so large, they were overloaded, slow to do their jobs, and generally poor performers. Solutions were urgently needed because the Internet was growing very fast and the problem was only getting worse.

The solutions came in things like VLSM-capable protocols, route summarization, a reassignment and redistribution of addresses, and the NAT service. These solutions have allowed the IPv4 address space to continue to function and serve as the address system for the Internet, but the second problem is one we can’t get around: The mathematical reality is that there are not enough IP addresses available to meet the demand (especially in Europe and Asia). More people want Internet addresses than there are addresses to hand out.

This is where IPv6 comes in. Whereas an IPv4 address is a 32-bit string, theoretically providing more than 4 billion IP addresses (for the sake of clarity I’ll ignore the fact that a large number of these addresses are not really usable). An IPv6 address is 128 bits long, providing about 3.4×10^{38} possible addresses, or as the story goes, 500,000,000,000,000,000,000,000,000,000 addresses for each of the 6.5 billion people on the planet. Running out of IPv6 addresses is not expected to be a problem.

Along with the sheer number of addresses available, IPv6 also cleans up a few of the issues with IPv4, making the operation and management of large inter-networks easier and more efficient, and adds some useful new functionality as well. So now we can easily envision a world where anything we want can have an Internet IP address (including silly things such as the fridge), where an Internet-enabled mobile phone can keep its IP address as it moves across the globe, and all the difficulty and headache caused by using VPNs through NAT disappears.

IPv6 Address Allocation

An organization called the Internet Corporation for Assigned Network Numbers (ICANN) has the overall responsibility for dividing up the IPv6 address space. They do so with the benefit of a better understanding of the global demand for Internet IP addresses and the luxury of a huge number of addresses to hand out.

The system works like this: First, remember that for the Internet to work well, we need to use route summarization so that the route tables don't get huge and slow the routers down. Route summarization works best if every router is responsible only for its "branch of the tree," with smaller branches feeding into larger and larger ones as we get closer to the core or trunk of the tree. This allows the possibility for a single router to advertise a summary that in effect says, "I can reach all North American routes." That big router connects to other routers that summarize routes for four major Internet service providers (ISPs). Each ISP router connects to smaller ISPs or large enterprise customers, who advertise the summaries that represent the addresses assigned to them. Figure 3.19 gives some idea of how this system works.

The beauty of the system is that it is organized, planned, and executed in advance, with efficient routing in mind. The large number of addresses available also means that changes at or below the ISP level, for example, because of mergers or large customers changing Internet providers, do not affect the global routing information at the core.




FIGURE 3.19 Global IPv6 address design.

IPv6 Address Notation

IPv6 addresses are different in appearance from IPv4. Of course, they are 128 bits long, so even in binary they would be four times longer than a 32-bit IPv4 address, but in notation that humans read and write the format is still different. Instead of using dotted-decimal in four octets, we use hexadecimal in eight sets of four characters separated by colons, like this:

```
2201:0FA0:080B:2112:0000:0000:0000:0001
```

The use of hex makes it a little easier to represent all those 128 bits in a shorter format because each character represents 4 bits. But it's still a long thing to type out, and remember that network people are generally lazy—so we have a couple of truncation methods to make the long addresses even shorter. The first method is that we are allowed to drop leading 0s (0s that appear at the beginning of each set), like so:

```
2201:FA0:80B:2112:0:0:0:1
```

That makes for a little less typing and a little more clarity. Pay attention to the fact that dropping 0s at the end of each set is not allowed! Dropping leading 0s does not change the value of the set; dropping 0s at the end does (like removing a 0 from the end of your paycheck amount—not good!)

The second truncation method we can use is to condense contiguous groups of all-0 sets. In our example, there are three sets that are all 0s. We can represent these by a double colon, like this:

```
2201:FA0:80B:2112::1
```

This is as short as it gets. We are only allowed to do the double-colon trick once in any address, so if you see an address with two double-colons in it, it is not valid. Here's an example:

```
2201::BCBF::1
```

One last piece of the addressing notation: the mask. We do not represent the mask as another set of hex characters. Instead, we identify the prefix length with slash notation. This is not as confusing as it seems: The slash notation simply identifies how many bits identify the network part, with the remainder being the host part.

For example, the North American registry ARIN (American Registry for Internet Numbers) was given the block of 2620:0000::/23 in September 2006. This indicates that the first 23 bits of 2620:0000:: identify the block of addresses that the North American routers will advertise to the rest of the

world. From this point, ARIN will assign chunks of that space to the Big ISPs; Big ISP1 might get 2620:0100::/24, and Big ISP2 might get 2620:0200::/24. Those ISPs then hand out pieces of their chunk to smaller ISP or big customers, and the prefix length will get bigger as the chunks gets smaller—this should feel familiar because what we are doing here is subnetting. Don't worry, you won't be expected to subnet in IPv6. Not yet at least.

Types of IPv6 Addresses

An IPv6 address will be one of the following types. Some will be familiar, but there is one brand-new one, too.

- ▶ **Unicast:** An IPv6 unicast address is the same as an IPv4 unicast address; it is an IP that is assigned to an interface on a host. It can be the source of an IP packet or the destination for one. A packet sent to a unicast address goes to the one host with that address.
- ▶ **Global unicast:** A global unicast IPv6 address is the equivalent of a public, registered IP address. They are Internet routable, globally registered IPs that must be leased from an ISP.
- ▶ **Unique local:** Equivalent to a private IPv4 address; not registered with an ISP and not Internet routable.
- ▶ **Link local:** Every IPv6 interface gives itself a link-local address. The address range is FE80::/10, and usually combines this prefix with the last 64 bits in EUI-64 format. It is roughly equivalent to the Automatic Private IP Address (APIPA) address range of 169.254.0.0/16.
- ▶ **Multicast:** Just like in IPv4, a single IPv6 multicast address is assigned to multiple hosts so that a packet sent to the address may be delivered to multiple hosts more or less at the same time. IPv6 multicast addresses always start with the prefix FF00::/8.
- ▶ **Anycast:** An anycast address is a single address that is assigned to multiple hosts. This is similar to a multicast, except that a packet for the anycast address will be delivered to the one host that is nearest according to the routing protocol's idea of distance. There is no special prefix for anycast addresses.

There is no such thing as a broadcast in IPv6. Ever. Any requirement for broadcasting is performed by a multicast instead.

ExamAlert

Know the IPv6 address types. Remember that IPv6 cannot broadcast, ever! Any answers with the word broadcast in them are invalid.

IPv6 Address Configuration

For hosts to use IPv6 addresses, an IPv6 protocol stack must be installed. This likely means that you will need to upgrade your router IOS to provide IPv6 support. Then you can choose one of four options for address assignment.

To understand the address assignment choices better, we need to examine the concepts of stateful versus stateless configuration and the EUI-64 address format.

In IPv6, we can use DHCP to assign IP addresses just like in IPv4. The admin must set up the server with a scope of IPv6 addresses to hand out. The mechanisms used to discover and assign addresses are a little different, but the net result is the same. This is called *stateful addressing*, where the DHCP server keeps track of what hosts have been assigned what IPv6 address—in other words, the state of the host DHCP-wise.

There is another option for dynamic addressing in IPv6 called *stateless auto-configuration*. This feature allows a host to choose and configure an address for itself. The host that wants an address learns what the /64 network prefix is on the local link, then appends its MAC address (in a special 64-bit format called EUI-64), thus generating a 128-bit IPv6 address that is unique to that host because it incorporates the unique MAC of the host.

The EUI-64 format is not so difficult to understand. We simply take the 48-bit MAC address and put a special pattern, FFFE, after the first 24 bits (the six OUI characters), followed by the rest of the six hex characters in the host MAC. The only trick is that according to IPv6 rules, the seventh bit in an EUI-64 address must be 1, which identifies that the burned-in MAC address has been modified. This is a little confusing, to be sure, but you can relax because the host determines and configures its EUI-64 address all by itself, if you tell it to. Here's what an EUI-64 address conversion looks like:

Original MAC:

00-15-C5-CB-42-2B

Original MAC in binary:

00000000-00010101-11000101-11001011-01000010-00101011

7th bit = 0

Change 7th bit to 1:

00000010-00010101-11000101-11001011-01000010-00101011

EUI-64 MAC now:

02-15-C5-CB-42-2B

EUI-64 Address = </64 net_prefix>:0215:C5FF:FECB:422B

So, back to the four choices. The following really simplifies the options:

- ▶ **Static configuration:** The administrator chooses and assigns a static IPv6 address to the host NIC. It is the admin's responsibility to choose an address that will function and be valid in the network to which the host is connected.
- ▶ **Static configuration using EUI-64:** The administrator manually configures the address with the local /64 network prefix followed by the host's MAC in EUI-64 format.
- ▶ **Dynamic configuration using DHCP to assign 128-bit address:** The host is set to obtain its address from DHCP, and the DHCP server is set up to hand out IPv6 addresses from a scope.
- ▶ **Dynamic configuration using stateless autoconfiguration with EUI-64:** The host is set to obtain its address automatically, but the DHCP server either does not exist (which works fine by the way), or if it does, it only informs the host of the /64 local network prefix.

IPv6 Router Configuration

Assuming your IOS provides IPv6 support, giving it an IPv6 address is really easy. The command is carried out at the interface configuration prompt:

```
interface fastethernet 1/0
ipv6 address 2001:AB00:00FF:1::/64 eui-64
```

Notice the **eui-64** switch; this tells the router to figure out its own EUI-64 address to follow the /64 prefix provided. Without that, you must provide a full 128-bit address in the command.

To verify your configuration, use the **show ipv6 interface** command at the interface configuration prompt. The following is a sample output (with different addresses applied). You can see multiple addresses in use by the interface for global unicast, link-local, and multiple multicast groups:

```
Router#show ipv6 interface
Serial1/0 is up, line protocol is up
```

```
IPv6 is enabled, link-local address is FE80::A8BB:C2FF:FE00:D200
Global unicast address(es):
  2001:1:33::3, subnet is 2001:1:33::/64 [TENTATIVE]
Joined group address(es):
  FE02::1
  FF02::1:FF00:3
  FF02::1:FF00:D200
MTU is 1500 bytes
ICMP error messages limited to one every 100 milliseconds
ICMP redirects are enabled
ND DAD is enabled, number of DAD attempts: 1
ND reachable time is 30000 milliseconds
Router#
```

IPv6 Features

IPv6 has a couple features that you should keep in mind:

- ▶ **IPsec:** Support for IPsec is built in for IPv6; this means that every packet can be protected by IPsec transport on every IPv6 host if so configured.
- ▶ **Mobility:** IP mobility is built in, but obviously not mandatory because some hosts are not mobile.
- ▶ **Fixed header size:** The IPv6 header is fixed at 40 bytes or 320 bits. Figure 3.20 (in the next section) shows the IPv6 header.
- ▶ **ICMP for IPv6 has changed, adding new functionality:** One example of the new tricks it has learned is path MTU (PMTU) discovery: Before transmitting a packet, a host can send an ICMP message to learn what the smallest maximum transmission unit (MTU) on any link is between the sender and the destination. Then, the host sends packets that are no larger than that value. This clever trick relieves routers of having to fragment and reassemble packets over a small-MTU link, which can be a real performance hog. According to RFC 1981, hosts not using PMTU will transmit packets at the minimum IPv6 link MTU, which is actually quite small and likely to be inefficient.
- ▶ **IPv6 makes extensive use of Router Solicitation (RS) and Router Advertisement (RA) messages:** These are multicast messages to the addresses FF02::1 and FF00::2, respectively. The RS is sent from a host to all routers on the link as a multicast, and the RA message is sent from a router to all hosts on the link, also as a multicast. This is one way that the hosts learn whether DHCP is supported on the link, and possibly the DHCP server address.

The IPv6 Header

As mentioned in the preceding section, the IPv6 header is fixed at 40 bytes (320 bits) in length. Figure 16.11 shows the header fields and their sizes, and this section identifies what the fields are for.

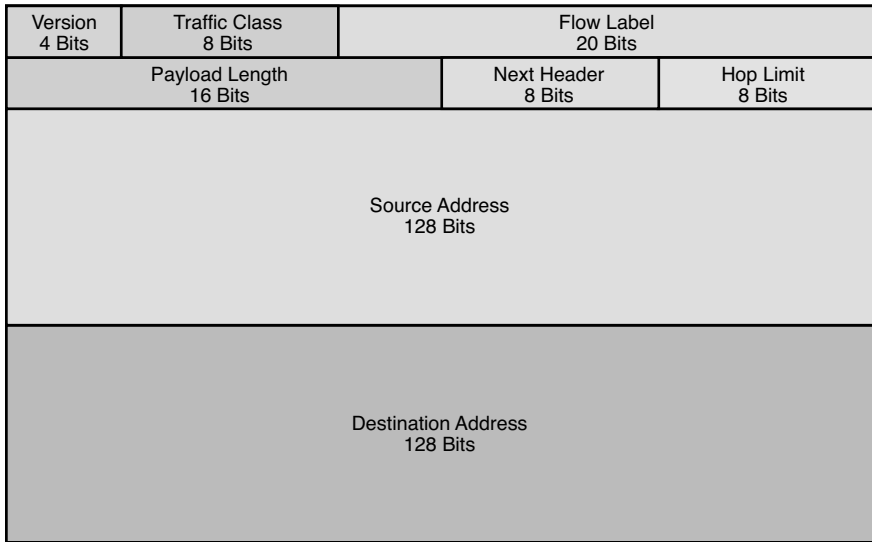


FIGURE 3.20 The IPv6 header.

The Version field identifies the IP version of this packet; for IPv6, obviously the version will be 6.

The Traffic Class field is where quality of service (QoS) marking for Layer 3 can be identified. In a nutshell, the higher the value of this field, the more important the packet. Your Cisco routers (and some switches) can be configured to read this value and send a high-priority packet sooner than other lower ones during times of congestion. This is very important for some applications, especially VoIP.

The Flow Label field is a number that identifies this packet as one of a flow of packets in a stream from sender to receiver; a good example is a VoIP call. It's best for VoIP if all the packets in a given call get sent along exactly the same path to the receiving phone, so that they arrive in the same order they were sent. The flow label is one mechanism that IPv6 routers can use to keep track of different application flows and try to make sure that all the packets within a flow get treated the same way.

The Payload Length field indicates how big the payload of this packet is; it can be variable, so the router needs to know where the packet is supposed to end. That way it knows if anything went missing. This is especially important because there is no header checksum, as there used to be in IPv4.

The Next Header field takes over the Options header functions in IPv4. Short codes for extension headers are listed in the Next Header field, and additional information is appended in additional headers after the primary IPv6 header. All of this is designed to speed up the routing of IPv6 packets by preserving the size and content of the primary header, so it can be routed in the “fast switching” path of the router.

The Hop Limit field is a cool one: Whereas in IPv4 there was a TTL field that limited the life of a packet to 255 hops. (The IPv4 TTL value typically starts at 255 and is decremented by at least 1 as a packet is processed by a router. If it reaches 0, the packet is dropped. This prevents the packet from being endlessly misrouted around the Internet, although it could be misrouted up to 255 hops.) IPv6 is smarter: The Hop Limit value is set to the actual number of hops the packet will go through to reach its destination. This hop information comes from the IPv6 routing protocols. The Hop Limit is still decremented by 1 at each router, but the more accurate value means that the packet can't be misrouted even by one hop.

The Source and Destination Address fields are self-explanatory; remember that the full 128-bit address for each is listed.

IPv6 Transition Strategies

Clearly, things are moving toward IPv6. The U.S. government once specified that all federal agencies must deploy IPv6 by 2008; then it was pushed back, ultimately to 2012. The process, though, is not going to affect every single host in these large networks overnight. Cisco wants you to be aware of their strategies for the transition to using IPv6 while still maintaining IPv4 functionality.

Tip

Putting CCENT exam studying aside for a moment, we strongly recommend that you start learning how to use IPv6 in your labs now. It is your big chance to be ready when the boss walks in and says, “We need to deploy IPv6 connectivity because of blah blah blah. Can you do it?” When you say, “Sure, no problem,” and gain massive respect, that's when you can send us an email and thank us.

The easiest IPv6 transition choice is called dual stack. Dual stacking means that the host (router, PC, printer, and so on) runs both the IPv4 and IPv6 protocol stacks and can send and receive both types of packets, probably (but not necessarily) on the same interface. The drawbacks here are the additional load on the host and whether an IPv6 stack for that device is available (your old router might not be able to run IPv6).

Tunneling mode creates a tunnel for one protocol through another. You can picture taking an IPv6 packet from the head office, encapsulating it inside an IPv4 packet to transition across the provider network, then decapsulating it on the other side and forwarding the IPv6 packet into the remote branch office. This is known as a 6-to-4 tunnel; these tunnels can be either automatic or manual. 6-to-4 tunnels have a special address range of 2002::/16. Other tunneling strategies include the following:

- ▶ **Teredo tunneling:** Named after a particularly ugly species of marine wood-boring clam that makes tunnels in wood, this technology encapsulates IPv6 packets in IPv4 UDP datagrams for routing through the IPv4 network (usually the Internet). Its chief benefit is that it can operate from behind NAT devices. It is considered a “last resort” transition strategy, meaning that you should implement IPv6 natively instead, if possible.
- ▶ **ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) tunneling:** This technology uses the IPv4 network (again, this would usually be the Internet) as a virtual NBMA data link layer. IPv6 link layer addressing is derived dynamically from IPv4 addresses, allowing dynamic neighbor discovery on top of IPv4 in addition to simple routability. ISATAP is a native capability in most Windows operating systems, Linux, and most Cisco IOS versions.

Translation means taking an IPv6 packet, removing the IP header, and replacing it with an IPv4 header that approximates the original IPv6 information as much as possible.

Translation is usually associated with a NAT router, and sometimes is known as NAT-PT (for Protocol Translation). What happens here is that the IPv6 packet header is removed and replaced with an IPv4 header (or vice versa), effectively changing from one protocol to the other. The big issues with NAT-PT are latency, performance loading, and the loss of header information in the translation process.

You need to know IPv6 configuration for your test; they may or may not ask you to actually do it (in a sim), but you definitely need to be able to recognize if the configuration they show you is valid.

Cram Quiz

1. Which is a valid alternate expression of FE80:0000:0000:0202:B3FF:0E1E:8329? Choose all that apply.
 - A. FE80::0202:B3FF:0E1E:8329
 - B. FE80::0202:B3FF::E1E:8329
 - C. FE80::202:B3FF:E1E:8329
 - D. FE80::0202:B8FF:0E1E:8329

2. If the router's MAC address is 0012.7feb.6b40, which of the following is the correct EUI-64 format for the IPv6 link-local interface address?
 - A. ::0012:7FEB:6B40
 - B. 2001:DB8::212:7FFF:FEEB:6B40
 - C. 2001:DB8::212:7FFF:FE80:6B40
 - D. 2001:DB8::2012:7FEB:6B40
 - E. 2001:DB8::212:7FFF:0000:6B40

3. Which of the following are valid transition strategies when moving from IPv4 to IPv6? Choose all that apply.
 - A. NAT-PT
 - B. SNARD encapsulation
 - C. 4-in-4-out tunneling
 - D. Short stacking
 - E. 6-to-4 tunneling
 - F. Dual stacking

Cram Quiz Answers

1. Answers A and C are correct. Answer B is incorrect because it uses the "::" twice. Answer D is incorrect because of the B8FF, which should be B3FF:.
 2. Answer B is correct. Answer A is incorrect because it omits the network number. Answer C is wrong because it uses the wrong value for the EUI expansion. (It uses FE80 instead of the correct FFFE.) Answer D is incorrect because it uses the wrong link-local network number, adding 0x2 to the end of the link-local identifier.
 3. Answers A, E, and F are correct. Answers B, C, and D are incorrect; I just made up those terms.
-

Review Questions

1. Which of the following are alternate representations of the decimal number 227? Choose two.
 - A. 0x227
 - B. 11100011
 - C. 0x143
 - D. 0xE3
 - E. 11100110
2. Which of the following are alternate representations of 0xB8? Choose two.
 - A. 10110100
 - B. 10111111
 - C. 10111000
 - D. 184
 - E. 0x184
3. You have been asked to create a subnet that supports 16 hosts. What subnet mask should you use?
 - A. 255.255.255.252
 - B. 255.255.255.248
 - C. 255.255.255.240
 - D. 255.255.255.224
4. Given the mask 255.255.254.0, how many hosts per subnet does this create?
 - A. 254
 - B. 256
 - C. 512
 - D. 510
 - E. 2
5. You are a senior network engineer at True North Technologies. Your boss, Mr. Martin, asks you to create a subnet with room for 12 IPs for some new managers. Mr. Martin promises that there will never be more than 12 managers, and he asks you to make sure that you conserve IP address space by providing the minimum number of possible host IPs on the subnet. What subnet mask will best meet these requirements?
 - A. 255.255.255.12
 - B. 255.255.255.0

- C. 255.255.240.0
 - D. 255.255.255.240
 - E. 255.255.255.224
6. Your boss, Duncan, does not seem to be able to grasp subnetting. He comes out of a management meeting and quietly asks you to help him with a subnetting issue. He needs to divide the Class B address space the company uses into six subnets for the various buildings in the plant, while keeping the subnets as large as possible to allow for future growth. What is the best subnet mask to use in this scenario?
- A. 255.255.0.0
 - B. 255.255.248.0
 - C. 255.255.224.0
 - D. 255.255.240.0
 - E. 255.255.255.224
7. You have purchased several brand-new Cisco routers for your company. Your current address space is 172.16.0.0 /22. Because these new routers support the **ip subnet zero** command, you realize you are about to gain back two subnets that you could not use with the old gear. How many subnets total will be available to you once the upgrades are complete?
- A. 4
 - B. 2
 - C. 32
 - D. 62
 - E. 64
8. Which of the following are true about the following address and mask pair: 10.8.8.0 /24? Assume that all subnets are available. Choose all that apply.
- A. This is a Class B address.
 - B. This is a Class A address.
 - C. This is a Class C address.
 - D. 16 bits were stolen from the host field.
 - E. 24 bits were stolen from the host field.
 - F. The default mask for this address is 255.0.0.0.
 - G. The mask can also be written as 255.255.255.0.
 - H. The mask creates 65,536 subnets total from the default address space.
 - I. Each subnet supports 256 valid host IPs.
 - J. Each subnet supports 254 valid host IPs.

9. Indy and Greg have configured their own Windows 8 PCs and connected them with crossover cables. They can't seem to share their illegally downloaded MP3 files, however. Given their configurations, what could be the problem?

Indy's configuration:

IP:192.168.0.65

Mask:255.255.255.192

Greg's configuration:

IP:192.168.0.62

Mask:255.255.255.192

- A. Indy is using a broadcast ID for his IP.
 - B. Greg is using an invalid mask.
 - C. Indy's IP is in one of the zero subnets.
 - D. Greg and Indy are using IPs in different subnets.
10. You are given an old router to practice for your CCNA. Your boss, Dave, has spent a lot of time teaching you subnetting. Now he challenges you to apply your knowledge. He hands you a note that says the following:

"Given the subnetted address space of 192.168.1.0 /29, give the E0 interface the first valid IP in the eighth subnet. Give the S0 interface the last valid IP in the 12th subnet. The zero subnets are available. You have 10 minutes. Go."

Which two of the following are the correct IP and mask configurations? Choose two.

- A. E0: 192.168.1.1 255.255.255.0
 - B. E0: 192.168.1.56 255.255.255.248
 - C. E0: 192.168.1.57 255.255.255.248
 - D. S0: 192.168.1.254 255.255.255.0
 - E. S0: 192.168.1.95 255.255.255.248
 - F. S0: 192.168.1.94 255.255.255.248
11. The following questions are part of a Subnetting SuperChallenge. This monster question will stretch your subnetting skills, especially if you give yourself a time limit. Start with 10 minutes and see if you can get down to 5.

The Vancouver Sailing Company has four locations: a head office and three branch offices. Each of the branches is connected to the head office by a point-to-point T1 circuit. The branches have one or more LANs connected to their routers. The routers are called Main, Jib, Genoa, and Spinnaker. The company has been assigned the 172.16.0.0/20 address space to work within.

Your task is to choose the correct IP address and mask for each interface, based on the information provided. Remember that no IP address may overlap with any address in another subnet, and that the required number of hosts for each subnet will affect your decision as to which address to use.

Here are the known IP configurations for the routers:

Main:

S0/0:172.16.0.1 /30

S0/1:172.16.0.5 /30

S0/2:172.16.0.9 /30

Fa1/0:172.16.4.1 /23

Fa1/1:172.16.6.1 /23

Jib:

S0/0: Connects to Main S0/0

Fa1/0: 172.16.8.33/27

Fa1/1: 30 hosts needed

Genoa:

S0/0: Connects to Main S0/1

Fa1/0: 172.16.8.129/26

Fa1/1: 100 hosts needed

Fa2/0: 100 hosts needed

Fa2/1: 172.16.13.0/24.

Spinnaker:

S0/0: Connects to Main S0/2

Fa1/0: 500 hosts needed

Choose the correct IP and mask assignments for each router:

- A. Jib Fa1/1: 172.16.8.62/27
- B. Jib Fa1/1: 172.16.8.64/27
- C. Jib Fa1/1: 172.16.8.65/28
- D. Jib Fa1/0: 172.16.8.65/27
- E. Jib Fa1/1: 172.16.8.65/27
- F. Genoa S0/0: 172.16.0.2/30
- G. Genoa S0/0:172.16.0.6/30
- H. Genoa Fa1/1:172.16.12.1/26
- I. Genoa Fa1/1:172.16.12.1/25
- J. Genoa Fa2/0:172.16.12.129/24
- K. Genoa Fa2/0:172.16.12.129/25
- L. Genoa Fa2/1:172.16.12.193/25
- M. Genoa Fa2/1:172.16.13.1/25
- N. Genoa Fa0/2:172.16.13.1/24
- O. Spinnaker S0/0: 172.16.0.10/30

- P. Spinnaker S0/0: 172.16.0.12/30
 - Q. Spinnaker Fa1/0: 172.16.13.0/23
 - R. Spinnaker Fa1/0: 172.16.14.0/23
 - S. Spinnaker Fa1/0: 172.16.14.1/23
12. Which of the following is the best summary statement for the following range of networks?
- 192.168.1.0 /24–192.168.15.0 /24
- A. 192.168.1.0
 - B. 192.168.1.0 255.255.240.0
 - C. 192.168.1.0 0.0.15.0
 - D. 192.168.1.0 255.255.248.0
 - E. 192.168.0.0 255.255.240.0
13. Which of the following is the best summary statement for the following range of networks?
- 192.168.24.0 /24–192.168.31.0 /24
- A. 192.168.24.0 255.255.240.0
 - B. 192.168.24.0 /28
 - C. 192.168.24.0 /21
 - D. 192.168.0.0 /27
14. Which of the following networks are included in the summary 172.16.0.0 /13? Choose all that apply.
- A. 172.0.0.0 /16
 - B. 172.16.0.0 /16
 - C. 172.24.0.0 /16
 - D. 172.21.0.0 /16
 - E. 172.18.0.0 /16
15. What are the advantages of route summarization? Choose three.
- A. Ensures job security for network admins because of difficulty of configuration
 - B. Reduces routing update traffic overhead
 - C. Reduces the impact of discontinuous subnets
 - D. Reduces CPU and memory load on routers
 - E. Identifies flapping interfaces
 - F. Hides network instability

16. Which of the following is a valid IPv6 address format?
- A. G412:AFFA:2001:0000:0000:0000:0001
 - B. 2001:8888:EEEE:1010:0000:0000:0000:0001
 - C. 2001::8888::1
 - D. 2010:2112:5440:1812:1867
17. Which of the following is a valid IPv6 unicast address format?
- A. 2001:8888:EEEE:1010:0000:0000:0000:0001
 - B. FF00:0000:0000:0002:00CO:00A8:0001:0042
 - C. 2001:8888::2FFE::00A8
 - D. FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF
18. Which of the following are valid IPv6 address formats for the same address?
Choose three.
- A. 2001:0000:0000:0200:0222:0000:0000:0001
 - B. 2001:0000:0000:02:0222:0000:0000:0001
 - C. 2001:0:0:200:222:0:0:1
 - D. 2001::200:222::1
 - E. 2001:0:0:200:222::1
19. Which of the following is a valid IPv6 unicast address?
- A. FF00:2112:1812:5440::1
 - B. 1812:2112:5440:1
 - C. 255:255:255:255:255:255:255:255
 - D. None of the above

Answers to Review Questions

1. Answers B and D are correct. Answer A in decimal would be 551. Answer C in decimal would be 323. Answer E in decimal is 230.
2. Answers C and D are correct. Answer A in hex is 0xB4. Answer B in hex is 0xBF. Answer E is simply an attempt to trick you—the correct decimal answer is incorrectly expressed as a hex value.
3. Answer D is correct. A will only support 2 hosts; B only 6, and C only 14. Watch out for the minus 2 in the host calculation! Answer C creates 16 IP addresses on the subnet, but we lose 2—one for the network ID and one for the broadcast ID.
4. Answer D is correct. The mask 255.255.254.0 gives us nine 0s at the end of the mask; $2^9 - 2 = 510$. Answer A is checking to see if you missed the 254 in the third octet because you are used to seeing 255. Answer B does the same thing plus tries to catch you on not subtracting 2 from the host calculation. Answer C tries to catch you on not subtracting 2, and Answer E is the increment of the given mask that you might pick if you were really off track.
5. Answer D is correct. Disregarding for the moment the possibility that Mr. Martin might be wrong, let's look at the requirements. He says make room for 12 managers, and make the subnets as small as possible while doing so. You need to find the mask that has sufficient host IP space without making it bigger than necessary. Answer A is invalid; 12 is not a valid mask value. Remember, a mask is a continuous string of 1s followed by a continuous string of 0s. In Answer B, the mask is valid, but it is not correct. This mask has eight 0s at the end, which, when we apply the formula $2^8 - 2$ gives us 254 hosts. That makes more than enough room for the 12 managers, but does not meet the “as small as possible” requirement. Answer C has the correct mask value in the wrong octet. That mask gives us eight 0s in the fourth octet, plus another four in the third octet; that would give us 4094 hosts on the subnet. Answer E gives us 30 hosts per subnet, but that only meets half the requirement. This mask does not provide the minimum number of hosts.
6. Answer C is correct. The default mask for a Class B is 255.255.0.0. Answer C extends that mask by three bits, creating eight subnets ($2^3 = 8$). Although we only need six of them, we have to use the mask that creates eight because the next smaller mask would only create four, and that isn't enough. Answer A is incorrect because it is the default mask for a Class B and not subnetted at all. Answers B and D are incorrect because although they create sufficient subnets, they do not maximize the number of hosts per subnet and so are not the best answer. Answer E uses the correct mask in the wrong octet.
7. Answer E is correct. With **ip subnet zero** enabled, all 64 subnets created by the mask in use become available. Answers A, B, and C are not even close and are simply distracters. Answer D wants to catch you by subtracting the zero subnets.
8. Answers B, D, F, G, H, and J are correct. Answers A and C are incorrect because this is a Class A address. Answer E is incorrect because only 16 bits were stolen. Answer I is incorrect because it does not subtract the two IPs for the network ID and broadcast ID.

9. Answer D is correct. With that mask, the increment is 64. Greg is in the first subnet, and Indy is in the second. Without a router between them, their PCs will not be able to communicate above Layer 2. Answer A is incorrect; the broadcast ID for Indy would be .63. Answer B is incorrect; nothing is wrong with the mask. Answer C is incorrect; the zero subnets are the first and last created, and Indy is in the second subnet. The question does mention the zero subnets, so we can use them, and in any case Windows 8 fully supports them.
10. Answers C and F are correct. This is an increment question. The increment here is 8, so you should start by jotting down the multiples of 8 (those are all the network IDs), and then noting what 1 less than each of the network IDs is (those are the broadcast IDs). From there, it is easy to find what the first and last IPs in each subnet are. (Remember that Dave says we can use the zero subnets.)

The eighth subnet network ID is 192.168.1.56; the first valid IP is 192.168.1.57. The twelfth subnet network ID is 192.168.1.88; the last valid IP is 192.168.1.94. Answers A and D are incorrect because they do not use the subnetted address space Dave requested. Answer B is incorrect because it is a network ID. Answer E is incorrect because it is a broadcast ID.

11. SuperChallenge answers:

- A. Incorrect (same subnet as Fa1/0)
- B. Incorrect (network ID)
- C. Incorrect (not enough hosts)
- D. Incorrect (Fa1/0 IP already assigned)
- E. Correct
- F. Incorrect (wrong subnet—not on the same network as the connected interface on Main)
- G. Correct
- H. Incorrect (not enough hosts)
- I. Correct
- J. Incorrect (overlaps with Fa1/1)
- K. Correct
- L. Incorrect (overlaps with Fa2/0)
- M. Incorrect (wrong mask)
- N. Incorrect (no Fa0/2 interface on Genoa)
- O. Correct
- P. Incorrect (network ID)
- Q. Incorrect (overlaps with Genoa)
- R. Incorrect (network ID)
- S. Correct

12. Answer E is correct. Answers A and C use incorrect syntax; Answer D uses the wrong mask. Answer B looks correct, but it does not use the correct network ID; the range should always start at a binary increment (in this case 0, not 1). (In other words, this is a case of “best answer.”) Note that the correct summary does include the 192.168.0.0/24 network as well (not just 192.168.1–15.0/24). This is intended to confuse and distract you!
13. Answer C is correct. Answer A uses the wrong mask and summarizes more than the specified networks. Answer B subnets instead of summarizes. Answer D uses the wrong address and mask.
14. Answers B, D, and E are correct. The networks in Answers A and C are out of the range, which is 172.16.0.0 through 172.23.0.0.
15. Answers B, D, and F are correct. Answer A might have an element of truth, but Cisco does not have much of a sense of humor. Answer C is incorrect because discontinuous subnets are a real problem if you intend to summarize. Answer E is incorrect; route summarization does not identify but rather hides the effects of flapping interfaces.
16. Answer B is correct. IPv6 addresses must have (or at least indicate, perhaps with ::) eight sets of four valid hex characters. Answer A is wrong because G is not a valid hex character. Answer C is wrong because it uses the :: notation twice, which is invalid. Answer D is wrong because it uses only five sets.
17. Answer A is correct. IPv6 addresses must have eight sets of four valid hex characters. Answer B is wrong because it starts with FF, which indicates a multicast address, not a unicast. Answer C is wrong because it uses :: twice, which is invalid. Answer D is wrong because there are only seven sets.
18. Answers A, C, and E are correct. These are the same address, represented in three valid notations: A is not truncated, C has dropped leading 0s, and E has compressed the contiguous all-zero groups with the ::. Answer B is wrong because it drops trailing 0s, not leading 1s. D is wrong because it uses the :: twice, which is invalid.
19. Answer D is correct. None of these is a valid unicast format. Answer A is an IPv6 multicast (starts with FF00/8). Answer B has only four sets instead of eight. Note that if there had been a double colon before the last 1, it could have been correct. Answer C uses the decimal 255 to confuse you; it could have been correct except that there are nine sets.

What Next?

If you want more practice on this chapter's exam topics before you move on, remember that you can access all of the Cram Quiz questions on the CD. You can also create a custom exam by topic with the practice exam software. Note any topic you struggle with and go to that topic's material in this chapter.

This page intentionally left blank

Index

Symbols

0x (hexadecimal characters), 144

802.1Q trunk encapsulation standard, 214

A

AAA protocols, 177

access layer (Cisco hierarchical networks), 40

access (off) mode (switch ports), switch trunking configuration, 215

acknowledgment numbers, 48

ACL (Access Control Lists), 295-297, 366

DDR, 289

established keyword, 304

extended ACL, 293-294

configuring, 300-302

troubleshooting configurations, 308-309

verifying configurations, 308-309

log-input keyword, 304

log keyword, 304

named ACL, 294

configuring, 305

troubleshooting configurations, 308-309

verifying configurations, 308-309

NAT, 289

packet filtering, 290-292

established keyword, 304

extended ACL, 293-294, 300-302, 308-309

log-input keyword, 304

log keyword, 304

named ACL, 294, 305, 308-309

SSH filtering, 303

standard ACL, 293, 297-299, 308-309

Telnet filtering, 303

vtv lines, 303-304

purposes of, 289

SSH filtering, 303

standard ACL, 293

configuring, 297-299

troubleshooting configurations, 308-309

verifying configurations, 308-309

wildcard masks, 298-299

Telnet filtering, 303

troubleshooting, 365-366

vtv lines, 303-304

address classes

default subnet masks, 79

IP addressing, 77-78

subnetting, 83-84

address learning, switches and, 192

address ranges

IP addressing, 77-78

private IP addresses, 78

administrative distance, 232-234

administratively down. See shutdown state

Ambiguous command error message, 148

answers

practice exam 1, 413-422

practice exam 2, 459-467

anycast IPv6 addressing, 114

application layer (OSI model), 45-46

ARP (Address Resolution Protocol), 56

authentication, 177

Auto mode (switch ports), switch trunking configuration, 216

B

backoff algorithm, defining, 29

bandwidth, 161, 250

BDR (Backup Designated Routers), OSPF and, 252-253

binary numbering systems, 70-76

boot field (configuration register), 144

boot process, 143-144

bridges, 190

broadcast domains, 207

cut-through switching mode, 197

defining, 191

duplex configuration, 199-200

fragment-free switching mode, 197

full duplex switch connections, 198-199

half duplex switch connections, 198

store-and-forward switching mode, 197

switches versus, 195-196

transparent bridges, 192

broadcasts

broadcast domains, defining, 207

broadcast ID, 88-89

broadcast storms, 200

frame forwarding, 193

bus topologies, 27

C

cabling

crossover cable, trunking, 213

Layer 1 cabling, troubleshooting, 348

CDP (Cisco Discovery Protocol), 155-156

CEF (Cisco Express Forwarding), routing, 239-240

cells, defining, 32

cell-switching, defining, 32

CIDR (Classless Interdomain Routing) notation, 80

circuit-switching, defining, 31

Cisco devices

boot process, 143-144

configuring

command-line modes, 146-147

command shortcuts, 148

configuration register, 144

password recovery, 145-146

routers, 160-162

setup mode, 144

switches, 159-160

connecting, 141

exec sessions, 146

IOS startup process, 143-144

memory, 142

rebooting, 146

remote management, 152

CDP, 155-156

ICMP, 156-158

NetFlow, 159

NTP, 158

SNMP, 158

SSH, 154-155

Syslog, 159

Telnet, 152-153

routers, configuring, 160-162

switches, configuring, 159-160

Cisco hierarchical network design, 43

- access layer, 40
- advantages of, 42
- core layer, 41
- costs, 42
- distribution layer, 41
- scalability, 42
- troubleshooting, 42

collision domains, 29**collisions, 29**

- backoff algorithm, 29
- CSMA/CD, 29
- defining, 28
- microsegmented switches, 30
- troubleshooting, 350-351

command-line modes

- Cisco devices, configuring, 146
- interface configuration mode, 147
- line configuration mode, 147
- router configuration mode, 147

command shortcuts, Cisco device configuration, 148**configuring**

- Cisco devices
 - command-line modes, 146-147
 - command shortcuts, 148
 - configuration register, 144
 - password recovery, 145-146
 - routers, 160-162
 - setup mode, 144
 - switches, 159-160
- duplexes, 199-200
- extended ACL, 300-302, 308-309
- hosts, troubleshooting, 361-362
- IPv6 addresses, 115-116
- named ACL, 305, 308-309
- NAT, verifying configurations, 330
- NTP, verifying configurations, 336-337

OSPF

- single-area OSPFv2, 254-258
- single-area OSPFv3, 258-260
- verifying configuration, 260
- verifying OSPFv3 configuration, 261-263

PAT, 329-330**routers**

- enabling interfaces, 161-162
- serial port configuration, 160
- standard ACL, 297
- troubleshooting configurations, 308-309
- verifying configurations, 308-309
- wildcard masks, 298-299

switches

- default gateway configuration, 159-160
- Management IP configuration, 159-160
- setting Ethernet port speed and duplex, 160
- setting hostnames, 159
- trunking, verifying configuration/operation, 357-360

connected networks

- administrative distance, 234
- routing and, 228-229

connection-oriented data transmission, 47**contiguous subnetting, 101****converting**

- decimal to hexadecimal numbering systems (IP addressing), 75
- hexadecimal to binary numbering systems (IP addressing), 75-76
- hexadecimal to decimal numbering systems (IP addressing), 75-76

core layer (Cisco hierarchical networks), 41

costs

- Cisco hierarchical networks, 42
- cost metric
 - OSPF, 249
 - routing, 235

CRC (Cyclic Redundancy Checks)

- errors, troubleshooting, 348, 351
- store-and-forward switching mode, 197

crossover cable, trunking, 213**CSMA/CD (Carrier Sense Multiple Access with Collision Detection), 29****cut-through switching mode, 197****D****data flow visualization (networks), 271**

- host-to-host direct communications, 272-274
- host-to-host through router communications, 278-282
- host-to-host through switch communications, 275-277

data link layer (OSI model), MAC addresses, 56-57**data transmission**

- acknowledgment numbers, 48
- connection-oriented transmission, 47
- destination ports, 51-52
- OSI network design model, 58-59
- PAR, 48
- reliable transmission
 - TCP, 47-58, 305
 - UCP, 58
 - UDP, 51-54
- sequence numbers, 48
- source ports, 51-52
- three-way handshakes, 48
- unreliable transmission, UDP, 47-48

DDR (Dial-on-Demand Routing), 289**decapsulated data, OSI network design model, 59****decimal numbering systems, IP addressing, 74**

- decimal to hexadecimal conversions, 75
- hexadecimal to decimal conversions, 75-76

dedicated leased line connections. See leased line connections**default gateways**

- defining, 160
- host configurations, verifying, 362
- switch configuration, 159-160

default routes, 230-232**default subnet masks, IP addressing, 78-79****deleting VLAN, 209****Desirable mode (switch ports, switch trunking configuration), 216****Destination Address field (IPv6 headers), 119****destination ports, 51-52****determination (path), 54****DHCP (Dynamic Host Configuration Protocol), 338**

- dynamic IPv6 address configuration, 116
- IOS routers, configuring as
 - DHCP clients, 338
 - DHCP servers, 339-340
- verifying operation of, 340-341

directed broadcasts. See broadcasts, broadcast ID**distribution layer (Cisco hierarchical networks), 41****DNS (Domain Naming System), 46****DR (Designated Routers), OSPF and, 251-253****DTP (Dynamic Trunking Protocol), 215****duplexes**

- configuring, 199-200
- full duplex switch connections, 198-199
- half duplex switch connections, 198
- troubleshooting, 349-351

duplicate frames, switching loops, 200
dynamic IPv6 address configuration, 116
dynamic NAT (Network Address Translation), 326-327
dynamic routes, 232
dynamic VLAN membership, 211

E

EIGRP (Enhanced Interior Gateway Routing Protocol), 56
 external, administrative distance, 234
 feasible distance metric, 235
 internal, administrative distance, 234

encapsulated data, OSI network design model, 58-59

error messages, 148

established keyword, ACL advanced options, 304

Ethernet
 collisions, 28-29
 CSMA/CD, 29
 development of, 28
 giants, defining, 351
 implementation comparison table, 30
 ports, switch configuration, 160
 runts, defining, 351

EUI-64, dynamic and static IPv6 address configuration, 116

eui-64 switch (routers), IPv6 router configuration, 116

exams
 practice exam 1
 answers, 413-422
 questions, 379-411
 practice exam 2
 answers, 459-467
 questions, 425-457
 strategies for taking, 377-378, 423-424

exec sessions, 146

exit interface (next hop concept), 227
extended ACL (Access Control Lists), 293-294, 300-302
extended NAT entry. See PAT
external authentication, 177

F

fast switching, 239-240
feasible distance metric, 235
Flash memory, 142-144
floating static routes, 233
Flow Label field (IPv6 headers), 118
fragment-free switching mode, 197
frames, 58
 defining, 32
 filtering, switches and, 194
 forwarding, switches and, 193-194

FTP (File Transfer Protocol), 46

full duplex switch connections, 198-199

full mesh topologies, 27

G - H

giants (Ethernet), defining, 351

half duplex switch connections, 198

headers
 defining, 52
 IPv6 addressing, 117-119
 packets
 defining, 55
 headers, 58
 TCP headers, 52-54
 UDP, 53-54

help, 149

hexadecimal characters, 0x notation, 144

hexadecimal numbering systems, IP addressing, 74
 decimal to hexadecimal conversions, 75

hexadecimal numbering systems, IP addressing

hexadecimal to binary conversions,
75-76

hexadecimal to decimal conversions,
75-76

hierarchical network design (Cisco), 43

access layer, 40

advantages of, 42

core layer, 41

costs, 42

distribution layer, 41

scalability, 42

troubleshooting, 42

hierarchical routing, OSPF and, 250-251

high-order bits, 71

hop count metric, 235

Hop Limit field (IPv6 headers), 119

host devices, defining, 26

host field, IP addressing, 80

hostnames, switch configuration, 159

host routes, 236

hosts

configuring, troubleshooting,
361-362

defining, 86

IP addressing

host subnetting requirements, 91

number of hosts, 86-87

network data flow visualization, 271

host-to-host direct
communications, 272-274

host-to-host through router
communications, 278-282

host-to-host through switch
communications, 275-277

subnetting

host requirements, 91

number of hosts, 86-87

troubleshooting configurations,
361-362

HTTP (Hypertext Transfer Protocol), 46

ICANN (Internet Corporation for Assigned Network Numbers), IPv6 address allocation, 112

ICMP (Internet Control Message Protocol), 56, 117, 156-158

IEEE 802.1Q trunk encapsulation standard, 214

IMAP (Internet Message Access Protocol), 46

Incomplete command error message, 148

increments (subnetting), 84-86

Inside (NAT terminology), 322

Inside global (NAT terminology), 322

Inside local (NAT terminology), 322

interesting octets, subnetting, 85

interface configuration mode, 147

Internet as WAN, 32

inter-VLAN routing, 242

L3 switching, 245-246

router on a stick, 243-244

upstream routing, 246

Invalid input command error message, 148

inverse masks. See wildcard masks

IOS (Internetwork Operating System)

error messages, 148

help, 149

licensing, 150-151

routers, configuring as

DHCP clients, 338

DHCP servers, 339-340

NTP clients, 335

startup process, 143-144

IP (Internet Protocol), 56

IP access lists. See ACL

IP addressing

address classes, 77-78

default subnet masks, 79

subnetting, 83-84

- address ranges, 77-78
- binary numbering systems, 70-76
- CIDR notation, 80
- decimal numbering systems, 74
 - decimal to hexadecimal conversions, 75
 - hexadecimal to decimal conversions, 75-76
- hexadecimal numbering systems, 74
 - decimal to hexadecimal conversions, 75
 - hexadecimal to binary conversions, 75-76
 - hexadecimal to decimal conversions, 75-76
- host field, 80
- hosts, number of, 86-87
- IPv6 addressing, 110-111
 - address allocation, 112
 - address configuration, 115-116
 - address notation, 113-114
 - anycast IPv6 addresses, 114
 - headers, 117-119
 - ICMP, 117
 - IPsec, 117
 - mobility, 117
 - multicast IPv6 addresses, 114
 - RA messages, 117
 - router configuration, 116
 - RS messages, 117
 - transitioning strategies, 119-120
 - translation, 120
 - tunneling strategies, 120
 - unicast IPv6 addresses, 114
- network field, 79-80
- private IP addresses, 78
- public IP addresses, 78
- remotely changing, 362
- route summarization, 100
 - guidelines, 102-103
 - summary addresses, 104-109
- slash notation, 80
- subnet field, 81
- subnet masks
 - default masks, 78-79
 - nondefault masks, 80
- subnetting, 82
 - address classes, 83-84
 - broadcast ID, 88-89
 - contiguous subnetting, 101
 - host requirements, 91
 - increments, 84-86, 92-93
 - interesting octets, 85
 - number of hosts, 86-87
 - number of subnets formula, 89-91
 - subnet masks, 83-84
 - subnet requirements, 92
 - subnetting chart, 93-95
 - Subnet Zero Rule, 90-94
- troubleshooting, 361
 - NAT, 363
 - remotely changing IP addresses, 362
 - verifying hosts have correct IP address, 362
 - verifying host configurations on default gateways, 362
- VLSM, 96-98
- IP route table, 237**
- IPsec, 117**
- IP services**
 - DHCP
 - configuring IOS routers as DHCP clients, 338
 - configuring IOS routers as DHCP servers, 339-340
 - verifying operations of, 340-341
 - NAT, 320
 - advantages of, 324
 - applications of, 323-324
 - disadvantages of, 324
 - dynamic NAT, 326-327
 - NAT pools, 326
 - overlapping address space, 323

- PAT, 328-333
- static NAT, 324-326
- terminology, 321-322
- troubleshooting, 330-333
- verifying configurations, 330
- well-meaning admin errors, 323

NTP, 334

- configuring IOS routers as NTP clients, 335
- organizing clocks by stratum, 335
- true time, 335
- UTC, 335
- verifying configurations, 336-337

ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) tunneling, IPv6 addressing, 120

ISIS (Intermediate System to Intermediate System), 56

J - K - L

L3 switching, 245-246

LAN (Local-Area Networks)

- defining, 28
- Ethernet
 - collisions, 28-29
 - CSMA/CD, 29
 - development of, 28
 - implementation comparison table, 30
- wireless networks, 32

late collisions, troubleshooting, 350-351

Layer 1 physical layer (OSI model), 57

- cabling, 348
- duplexes, 349-351
- link state, 350
- troubleshooting, 348-352

Layer 2 data link layer (OSI model), MAC addresses, 56-57

Layer 3 network layer (OSI model), 54-55

Layer 4 transport layer (OSI model)

- multiplex connections, 52
- TCP, 47-48
 - headers, 52-54
 - port numbers, 51
 - segments, 58
 - sliding window, 50
- UCP segments, 58
- UDP, 47-48
 - headers, 53-54
 - port numbers, 51

Layer 5 session layer (OSI model), 47

Layer 6 presentation layer (OSI model), 47

Layer 7 application layer (OSI model), 45-46

leased line connections, defining, 31

licensing (IOS), managing, 150-151

line configuration mode, 147

link state, determining, 350

log-input keyword, ACL advanced options, 304

log keyword, ACL advanced options, 304

longest match analysis, routing and, 237

loops (switching), 200

low-order bits, 71

LSA (Link State Advertisements), 249-250

M

MAC addresses

- defining, 56
- OUI, 57
- switches, 192

MAC table, switching loops, 200

Management IP, switch configuration, 159-160

managing

- IOS licensing, 150-151
- remote management

- CDP, 155-156
- ICMP, 156-158
- NetFlow, 159
- NTP, 158
- SNMP, 158
- SSH, 154-155
- Syslog, 159
- Telnet, 152-153

memory, 142-144

mesh topologies, 27

metrics

- cost, 235
- cost metric, 249
- feasible distance, 235
- hop count, 235
- routing metrics, 234

microsegmentation. See full duplex switch connections

microsegmented switches, 30

mobility, IPv6 addressing, 117

multicasts

- frame forwarding, 194
- IPv6 addressing, 114

multiplex transport layer connections, 52

N

named ACL (Access Control Lists), 294, 305

naming VLAN, 209

NAT (Network Address Translation), 289, 320-321, 363

- advantages of, 324
- applications of, 323-324
- disadvantages of, 324
- dynamic NAT, 326-327
- extended NAT entry. *See* PAT
- NAT pools, 326
- overlapping address space, 323
- PAT, 328
 - configuring, 329
 - troubleshooting, 330-333

- verifying configurations, 330
- static NAT, 324-326
- terminology, 321-322
- troubleshooting, 330-333, 363
- verifying configurations, 330
- well-meaning admin errors, 323

NetFlow, 159

network devices, defining, 26

network field, IP addressing, 79-80

network layer (OSI model), 54-55

networks

- connected networks
 - administrative distance, 234
 - routing and, 228-229
- defining, 26
- hierarchical network design, 43
 - access layer, 40
 - advantages of, 42
 - core layer, 41
 - costs, 42
 - distribution layer, 41
 - scalability, 42
 - troubleshooting, 42
- host data flow visualization, 271
 - host-to-host direct communications, 272-274
 - host-to-host through router communications, 278-282
 - host-to-host through switch communications, 275-277

LAN

- defining, 28
- Ethernet, 28-30
- wireless networks, 32

OSI network design model, 44-45, 60

- data transmission, 58-59
- decapsulated data, 59
- encapsulated data, 58
- Layer 1 physical layer, 57
- Layer 2 data link layer, 56-57
- Layer 3 network layer, 54-55

- Layer 4 transport layer, 47-54
- Layer 5 session layer, 47
- Layer 6 presentation layer, 47
- Layer 7 application layer, 45-46
- TCP/IP model comparisons, 62

TCP/IP network design model, 61-62

topologies, 27

VLAN, 204-205, 211

- adding to switch ports, 210
- benefits of, 206-207
- creating, 209
- defining, 206
- deleting, 209
- dynamic VLAN membership, 211
- implementing, 208-210
- naming, 209
- port assignments, 354-356
- static VLAN membership, 210
- switches with preconfigured VLAN, 210
- troubleshooting, 353-356
- trunking, 206-207, 213-216, 357-360
- verifying, 353-356
- VMPS, 211

WAN

- cell-switching, 32
- circuit-switching, 31
- defining, 31
- Internet as, 32
- leased line connections, 31
- packet-switching, 32
- wireless networks, 32

Next Header field (IPv6 headers), 119

next hop concept (routing), exit interface, 227

NMS (Network Management Stations), SNMP, 158

noise, troubleshooting, 350-351

nondefault subnet masks, IP addressing, 80

NoNegotiate mode (switch ports), switch trunking configuration, 216

notation (IPv6 address), 113-114

NTP (Network Time Protocol), 46, 158, 334

IOS routers, configuring as NTP clients, 335

stratum, organizing clocks by, 335

true time, 335

UTC, 335

verifying configurations, 336-337

NVRAM (NonVolatile Random Access Memory), Cisco devices, 142

O

OSI network design model, 44, 60

data transmission, 58-59

decapsulated data, 59

encapsulated data, 58-59

Layer 1 physical layer, 57

Layer 2 data link layer, MAC addresses, 56-57

Layer 3 network layer, 54-55

Layer 4 transport layer
multiplex connections, 52
TCP, 47-54

UDP, 47-48, 51-54

Layer 5 session layer, 47

Layer 6 presentation layer, 47

Layer 7 application layer, 45-46

TCP/IP model comparisons, 62

OSP (Open Settlement Protocol), cost metric, 235

OSPF (Open Shortest Path First), 56, 248

administrative distance, 234

BDR, 252-253

characteristics of, 249-250

cost metric, 249

DR, 251-253

hierarchical routing, 250-251

LSA, 249-250

- single-area OSPFv2, configuring, 254-258
- single-area OSPFv3, configuring, 258-260
- syntax of, 257
- troubleshooting, 261
- verifying
 - OSPFv2 configuration, 260
 - OSPFv3 configuration, 261-263

OUI (Organizational Unique Identifiers), 57

Outside global (NAT terminology), 322

Outside local (NAT terminology), 322

Outside (NAT terminology), 322

overlapping address space, NAT and, 323

P

packets, 58

- defining, 55
- filtering, ACL, 290-292
 - established keyword, 304
 - extended ACL, 293-294, 300-302, 308-309
 - log-input keyword, 304
 - log keyword, 304
 - named ACL, 294, 305, 308-309
 - SSH filtering, 303
 - standard ACL, 293, 297-299, 308-309
 - Telnet filtering, 303
 - vtv lines, 303-304
- switching, defining, 32

PAR (Positive Acknowledgment and Retransmission), 48

partial mesh topologies, 27

passive interfaces, routing and, 240

password recovery, configuring Cisco devices, 145-146

path determination, 54

PAT (Port Address Translation), 328

- configuring, 329

- troubleshooting, 330-333
- verifying configurations, 330

Payload length field (IPv6 headers), 119

PDU (Protocol Data Units), 59

physical layer (OSI model), 57

point-to-point topologies, 27

pools (NAT), 326

POP3 (Post Office Protocol 3), 46

ports

- destination ports, 51-52

- Ethernet ports, switch configuration, 160

- PAT, 328

- configuring, 329

- troubleshooting, 330-333

- verifying configurations, 330

- port-based VLAN membership.

- See* static VLAN membership

- port numbers, 51

- serial ports, router configuration, 160

- source ports, 51-52

- switch ports

- access (off) mode, 215

- adding VLAN to, 210

- Auto mode, 216

- Desirable mode, 216

- NoNegotiate mode, 216

- security, 183

- trunking configuration, 215-216

- trunk (on) mode, 216

- VLAN, verifying port assignments, 354-356

practice exams

- exam 1

- answers, 413-422

- questions, 379-411

- exam 2

- answers, 459-467

- questions, 425-457

- strategies for taking, 377-378, 423-424

presentation layer (OSI model)

presentation layer (OSI model), 47

private IP addresses, 78

privileged exec access, 146

process switching, 239-240

public IP addresses, 78

Q - R

questions

practice exam 1, 379-411

practice exam 2, 425-457

RAM (Random Access Memory), 142

RA (Router Advertisement) messages, IPv6 addressing, 117

rebooting Cisco devices, 146

redundancy, CRC and store-and-forward switching mode, 197

reliable data transmission

TCP, 47, 305

headers, 52-54

port numbers, 51

segments, 58

sliding window, 50

three-way handshakes, 48

UCP segments, 58

UDP

headers, 53-54

port numbers, 51

remote management

CDP, 155-156

ICMP, 156-158

NetFlow, 159

NTP, 158

SNMP, 158

SSH, 154-155

Syslog, 159

Telnet, 152-153

RID (router ID), 252-253

ring topologies, 27

RIP (Routing Information Protocol), 56

administrative distance, 234

hop count metric, 235

ROM (Read-Only Memory), 142

routers

boot process, 143-144

configuring

command-line modes, 146-147

command shortcuts, 148

configuration register, 144

enabling interfaces, 161-162

password recovery, 145-146

serial port configuration, 160

setup mode, 144

connecting, 141

default gateways, defining, 160

eui-64 switches, IPv6 router configuration, 116

exec sessions, 146

host to host through router communications (network data flow visualization), 278-282

IOS startup process, 143-144

IPv6 configuration, 116

memory, 142

RA messages, IPv6 addressing, 117

rebooting, 146

remote management

CDP, 155-156

ICMP, 156-158

NetFlow, 159

NTP, 158

SNMP, 158

SSH, 154-155

Syslog, 159

Telnet, 152-153

router configuration mode, 147

RS messages, IPv6 addressing, 117

shutdown state, 161

route summarization (IP addressing), 100

guidelines, 102-103

summary addresses, 104-109

routing, 226

- administrative distance, 232-234
- BDR, OSPF and, 252-253
- CEF, 239-240
- connected networks, 228-229
- default routes, 230-232
- defining, 227
- DR, OSPF and, 251-253
- dynamic routes, 232
- exit interface, 227
- fast switching, 239-240
- floating static routes, 233
- hierarchical routing, OSPF and, 250-251
- host routes, 236
- inter-VLAN routing, 242
 - L3 switching, 245-246
 - router on a stick, 243-244
 - upstream routing, 246
- IOS routers, configuring as
 - DHCP clients, 338
 - DHCP servers, 339-340
 - NTP clients, 335
- IP route table, 237
- L3 switching, 245-246
- longest match analysis, 237
- metrics, 234
 - cost, 235
 - cost metric, 249
 - feasible distance, 235
 - hop count, 235
- next hop concept, exit interface, 227
- OSPF, 248
 - BDR, 252-253
 - characteristics of, 249-250
 - configuring single-area OSPFv2, 254-258
 - configuring single-area OSPFv3, 258-260
 - cost metric, 249
 - DR, 251-253
 - hierarchical routing, 250-251

- LSA, 249-250
- summary LSA, 250
- syntax of, 257
- troubleshooting, 261
- verifying configuration, 260
- verifying OSPFv3 configuration, 261-263
- passive interfaces, 240
- process switching, 239-240
- RID, 252-253
- router lookup process, 235-238
- router on a stick, 243-244
- static routes, 229-230
 - administrative distance, 233-234
 - floating static routes, 233
- upstream routing, 246
- VLAN, inter-VLAN routing, 242-246
- VLSM, 239

RS (Router Solicitation) messages, IPv6 addressing, 117**runless switching. See fragment-free switching mode****runts (Ethernet), 198, 351****S****scalability, Cisco hierarchical networks, 42****security**

- authentication, 177
- password recovery, Cisco device configuration, 145-146
- switch port security, 183
- switches

segments, 58**sequence numbers, 48****serial port, routers configuration, 160****servers (DHCP), configuring IOS routers as, 339-340****session layer (OSI model), 47****setup mode, configuring Cisco devices, 144**

shortcuts (command), Cisco device configuration

shortcuts (command), Cisco device configuration, 148

shutdown state, routers, 161

slash notation, 80

sliding window (TCP), 50

SMTP (Simple Mail Transfer Protocol), 46

SNMP (Simple Network Management Protocol), 46, 158

Source field (IPv6 headers), 119

source ports, 51-52

SSH (Secure Shell), 154-155, 303

standard ACL (Access Control Lists), 293

configuring, 297-299

wildcard masks, 298-299

star topologies, 27

startup process (IOS), 143-144

static IPv6 address configuration, 116

static NAT (Network Address Translation), 324-326

static routes, 229-230

administrative distance, 233-234

floating static routes, 233

static VLAN membership, 210

store-and-forward switching mode, 197

STP (Spanning Tree Protocol), switches and Layer 2 loop removal, 200-201

stratum, NTP clock organization, 335

subnet field, IP addressing, 81

subnet masks

default masks, 78-79

nondefault masks, 80

route summarization, 100

guidelines, 102-103

summary addresses, 104-109

VLSM, 96-98, 239

subnetting, 82, 91

address classes, 83-84

broadcast IP, 88-89

contiguous subnetting, 101

host requirements, 91

hosts, number of, 86-87

increments, 84-86, 92-93

interesting octets, 85

number of subnets formula, 89-91

requirements, 92

subnet masks, 83-84

subnetting chart, 93-95

Subnet Zero Rule, 90-94

summarization (route), IP addressing, 100

guidelines, 102-103

summary addresses, 104-109

summary LSA (Link State Advertisements), 250

switches, 190-191

address learning, 192

bridges versus, 195-196

broadcast domains, 207

configuring

command-line modes, 146-147

command shortcuts, 148

default gateway configuration, 159-160

Management IP configuration, 159-160

setting Ethernet port speed and duplex, 160

setting hostnames, 159

cut-through switching mode, 197

defining, 191

duplex configuration, 199-200

exec sessions, 146

fragment-free switching mode, 197

frame filtering, 194

frame forwarding, 193-194

full duplex switch connections, 198-199

functions of, 191

half duplex switch connections, 198

host to host through switch communications (network data flow visualization), 275-277

- L3 switching, 245-246
- Layer 2 loop removal, 200-201
- loops, 200
- MAC addresses, 192
- microsegmented switches, 30
- ports
 - access (off) mode, 215
 - adding VLAN to, 210
 - Auto mode, 216
 - Desirable mode, 216
 - NoNegotiate mode, 216
 - security, 183
 - trunking configuration, 215-216
 - trunk (on) mode, 216
- store-and-forward switching mode, 197
- STP, 200-201
- trunking, 212-213
 - 802.1Q trunk encapsulation standard, 214
 - port modes, 215-216
- VLAN, 210

synchronizing time, NTP, 334

Syslog, 159

T

TCP (Transmission Control Protocol), 47, 305

- headers, 52-54
- port numbers, 51
- segments, 58
- sliding window, 50
- three-way handshakes, 48

TCP/IP network design model, 61-62

Telnet, 46, 152-153, 303

teredo tunneling, IPv6 addressing, 120

tests. See exams

TFTP (Trivial File Transfer Protocol), 46

three-way handshakes, 48

time stamps, VoIP, 335

time synchronization, NTP, 334

topologies, 27

Traffic Class field (IPv6 headers), 118

translation, IPv6 addressing, 120

transparent bridges, 192

transport layer (OSI model)

- multiplex connections, 52
- TCP, 47-48
 - headers, 52-54
 - port numbers, 51
 - segments, 58
 - sliding window, 50
- UCP segments, 58
- UDP, 47-48
 - headers, 53-54
 - port numbers, 51

troubleshooting

- ACL, 308-309, 365-366
- Cisco hierarchical networks, 42
- collisions, 350-351
- CRC errors, 348, 351
- duplexes, 349-351
- host configurations, 361-362
- IP addressing, 361
 - NAT, 363
 - remotely changing IP addresses, 362
 - verifying host configurations on default gateways, 362
 - verifying hosts have correct IP address, 362
- Layer 1, 352
 - cabling, 348
 - collisions, 350-351
 - duplexes, 349-351
 - link state, 350
 - noise, 350-351
- link state, 350
- NAT, 330-333, 363
- noise, 350-351
- OSPF, 261
- PAT, 330-333

troubleshooting

trunking

- administrative mode, 358-360
- links, 357
- operational mode, 358-360
- VLAN existence, 357

VLAN, 353-356

true time, NTP and, 335**trunking, 212-213**

- 802.1Q trunk encapsulation standard, 214
- cabling, 213
- configuring, verifying configuration/operation, 357-360
- DTP, 215
- switch configuration, port modes, 215-216
- troubleshooting
 - administrative mode, 358-360
 - links, 357
 - operational mode, 358-360
- VLAN, 206-207, 357

trunk (on) mode (switch ports), switch trunking configuration, 216**tunneling strategies (IPv6 addressing), 120****U****UDP (User Datagram Protocol), 47-48**

- headers, 53-54
- port numbers, 51
- segments, 58

unicasts

- frame forwarding, 193
- IPv6 addressing, 114

unreliable data transmission, UDP, 47-48**upstream routing, 246****user exec access, 146****UTC (Universal Time Coordinated), NTP and, 335****V****verifying**

- ACL configurations, 308-309
- DHCP operation, 340-341
- host configurations on default gateways, 362
- NAT configurations, 330
- NTP configurations, 336-337
- OSPF
 - OSPFv2 configuration, 260
 - OSPFv3 configuration, 261-263
- PAT configurations, 330
- trunking configuration/operation, 357-360
- VLAN, 353-356

Version field (IPv6 headers), 118**VLAN (Virtual Local Area Networks), 204-205, 211, 353**

- benefits of, 206-207
- creating, 209
- defining, 206
- deleting, 209
- implementing, 208-210
- inter-VLAN routing, 242
 - L3 switching, 245-246
 - router on a stick, 243-244
 - upstream routing, 246
- memberships
 - dynamic VLAN membership, 211
 - static VLAN membership, 210
 - VMPS, 211
- naming, 209
- ports, verifying assignments, 354-356
- switches, preconfigured VLAN, 210
- troubleshooting, 353-356
- trunking, 206-207, 213
 - 802.1Q trunk encapsulation standard, 214
 - switch configuration, 215-216
 - verifying configuration/operation, 357-360

VLSM (Variable-Length Subnet Masking), 96-98, 239

VMPS (VLAN Membership Policy Server), 211

VoIP (Voice over IP), time stamps, 335

vty (Virtual Teletype) lines, ACL, 303-304

W - X - Y - Z

WAN (Wide-Area Networks)

cell-switching, defining, 32

circuit-switching, defining, 31

defining, 31

Internet as, 32

leased line connections, defining, 31

packet-switching, defining, 32

wireless networks, 32

well-meaning admin errors, NAT and, 323

Wi-Fi, defining, 32

wildcard masks, 255-256, 298-299

wireless networks, 32