



CCNP ROUTE Exam Preparation



CCNP ROUTE 642-902

Official Certification Guide

- ✓ Master the **CCNP® Route 642-902** exam with this official study guide
- ✓ Assess your knowledge with **chapter-opening quizzes**
- ✓ Review key concepts with **Exam Preparation Tasks**
- ✓ Practice with **realistic exam questions** on the CD-ROM

ciscopress.com

Wendell Odom, CCIE® No. 1624

FREE SAMPLE CHAPTER



SHARE WITH OTHERS

CCNP ROUTE 642-902

Official Certification Guide

Wendell Odom, CCIE No. 1624

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

CCNP ROUTE 642-902 Official Certification Guide

Wendell Odom

Copyright© 2010 Pearson Education, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

Ninth Printing: May 2014

Odom, Wendell.

CCNP Route 642-902 official certification guide / Wendell Odom.

p. cm.

ISBN 978-1-58720-253-7 (hardback w/cd)

1. Routers (Computer networks)--Examinations--Study guides. 2.

Routing protocols (Computer network protocols)--Examinations--Study

guides. 3. Internetworking (Telecommunication)--Examinations--Study

guides. 4. Telecommunications

engineers--Certification--Examinations--Study guides. I. Title.

TK5105.543.O36 2010

004.62--dc22

2009049908

ISBN-13: 978-1-58720-253-7

ISBN-10: 1-58720-253-0

Warning and Disclaimer

This book is designed to provide information about the Cisco ROUTE exam (642-902). 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 is provided on an "as is" basis. The authors, Cisco Press, and Cisco Systems, Inc. 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 discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., 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.

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: U.S. Corporate and Government Sales 1-800-382-3419 corpsales@pearsontechgroup.com

For sales outside the United States please contact: International Sales international@pearsoned.com

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Executive Editor: Brett Bartow

Managing Editor: Patrick Kanouse

Development Editor: Dayna Isley

Project Editor: Mandie Frank

Book Designer: Louisa Adair

Composition: Mark Shirar

Indexer: Ken Johnson

Business Operation Manager, Cisco Press: Anand Sundaram

Manager Global Certification: Erik Ullanderson

Technical Editors: Michelle Plumb, Jerold Swan, Rick Graziani

Copy Editor: Apostrophe Editing Services

Proofreader: Barbara Hacha

Editorial Assistant: Vanessa Evans



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Author

Wendell Odom, CCIE No. 1624, has been in the networking industry since 1981. He has worked as a network engineer, consultant, systems engineer, instructor, and course developer; he currently works writing and creating certification tools. He is author of all the previous editions of the *CCNA Exam Certification Guide* series from Cisco Press, as well as the *CCNP ROUTE 642-902 Official Certification Guide*, the *CCIE Routing and Switching Official Exam Certification Guide*, *Computer Networking First Step*, the *CCNA Video Mentor*, and he is the primary networking consultant for the *CCNA 640-802 Network Simulator*. He maintains lists of Cisco certification tools, suggestions, and links to his Cisco certification blog at www.certskills.com.

About the Technical Reviewers

Michelle Plumb is a full-time Cisco certified instructor for Skillssoft. Michelle has more than 19 years experience in the field as an IT professional and telephony specialist. She maintains a high level of Cisco and Microsoft certifications. Michelle has been a technical reviewer for numerous books related to the Cisco CCNP and CCVP course material track. Michelle currently lives in Scottsdale, Arizona, with her husband and two dogs.

Jerold Swan, CCIE No. 17783, CCSP, works as a senior network engineer for the Southern Ute Indian Tribe Growth Fund in southwest Colorado. Prior to that he was a Cisco instructor for Global Knowledge. He has also worked in IT in the service provider and higher education sectors. His areas of interest include routing protocols, security, and network monitoring. He is a graduate of Stanford University. His other interests include trail running, mountain biking, and volunteer search and rescue.

Rick Graziani teaches computer science and computer networking courses at Cabrillo College in Aptos, California. Rick has worked and taught in the computer networking and information technology field for almost 30 years. Prior to teaching Rick worked in IT for various companies including Santa Cruz Operation, Tandem Computers, and Lockheed Missiles and Space Corporation. He holds an M.A. degree in computer science and systems theory from California State University Monterey Bay. Rick also does consulting work for Cisco and other companies. When Rick is not working he is most likely surfing. Rick is an avid surfer who enjoys surfing at his favorite Santa Cruz breaks.

Dedications

For Jeffrey Lanier Odom. My favorite brother. Gentle soul. Lover of stupid jokes (“baby bigger,” “tankety-tankety-tank,” “supplies”...) Nice guy. Good friend. Miss you, bro.
10/7/1959—6/15/2009.

Acknowledgments

As usual, Brett Bartow, executive editor, deserves thanks for allowing me to be involved with this book. Brett continually keeps an eye on the horizon for the right projects for me, essentially completing a run of books from the basics, to CCENT, CCNA, now CCNP, and CCIE. My work life wouldn't be possible without Brett keeping me pointed in the right direction. Thanks, Brett!

Jay Swan and Michelle Plumb did a nice job for us with technical edits of the book. Jay was particularly helpful with both ends of the tech edit spectrum, noticing specific and easy-to-overlook errors, while keeping an eye out for the big picture of how the text in one section impacted other sections. Michelle's diligent work helped us uncover several specific issues and make this a better book. Thanks to you both for helping make this book much better!

Rick Graziani deserves thanks with this book for several reasons. First, Rick wrote the questions on the CD with this book, a task that can be laborious—but Rick did a great job and with a positive outlook. Additionally, Rick gave us an additional set of experienced and thoughtful technical editor eyes on the BGP chapters. And while he was working on the CD questions, Rick gladly went the extra mile to point out technical edits to the other book chapters as well. Rick's great attitude toward helping with the book was very impressive.

Dayna Isley worked as the development editor for this book. Dayna and I have worked very well together for a long time, and having such a trusted editor look over every word on this new book has helped quite a bit. Dayna's attention to detail helps keep me on the authoring straight-and-narrow, this time while navigating a sometimes fluid set of processes. Dayna, thanks for sifting through this process and making me look good on paper!

Patrick Kanouse, managing editor, led us through many new production tools (WriteRAP) and processes. Additionally, Patrick happily agreed to continue several additional production tasks at my request (translated: more work for him and his team), while allowing me to manage the entire illustration process for the first time on one of my books—none of which he had to do. Patrick, thanks for your great attitude and willingness to work with me on so many extras.

The folks on Patrick's production team probably had the biggest challenge with this book compared to my other books. Mandie Frank worked as project editor, guiding the book through the various back-end processes to complete the book. Mandie got to sift through all the changing processes, help figure out when we were doing which tasks, and keep us all on track. Thanks, Mandie! San Dee Phillips retired last year so she could work even more, coming back to do the copyedit work—thanks for jumping in again, San Dee! And for Mark Shirar, Ken Johnson, and Barbara Hacha, doing the composition, indexing, and proofreading, thanks so much for handling these details—I do see the difference with

having professionals working on every step of the book creation process, and I do appreciate the results.

Thanks to Rich Bennett, good friend and part-time do-everything guy for my books and other projects. Thanks for doing all the Illustrator drawings and editing them all so many times!

The old expression “my better half” is lived out every day here at the Odom house in the person of my wife Kris. Many thanks to Kris, who listens when I need to talk through something in the book, and lets me go hide in the basement for a few weeks to meet the latest writing deadline. Thanks, doll!

Finally and foremost, many thanks to Jesus Christ, for demonstrating your love, and for helping me and my family learn better each day how to not be a clanging symbol, but instead to show others your love.

Contents at a Glance

Foreword xxiv

Introduction xxv

Part I Perspectives on Network Planning

Chapter 1 Planning Tasks for the CCNP Exams 3

Part II EIGRP

Chapter 2 EIGRP Overview and Neighbor Relationships 19

Chapter 3 EIGRP Topology, Routes, and Convergence 57

Chapter 4 EIGRP Route Summarization and Filtering 97

Part III OSPF

Chapter 5 OSPF Overview and Neighbor Relationships 137

Chapter 6 OSPF Topology, Routes, and Convergence 175

Chapter 7 OSPF Route Summarization, Filtering, and Default Routing 221

Chapter 8 OSPF Virtual Links and Frame Relay Operations 257

Part IV Path Control

Chapter 9 Basic IGP Redistribution 289

Chapter 10 Advanced IGP Redistribution 329

Chapter 11 Policy-Based Routing and IP Service Level Agreement 363

Part V BGP

Chapter 12 Internet Connectivity and BGP 387

Chapter 13 External BGP 419

Chapter 14 Internal BGP and BGP Route Filtering 455

Chapter 15 BGP Path Control 491

Part VI IPv6

- Chapter 16 IP Version 6 Addressing 529
- Chapter 17 IPv6 Routing Protocols and Redistribution 569
- Chapter 18 IPv4 and IPv6 Coexistence 607

Part VII Branch Office Networking

- Chapter 19 Routing over Branch Internet Connections 647

Part VIII Final Preparation

- Chapter 20 Final Preparation 673

Part IX Appendixes

- Appendix A Answers to “Do I Know This Already?” Quizzes 681
- Appendix B Conversion Tables 701
- Appendix C Route Exam Updates 705
- Index 708

CD-Only Appendixes and Glossary

- Appendix D Memory Tables
- Appendix E Memory Tables Answer Key
- Appendix F Completed Planning Practice Tables
- Glossary

Contents

Foreword xxiv

Introduction xxv

Part I Perspectives on Network Planning

Chapter 1 Planning Tasks for the CCNP Exams 3

Perspectives on CCNP Exam Topics Related to Planning	3
CCNP Route Exam Topics That Do Not Require the CLI	4
Impressions on the Planning Exam Topics	5
Relating the Exam Topics to a Typical Network Engineer's Job	6
A Fictitious Company and Networking Staff	6
The Design Step	7
Implementation Planning Step	7
Verification Planning Step	9
Documenting the Results of the Implementation	10
Summary of the Role of Network Engineer	10
How to Prepare for the Planning Topics on the Exams	10
Planning Preparation: Design Review Table	12
Planning Preparation: Implementation Plan Peer Review Table	12
Create an Implementation Plan Table	13
Choose Commands for a Verification Plan Table	13
Background Information on Implementation and Verification Plans	13
No Single Plan Style	13
Typical Elements in an Implementation Plan	14
Focus for Implementation Plans for CCNP	15
Structured Implementation Planning Methodologies	15
Typical Verification Plan Components	16
Conclusions	16

Part II EIGRP

Chapter 2 EIGRP Overview and Neighbor Relationships 19

“Do I Know This Already?” Quiz	20
Foundation Topics	23
EIGRP CCNA Review	23
Configuration Review	23
Verification Review	25
Internals Review	29

EIGRP Neighborships	32
Manipulating EIGRP Hello and Hold Timers	32
Preventing Unwanted Neighbors Using Passive Interfaces	36
Controlling Neighborships Using EIGRP Authentication	39
Controlling Neighborships with Static Configuration	43
Configuration Settings That Could Prevent Neighbor Relationships	46
Neighborship over WANs	48
Neighborship on Frame Relay	49
Neighborship on MPLS VPN	50
Neighborship on Metro Ethernet	51
Exam Preparation Tasks	52
Planning Practice	52
Design Review Table	52
Implementation Plan Peer Review Table	52
Create an Implementation Plan Table	53
Choose Commands for a Verification Plan Table	53
Review All the Key Topics	55
Complete the Tables and Lists from Memory	55
Define Key Terms	55

Chapter 3 EIGRP Topology, Routes, and Convergence 57

“Do I Know This Already?” Quiz	57
Foundation Topics	60
Building the EIGRP Topology Table	60
Seeding the EIGRP Topology Table	60
The Content of EIGRP Update Message	61
The EIGRP Update Process	64
WAN Issues for EIGRP Topology Exchange	65
Building the IP Routing Table	69
Calculating the Metrics: Feasible Distance and Reported Distance	69
EIGRP Metric Tuning	72
Optimizing EIGRP Convergence	78
Fast Convergence to Feasible Successors	78
Converging by Going Active	83
Unequal Metric Route Load Sharing	88
Exam Preparation Tasks	92
Planning Practice	92
Design Review Table	92

	Implementation Plan Peer Review Table	92
	Create an Implementation Plan Table	93
	Choose Commands for a Verification Plan Table	94
	Review all the Key Topics	94
	Complete the Tables and Lists from Memory	95
	Define Key Terms	95
Chapter 4	EIGRP Route Summarization and Filtering	97
	“Do I Know This Already?” Quiz	97
	Foundation Topics	101
	Route Filtering	101
	Filtering by Referencing ACLs	102
	Filtering by Referencing IP Prefix Lists	105
	Filtering by Using Route Maps	110
	Route Summarization	114
	Route Summarization Design	114
	Configuring EIGRP Route Summarization	120
	Auto-summary	124
	Default Routes	126
	Default Routing to the Internet Router	126
	Default Routing Configuration with EIGRP	127
	Exam Preparation Tasks	132
	Planning Practice	132
	Design Review Table	132
	Implementation Plan Peer Review Table	132
	Create an Implementation Plan Table	133
	Choose Commands for a Verification Plan Table	134
	Review all the Key Topics	134
	Complete the Tables and Lists from Memory	135
	Define Key Terms	135
Part III	OSPF	
Chapter 5	OSPF Overview and Neighbor Relationships	137
	“Do I Know This Already?” Quiz	137
	Foundation Topics	140
	OSPF Review	140
	OSPF Link State Concepts	140
	OSPF Configuration Review	144

OSPF Verification Review	146
OSPF Feature Summary	149
OSPF Neighbors and Adjacencies on LANs	149
Enabling OSPF Neighbor Discovery on LANs	150
Settings That Must Match for OSPF Neighborship	152
OSPF Neighbors and Adjacencies on WANs	162
OSPF Network Types	162
OSPF Neighborship over Point-to-Point Links	163
Neighborship over Frame Relay Point-to-Point Subinterfaces	166
Neighborship on MPLS VPN	166
Neighborship on Metro Ethernet	167
Exam Preparation Tasks	170
Planning Practice	170
Design Review Table	170
Implementation Plan Peer Review Table	170
Create an Implementation Plan Table	171
Choose Commands for a Verification Plan Table	172
Review All the Key Topics	173
Complete the Tables and Lists from Memory	173
Define Key Terms	173

Chapter 6 OSPF Topology, Routes, and Convergence 175

“Do I Know This Already?” Quiz	175
Foundation Topics	179
LSAs and the OSPF Link State Database	179
LSA Type 1: Router LSA	180
LSA Type 2: Network LSA	186
LSA Type 3: Summary LSA	191
Limiting the Number of LSAs	195
Summary of Internal LSA Types	195
The Database Exchange Process	196
OSPF Message and Neighbor State Reference	196
Exchange Without a Designated Router	197
Exchange with a Designated Router	200
Flooding Throughout the Area	203
Periodic Flooding	204
Choosing the Best OSPF Routes	204
OSPF Metric Calculation for Internal OSPF Routes	205

	Metric and SPF Calculations	211
	Metric Tuning	212
	Exam Preparation Tasks	215
	Planning Practice	215
	Design Review Table	215
	Implementation Plan Peer Review Table	215
	Create an Implementation Plan Table	216
	Choose Commands for a Verification Plan Table	216
	Review All the Key Topics	218
	Complete the Tables and Lists from Memory	218
	Define Key Terms	218
Chapter 7	OSPF Route Summarization, Filtering, and Default Routing	221
	“Do I Know This Already?” Quiz	221
	Foundation Topics	225
	Route Filtering	225
	Type 3 LSA Filtering	226
	Filtering OSPF Routes Added to the Routing Table	230
	Route Summarization	231
	Manual Summarization at ABRs	232
	Manual Summarization at ASBRs	235
	Default Routes and Stub Areas	236
	Domain-wide Defaults Using the default-information originate Command	237
	Stubby Areas	239
	Exam Preparation Tasks	251
	Planning Practice	251
	Design Review Table	251
	Implementation Plan Peer Review Table	251
	Create an Implementation Plan Table	252
	Choose Commands for a Verification Plan Table	253
	Review All the Key Topics	253
	Complete the Tables and Lists from Memory	254
	Define Key Terms	254
Chapter 8	OSPF Virtual Links and Frame Relay Operations	257
	“Do I Know This Already?” Quiz	257
	Foundation Topics	260
	Virtual Links	260

Understanding OSPF Virtual Link Concepts	260
Configuring OSPF Virtual Links with No Authentication	262
Verifying the OSPF Virtual Link	264
Configuring Virtual Link Authentication	265
OSPF over Multipoint Frame Relay	267
IP Subnetting Design over Frame Relay	267
OSPF Challenges When Using Multipoint	270
Configuring and Verifying OSPF Operations on Frame Relay	274
Exam Preparation Tasks	283
Planning Practice	283
Design Review Table	283
Implementation Plan Peer Review Table	283
Create an Implementation Plan Table	284
Choosing Commands for a Verification Plan Table	285
Review All the Key Topics	285
Complete the Tables and Lists from Memory	286
Define Key Terms	286

Part IV Path Control

Chapter 9 Basic IGP Redistribution 289

“Do I Know This Already?” Quiz	289
Foundation Topics	292
Route Redistribution Basics	292
The Need for Route Redistribution	292
Redistribution Concepts and Processes	294
Redistribution into EIGRP	297
EIGRP redistribute Command Reference	297
Baseline Configuration for EIGRP Redistribution Examples	298
Configuring EIGRP Redistribution with Default Metric Components	300
Verifying EIGRP Redistribution	302
Redistribution into OSPF	305
OSPF redistribute Command Reference	305
Configuring OSPF Redistribution with Minimal Parameters	306
Setting OSPF Metrics on Redistributed Routes	310
LSAs and Metrics for External Type 2 Routes	311
Redistributing into OSPF as E1 Routes	318
A Brief Comparison of E1 and E2 Routes	319
External Routes in NSSA Areas	320

Exam Preparation Tasks	324
Planning Practice	324
Design Review Table	324
Implementation Plan Peer Review Table	325
Create an Implementation Plan Table	326
Choosing Commands for a Verification Plan Table	326
Review all the Key Topics	327
Complete the Tables and Lists from Memory	327
Define Key Terms	327

Chapter 10 Advanced IGP Redistribution 329

“Do I Know This Already?” Quiz	329
Foundation Topics	332
Redistribution with Route Maps and Distribute Lists	332
Overview of Using route-maps with Redistribution	332
Filtering Redistributed Routes with Route Maps	334
Setting Metrics when Redistributing	339
Setting the External Route Type	343
Redistribution Filtering with the distribute-list Command	343
Issues with Multiple Redistribution Points	344
Preventing Routing Domain Loops with Higher Metrics	345
Preventing Routing Domain Loops with Administrative Distance	346
Domain Loop Problems with More than Two Routing Domains	349
Exam Preparation Tasks	358
Planning Practice	358
Design Review Table	358
Implementation Plan Peer Review Table	358
Create an Implementation Plan Table	359
Choose Commands for a Verification Plan Table	360
Review all the Key Topics	361
Complete the Tables and Lists from Memory	361
Define Key Terms	361

Chapter 11 Policy-Based Routing and IP Service Level Agreement 363

“Do I Know This Already?” Quiz	363
Foundation Topics	366
Policy-Based Routing	366
Matching the Packet and Setting the Route	367
PBR Configuration Example	368

How the default Keyword Impacts PBR Logic Ordering	370
Additional PBR Functions	371
IP Service-Level Agreement	372
Understanding IP SLA Concepts	373
Configuring and Verifying IP SLA	374
Tracking SLA Operations to Influence Routing	378
Exam Preparation Tasks	382
Planning Practice	382
Design Review Table	382
Implementation Plan Peer Review Table	382
Create an Implementation Plan Table	382
Choose Commands for a Verification Plan Table	383
Review all the Key Topics	384
Complete the Tables and Lists from Memory	385
Definitions of Key Terms	385

Part V BGP

Chapter 12 Internet Connectivity and BGP 387

“Do I Know This Already?” Quiz	388
Foundation Topics	390
The Basics of Internet Routing and Addressing	390
Public IP Address Assignment	391
Internet Route Aggregation	392
The Impact of NAT/PAT	393
Private IPv4 Addresses and Other Special Addresses	394
Introduction to BGP	396
BGP Basics	396
BGP ASNs and the AS_SEQ Path Attribute	397
Internal and External BGP	399
Public and Private ASNs	400
Outbound Routing Toward the Internet	402
Comparing BGP and Default Routing for Enterprises	402
Single Homed	404
Dual Homed	405
Single Multihomed	411
Dual Multihomed	412
Exam Preparation Tasks	414
Planning Practice	414

Design Review Table	414
Implementation Plan Peer Review Table	414
Create an Implementation Plan Table	415
Review all the Key Topics	415
Complete the Tables and Lists from Memory	416
Define Key Terms	416

Chapter 13 External BGP 419

“Do I Know This Already?” Quiz	419
Foundation Topics	423
External BGP for Enterprises	423
eBGP Neighbor Configuration	423
BGP Internals and Verifying eBGP Neighbors	430
Verifying the BGP Table	436
The BGP Update Message	436
Examining the BGP Table	438
Viewing Subsets of the BGP Table	440
Injecting Routes into BGP for Advertisement to the ISPs	443
Injecting Routes Using the network Command	443
The Effect of auto-summary on the BGP network Command	445
Injecting Routes Using Redistribution	446
Exam Preparation Tasks	449
Planning Practice	449
Design Review Table	449
Implementation Plan Peer Review Table	449
Create an Implementation Plan Table	450
Choose Commands for a Verification Plan Table	451
Review all the Key Topics	452
Complete the Tables and Lists from Memory	452
Define Key Terms	452

Chapter 14 Internal BGP and BGP Route Filtering 455

“Do I Know This Already?” Quiz	455
Foundation Topics	459
Internal BGP Between Internet-Connected Routers	459
Establishing the Need for iBGP with Two Internet-Connected Routers	459
Configuring iBGP	460
Verifying iBGP	463
Examining iBGP BGP Table Entries	464

Understanding Next-Hop Reachability Issues with iBGP	468
Avoiding Routing Loops when Forwarding Toward the Internet	471
Using an iBGP Mesh	472
IGP Redistribution and BGP Synchronization	475
Route Filtering and Clearing BGP Peers	476
BGP Filtering Overview	476
Inbound and Outbound BGP Filtering on Prefix/Length	478
Clearing BGP Neighbors	481
Displaying the Results of BGP Filtering	483
Exam Preparation Tasks	486
Planning Practice	486
Design Review Table	486
Implementation Plan Peer Review Table	487
Create an Implementation Plan Table	487
Choosing Commands for a Verification Plan Table	488
Review all the Key Topics	488
Complete the Tables and Lists from Memory	489
Definitions of Key Terms	489
Chapter 15 BGP Path Control	491
“Do I Know This Already?” Quiz	491
Foundation Topics	494
BGP Path Attributes and Best Path Algorithm	494
BGP Path Attributes	494
Overview of the BGP Best Path Algorithm	495
Perspectives on the Core 8 Best Path Steps	498
Memorization Tips for BGP Best Path	499
Influencing an Enterprise’s Outbound Routes	500
Influencing BGP Weight	500
Setting the Local Preference	507
IP Routes Based on BGP Best Paths	513
Increasing the Length of the AS_Path Using AS_Path Prepend	517
Influencing an Enterprise’s Inbound Routes with MED	519
MED Concepts	519
MED Configuration	521
Exam Preparation Tasks	523
Planning Practice	523
Design Review Table	523

Implementation Plan Peer Review Table	523
Create an Implementation Plan Table	524
Choose Commands for a Verification Plan Table	525
Review all the Key Topics	526
Complete the Tables and Lists from Memory	526
Define Key Terms	526

Part VI IPv6

Chapter 16 IP Version 6 Addressing 529

“Do I Know This Already?” Quiz	529
Foundation Topics	532
Global Unicast Addressing, Routing, and Subnetting	533
Global Route Aggregation for Efficient Routing	534
Conventions for Representing IPv6 Addresses	536
Conventions for Writing IPv6 Prefixes	537
Global Unicast Prefix Assignment Example	539
Subnetting Global Unicast IPv6 Addresses Inside an Enterprise	541
Prefix Terminology	543
Assigning IPv6 Global Unicast Addresses	544
Stateful DHCP for IPv6	545
Stateless Autoconfiguration	545
Static IPv6 Address Configuration	549
Survey of IPv6 Addressing	549
Overview of IPv6 Addressing	550
Unicast IPv6 Addresses	550
Multicast and Other Special IPv6 Addresses	553
Layer 2 Addressing Mapping and Duplicate Address Detection	554
Configuring IPv6 Addresses on Cisco Routers	556
Configuring Static IPv6 Addresses on Routers	557
Multicast Groups Joined by IPv6 Router Interfaces	559
Connected Routes and Neighbors	560
The IPv6 Neighbor Table	561
Stateless Autoconfiguration	561
Exam Preparation Tasks	563
Planning Practice	563
Design Review Table	563
Implementation Plan Peer Review Table	563
Create an Implementation Plan Table	564

Choose Commands for a Verification Plan Table 564

Review all the Key Topics 565

Complete the Tables and Lists from Memory 566

Define Key Terms 566

Chapter 17 IPv6 Routing Protocols and Redistribution 569

“Do I Know This Already?” Quiz 569

Foundation Topics 573

RIP Next Generation (RIPng) 573

RIPng—Theory and Comparisons to RIP-2 574

Configuring RIPng 575

Verifying RIPng 578

EIGRP for IPv6 581

EIGRP for IPv4 and IPv6—Theory and Comparisons 581

Configuring EIGRP for IPv6 582

Verifying EIGRP for IPv6 584

OSPF Version 3 588

Comparing OSPFv2 and OSPFv3 588

Configuring OSPFv3 590

Verifying OSPFv3 592

IPv6 IGP Redistribution 595

Redistributing without Route Maps 596

Redistributing with Route Maps 598

Static IPv6 Routes 599

Exam Preparation Tasks 602

Planning Practice 602

Implementation Plan Peer Review Table 602

Create an Implementation Plan Table 602

Choose Commands for a Verification Plan Table 603

Review all the Key Topics 604

Complete the Tables and Lists from Memory 604

Define Key Terms 604

Chapter 18 IPv4 and IPv6 Coexistence 607

“Do I Know This Already?” Quiz 607

Foundation Topics 611

IPv4 and IPv6 Migration and Coexistence Concepts 611

IPv4/IPv6 Dual Stacks 611

Tunneling 612

NAT Protocol Translation	617
Static Point-to-Point IPv6 Tunnels	619
Manually Configured Tunnels	620
GRE Tunnels	624
Point-to-Point IPv6 Tunnel Summary	625
Dynamic Multipoint IPv6 Tunnels	626
Automatic 6to4 Tunnels	627
IPv6 ISATAP Tunnels	634
Multipoint IPv6 Tunnel Summary	639
Exam Preparation Tasks	641
Planning Practice	641
Design Review Table	641
Implementation Plan Peer Review Table	642
Create an Implementation Plan Table	642
Choose Commands for a Verification Plan Table	643
Review all the Key Topics	644
Complete the Tables and Lists from Memory	644
Define Key Terms	644

Part VII Branch Office Networking

Chapter 19 Routing over Branch Internet Connections 647

“Do I Know This Already?” Quiz	647
Foundation Topics	650
Branch Office Broadband Internet Access	650
Broadband Internet Access Basics	650
Branch Router as DHCP Server and Client	652
Branch Office Security	653
Using IPsec Tunnels	654
Branch Routing for the Small Branch	656
Routing in Medium and Large Branches	657
Branch Router Configuration for Broadband Access	659
Understanding DSL Concepts	659
Configuring DSL	661
Configuring NAT	663
Configuring DHCP Server	664
VPN Configuration	664
Configuring an IPsec VPN	665

	Configuring GRE Tunnels	666
	Summary–Branch Routing from PC1 to Enterprise Server S1	667
	Exam Preparation Tasks	670
	Planning Practice	670
	Review all the Key Topics	671
	Define Key Terms	671
Part VIII	Final Preparation	
Chapter 20	Final Preparation	673
	Tools for Final Preparation	673
	Exam Engine and Questions on the CD	673
	Install the Software from the CD	674
	Activate and Download the Practice Exam	674
	Activating Other Exams	675
	The Cisco Learning Network	675
	Memory Tables	675
	Chapter-Ending Review Tools	676
	Suggested Plan for Final Review/Study	676
	Step 1: Review Key Topics and DIKTA Questions	677
	Step 3: Hands-On Practice	677
	Step 6: Subnetting Practice	677
	Step 7: Use the Exam Engine	678
	Summary	679
Part IX	Appendixes	
Appendix A	Answers to “Do I Know This Already?” Quizzes	681
Appendix B	Conversion Tables	701
Appendix C	Route Exam Updates	705
	Index	708
	CD-Only Appendixes and Glossary	
Appendix D	Memory Tables	
Appendix E	Memory Tables Answer Key	
Appendix F	Completed Planning Practice Tables	
	Glossary	

Icons Used in This Book



Router

Workgroup
Switch

Multilayer Switch



Firewall



Server



Network Cloud



Serial Cable




Line: Ethernet



VPN Tunnel



PC

Standing
Man

Scroll

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ([{ }]) indicate a required choice within an optional element.

Foreword

CCNP ROUTE 642-902 Official Certification Guide is an excellent self-study resource for the CCNP ROUTE exam. Passing this exam is a crucial step to attaining the valued CCNP Routing and Switching certification.

Gaining certification in Cisco technology is key to the continuing educational development of today's networking professional. Through certification programs, Cisco validates the skills and expertise required to effectively manage the modern enterprise network.

Cisco Press Certification Guides and preparation materials offer exceptional—and flexible—access to the knowledge and information required to stay current in your field of expertise or to gain new skills. Whether used as a supplement to more traditional training or as a primary source of learning, these materials offer users the information and knowledge validation required to gain new understanding and proficiencies.

Developed in conjunction with the Cisco certifications and training team, Cisco Press books are the only self-study books authorized by Cisco and offer students a series of exam practice tools and resource materials to help ensure that learners fully grasp the concepts and information presented.

Additional authorized Cisco instructor-led courses, e-learning, labs, and simulations are available exclusively from Cisco Learning Solutions Partners worldwide. To learn more, visit <http://www.cisco.com/go/training>.

I hope that you find these materials to be an enriching and useful part of your exam preparation.

Erik Ullanderson
Manager, Global Certifications
Learning@Cisco
January 2010

Introduction

This book focuses on one major goal: to help you prepare to pass the ROUTE exam (642-902). To help you prepare, this book achieves other useful goals as well: It explains a wide range of networking topics, shows how to configure those features on Cisco routers, and explains how to determine if the feature is working. As a result, you also can use this book as a general reference for IP routing and IP routing protocols. However, the motivation for this book, and the reason it sits within the Cisco Press Certification Guide series, is that its primary goal is to help you pass the ROUTE exam.

The rest of this introduction focuses on two topics: the ROUTE exam and a description of this book.

The CCNP ROUTE Exam

Cisco announced the ROUTE (642-902) exam in January 2010. The term ROUTE does not act as an acronym; instead, the name describes the content of the exam, which focuses on IP routing. Generally, the exam includes detailed coverage of the EIGRP, OSPF, and BGP IP routing protocols, IPv6, and a few other smaller topics related to IP routing.

Cisco first announced its initial Professional level certifications in 1998 with the CCNP Routing and Switching certification. CCNP Routing and Switching certification from its inception has included the same kinds of IP routing topics found in today's ROUTE exam, but the exam names changed over the years. The exam names have tracked the names of the associated Cisco authorized courses for the same topics: Advanced Cisco Router Configuration (ACRC) in the early days, Building Scalable Cisco Internetworks (BSCI) for much of the last 10 years, and now ROUTE, because the newly revised (in 2010) Cisco authorized course also goes by the name ROUTE.

Like its ancestors, the ROUTE exam is a part of the certification requirements for several Cisco certifications, as follows:

- Cisco Certified Networking Professional (CCNP)
- Cisco Certified Internetworking Professional (CCIP)
- Cisco Certified Design Professional (CCDP)

Each of these certifications emphasizes different perspectives on some similar topics. CCNP focuses on the skills needed by a network engineer working for an Enterprise—that is, a company that deploys networking gear for its own purposes. CCIP focuses on the skills required by network engineers deploying gear at a service provider, with the service provider then offering network services to customers. Finally, CCDP focuses more on design—but good design requires solid knowledge of the technology and configuration. So, although this book frequently refers to the most popular certification of these three—CCNP—the ROUTE exam does apply to several certifications.

Contents of the ROUTE Exam

Every student who ever takes an exam wants to know what's on the exam. As with all their exams, Cisco publishes a set of exam topics. These exam topics give general guidance as to what's on the exam.

You can find the exam topics at the Cisco website. The most memorable way to navigate is to go to www.cisco.com/go/ccnp, and look for the ROUTE exam. Also, you can go to the Cisco Learning Network website (www.cisco.com/go/learnnet)—a less memorable URL, but a great Cisco certification site. The Cisco Learning Network site hosts exam information, learning tools, and forums in which you can communicate with others and learn more about this and other Cisco exams.

Table I-1 lists the ROUTE exam topics, with a reference to the part of the book that covers the topic.

Table I-1 ROUTE Exam Topics

Book Part	Exam Topic
Implement an EIGRP based solution, given a network design and a set of requirements	
II	Determine network resources needed for implementing EIGRP on a network
II	Create an EIGRP implementation plan
II	Create an EIGRP verification plan
II	Configure EIGRP routing
II	Verify EIGRP solution was implemented properly using show and debug commands
II	Document results of EIGRP implementation and verification
Implement a multi-area OSPF Network, given a network design and a set of requirements	
III	Determine network resources needed for implementing OSPF on a network
III	Create an OSPF implementation plan
III	Create an OSPF verification plan
III	Configure OSPF routing
III	Verify OSPF solution was implemented properly using show and debug commands
III	Document results of OSPF implementation and verification plan
Implement an eBGP based solution, given a network design and a set of requirements	
V	Determine network resources needed for implementing eBGP on a network
V	Create an eBGP implementation plan
V	Create an eBGP verification plan
V	Configure eBGP routing
V	Verify eBGP solution was implemented properly using show and debug commands
V	Document results of eBGP implementation and verification plan

Table I-1 ROUTE Exam Topics

Book Part	Exam Topic
Implement an IPv6 based solution, given a network design and a set of requirements	
VI	Determine network resources needed for implementing IPv6 on a network
VI	Create an IPv6 implementation plan
VI	Create an IPv6 verification plan
VI	Configure IPv6 routing
VI	Configure IPv6 interoperation with IPv4
VI	Verify IPv6 solution was implemented properly using show and debug commands
VI	Document results of IPv6 implementation and verification plan
Implement an IPv4 or IPv6 based redistribution solution, given a network design and a set of requirements	
IV, VI	Create a redistribution implementation plan based upon the results of the redistribution analysis.
IV, VI	Create a redistribution verification plan
IV, VI	Configure a redistribution solution
IV, VI	Verify that a redistribution was implemented
IV, VI	Document results of a redistribution implementation and verification plan
IV, VI	Identify the differences between implementing an IPv4 and IPv6 redistribution solution
Implement Layer 3 Path Control Solution	
IV	Create a Layer 3 path control implementation plan based upon the results of the redistribution analysis.
IV	Create a Layer 3 path control verification plan
IV	Configure Layer 3 path control
IV	Verify that a Layer 3 path control was implemented
IV	Document results of a Layer 3 path control implementation and verification plan
Implement basic teleworker and branch services	
VII	Describe broadband technologies
VII	Configure basic broadband connections
VII	Describe basic VPN technologies
VII	Configure GRE
VII	Describe branch access technologies

How to Take the ROUTE Exam

As of the publication of this book, Cisco exclusively uses testing vendor Pearson Vue (www.vue.com) for delivery of all Cisco career certification exams. To register, go to www.vue.com, establish a login, and register for the 642-902 ROUTE exam. You also need to choose a testing center near to your home.

Who Should Take This Exam and Read This Book?

This book has one primary audience, with several secondary audiences. First, this book is intended for anyone wanting to prepare for the ROUTE 642-902 exam. The audience includes self-study readers—people who pass the test by studying 100 percent on their own. It includes Cisco Networking Academy students taking the CCNP curriculum, who use this book to round out their preparation as they get close to the end of the Academy curriculum.

The broader question about the audience may well be why you should take the ROUTE exam. First, the exam is required for the aforementioned CCNP, CCIP, and CCDP certifications from Cisco. These certifications exist at the midpoint of the Cisco certification hierarchy. These certifications have broader and deeper technology requirements as compared to the Cisco Certified Entry Network Technician (CCENT) and Cisco Certified Network Associate (CCNA) certifications.

The real question then about audience for this book—at least the intended audience—is whether you have motivation to get one of these Professional-level Cisco certifications. CCNP in particular happens to be a popular, well-respected certification. CCIP, although less popular in numbers, focuses on topics more important to service providers, so it gives you a good way to distinguish yourself from others looking for jobs at SP companies. CCDP has been a solid certification for a long time, particularly for engineers who spend a lot of time designing networks with customers, rather than troubleshooting.

Format of the CCNP ROUTE Exam

The ROUTE exam follows the same general format as the other Cisco exams. When you get to the testing center and check in, the proctor will give you some general instructions and then take you into a quiet room with a PC. When you're at the PC, you have a few things to do before the timer starts on your exam—for instance, you can take a sample quiz, just to get accustomed to the PC and to the testing engine. Anyone who has user-level skills in getting around a PC should have no problems with the testing environment.

When you start the exam, you will be asked a series of questions. You answer the question and then move on to the next question. *The exam engine does not let you go back and change your answer.* Yes, that's true—when you move on to the next question, that's it for the earlier question.

The exam questions can be in one of the following formats:

- Multiple choice (MC)
- Testlet
- Drag-and-drop (DND)
- Simulated lab (Sim)
- Simlet

The first three types of questions are relatively common in many testing environments. The multiple choice format simply requires that you point-and-click on a circle beside the correct answer(s). Cisco traditionally tells you how many answers you need to choose, and the testing software prevents you from choosing too many answers. Testlets are questions with one general scenario, with multiple MC questions about the overall scenario. Drag-and-drop questions require you to left-click and hold, move a button or icon to another area, and release the clicker to place the object somewhere else—typically into a list. So, for some questions, to get the question correct, you might need to put a list of five things into the proper order.

The last two types both use a network simulator to ask questions. Interestingly, the two types actually allow Cisco to assess two very different skills. First, Sim questions generally describe a problem, and your task is to configure one or more routers and switches to fix the problem. The exam then grades the question based on the configuration you changed or added. Interestingly, Sim questions are the only questions that Cisco (to date) has openly confirmed that partial credit is given.

The Simlet questions may well be the most difficult style of question on the exams. Simlet questions also use a network simulator, but instead of answering the question by changing the configuration, the question includes one or more MC questions. The questions require that you use the simulator to examine the current behavior of a network, interpreting the output of any **show** commands that you can remember to answer the question. Although Sim questions require you to troubleshoot problems related to a configuration, Simlets require you to both analyze working networks and networks with problems, correlating **show** command output with your knowledge of networking theory and configuration commands.

The Cisco Learning Network (learningnetwork.cisco.com) website has tools that let you experience the environment and see how each of these question types work. The environment should be the same as when you passed CCNA (a prerequisite for CCNP, CCIP, and CCDP).

CCNP ROUTE 642-902 Official Certification Guide

This section lists a general description of the contents of this book. The description includes an overview of each chapter, and a list of book features seen throughout the book.

Book Features and Exam Preparation Methods

This book uses several key methodologies to help you discover the exam topics on which you need more review, to help you fully understand and remember those details, and to help you prove to yourself that you have retained your knowledge of those topics. So, this book does not try to help you pass the exams only by memorization, but by truly learning and understanding the topics.

The book includes many features that provide different ways to study to be ready for the test. If you understand a topic when you read it, but do not study it any further, you probably will not be ready to pass the test with confidence. The book features included in this book give you tools that help you determine what you know, review what you know, better learn what you don't know, and be well prepared for the exam. These tools include

- **“Do I Know This Already?” Quizzes:** Each chapter begins with a quiz that helps you determine the amount of time you need to spend studying that chapter.
- **Foundation Topics:** These are the core sections of each chapter. They explain the protocols, concepts, and configuration for the topics in that chapter.
- **Exam Preparation Tasks:** The Exam Preparation Tasks section lists a series of study activities that should be done after reading the Foundation Topics section. Each chapter includes the activities that make the most sense for studying the topics in that chapter. The activities include
 - **Planning Tables:** The ROUTE exam topics includes some perspectives on how an engineer plans for various tasks. The idea is that the CCNP-level engineer in particular takes the design from another engineer, plans the implementation, and plans the verification steps—handing off the actual tasks to engineers working during change-window hours. Because the engineer plans the tasks, but may not be at the keyboard when implementing a feature, that engineer must master the configuration and verification commands so that the planned commands work for the engineer making the changes off-shift. The planning tables at the end of the chapter give you the chance to take the details in the Foundation Topics core of the chapter and think about them as if you were writing the planning documents.
 - **Key Topics Review:** The Key Topics icon is shown next to the most important items in the Foundation Topics section of the chapter. The Key Topics Review activity lists the Key Topics from the chapter, and page number. Although the contents of the entire chapter could be on the exam, you should definitely know the information listed in each key topic. Review these topics carefully.
 - **Memory Tables:** To help you exercise your memory and memorize some lists of facts, many of the more important lists and tables from the chapter are included in a document on the CD. This document lists only partial information, allowing you to complete the table or list. CD-only Appendix D holds the incomplete tables, and Appendix E includes the completed tables from which you can check your work.
 - **Definition of Key Terms:** Although the exams may be unlikely to ask a question such as “Define this term,” the ROUTE exam requires that you learn and know a lot of networking terminology. This section lists the most important terms from the chapter, asking you to write a short definition and compare your answer to the glossary at the end of the book.
- **CD-based practice exam:** The companion CD contains an exam engine (from Boson software, www.boson.com), which includes 100 unique multiple-choice questions. Chapter 20 gives two suggestions on how to use these questions: either as study questions, or to simulate the ROUTE exam.



- **Companion website:** The website <http://www.ciscopress.com/title/9781587202537> posts up-to-the-minute materials that further clarify complex exam topics. Check this site regularly for new and updated postings written by the author that provide further insight into the more troublesome topics on the exam.

Book Organization

This book contains 20 chapters, plus appendixes. The topics all focus in some way on IP routing and IP routing protocols, making the topics somewhat focused, but with deep coverage on those topics.

The book organizes the topics into seven major parts. Parts 1 and 7 include topics with less technical depth, and Parts 2 through 6 include the major technical topics in the book. The following list outlines the major part organization of this book:

Part I: “Perspectives on Network Planning”: This part includes a single chapter:

- **Chapter 1: “Planning Tasks for the CCNP Exams”:** This chapter discusses the CCNP ROUTE exam’s perspectives on the planning process, including network design, implementation plans, and verification plans.

Part II: “EIGRP”: This part starts with a CCNA-level EIGRP review and moves through EIGRP theory, configuration, authentication, route summarization, and more in the following chapters:

- **Chapter 2: “EIGRP Overview and Neighbor Relationships”:** This chapter reviews CCNA-level EIGRP topics and then closely examines the concepts, configuration, and verification of EIGRP neighbor relationships.
- **Chapter 3: “EIGRP Topology, Routes, and Convergence”:** This chapter examines the EIGRP topology database and the processes by which EIGRP processes this data to choose routes. It also examines the convergence process using feasible successors and with the Query process.
- **Chapter 4: “EIGRP Route Summarization and Filtering”:** This chapter discusses the theory behind route summarization and route filtering. It also shows how to configure and verify both features for EIGRP.

Part III: “OSPF”: Similar to Part II, this part starts with a CCNA-level OSPF review and moves through OSPF theory, configuration, authentication, metric tuning, default routing, route filtering, and route summarization, plus OSPF multiarea issues and different stubby area types, as follows:

- **Chapter 5: “OSPF Overview and Neighbor Relationships”:** This chapter reviews CCNA-level OSPF topics and then closely examines the concepts, configuration, and verification of OSPF neighbor relationships.
- **Chapter 6: “OSPF Topology, Routes, and Convergence”:** This chapter examines the OSPF topology database for routes internal to OSPF. The chapter also discusses how OSPF routers choose the best internal OSPF routes and how OSPF converges when a change occurs.

- **Chapter 7: “OSPF Route Summarization, Filtering, and Default Routing”:** This chapter discusses the design, configuration, and verification of OSPF route summarization and route filtering. It also discusses default routes and how to manage the size of the OSPF database and IP routing tables by using stubby areas.
- **Chapter 8: “OSPF Miscellany”:** This chapter discusses two additional OSPF topics: OSPF virtual links and OSPF issues when using NBMA networks (such as Frame Relay).

Part IV: “Path Control”: The term path control refers to a wide variety of topics related to IP routing and IP routing protocols. This part examines the path control topics not specifically included in the other parts of the book:

- **Chapter 9: “Basic IGP Redistribution”:** This chapter examines the concepts, configuration, and verification of IGP route redistribution. In particular, this chapter looks at the mechanics of redistribution without the use of route maps for any purpose.
- **Chapter 10: “Advanced IGP Redistribution”:** This chapter essentially continues Chapter 9, in this case focusing on the more complex configuration and issues. In particular, this chapter shows how to manipulate and filter routes at the redistribution function by using route maps, and how to avoid routing loops and inefficient routes when multiple redistribution points exist.
- **Chapter 11: “Policy Routing and IP Service Level Agreement”:** This chapter picks up two small path control topics that simply do not fit into any other broader chapter in this book: Policy Based Routing (PBR) and IP Service Level Agreement (IP SLA).

Part V: “BGP”: This part assumes no prior knowledge of BGP. It first examines BGP design issues, to give perspective on why BGP works differently than its IGP cousins OSPF and EIGRP. This part examines basic BGP concepts, configuration, and verification, including the path control functions of influencing both inbound and outbound BGP routes:

- **Chapter 12: “Internet Connectivity and BGP”:** This chapter introduces BGP. It begins with a review of Internet connectivity from a Layer 3 perspective. It then looks at the basics of how BGP works. It also examines some Internet access design issues, discussing the cases in which BGP can be helpful and the cases in which BGP has no practical use.
- **Chapter 13: “External BGP”:** This chapter examines the configuration and verification of BGP between an Enterprise and its ISP(s).
- **Chapter 14: “Internal BGP and BGP Route Filtering”:** This chapter examines the cases in which routers in the same ASN need to become BGP peers, creating an Internet BGP connection. It also discusses the need for BGP filtering and the mechanics of configuring BGP filtering.
- **Chapter 15: “BGP Path Control”:** This chapter discusses the concept of the BGP

Best Path Algorithm to choose the best BGP routes and how to influence those choices. In particular, this chapter shows the basic configuration for BGP weight, Local Preference, AS_Path length, and Multi-Exit Discriminator (MED).

Part VI: “IPv6”: This part assumes no prior knowledge of IPv6. The chapters in this part work through IPv6 addressing and IGP configuration (RIPng, EIGRP for IPv6, and OSPFv3). It also discusses route redistribution for IPv6 and IPv6/IPv4 coexistence mechanisms:

- **Chapter 16: “IP Version 6 Addressing”:** This chapter begins with an overview of IP Version 6 (IPv6). It then dives into IPv6 addressing concepts, plus the related protocols, including address assignment options and neighbor discovery. The chapter shows how to configure and verify IPv6 addresses on Cisco routers.
- **Chapter 17: “IPv6 Routing Protocols and Redistribution”:** This chapter introduces three IPv6 IGPs: RIP Next Generation (RIPng), EIGRP for IPv6, and OSPF Version 3 (OSPFv3). The chapter focuses on basic configuration and verification. It also discusses IPv6 redistribution in comparison with IPv4 IGP redistribution.
- **Chapter 18: “IPv4 and IPv6 Coexistence”:** This chapter discusses the many options to use during the potentially long migration from a purely IPv4 network to a future purely IPv6 network.

Part VII: “Branch Office Networking”: This short part includes one chapter that addresses a few small topics related to branch offices that connect to their Enterprise networks using the Internet:

- **Chapter 19: “Routing over Branch Internet Connections”:** Branch office routers can be configured to use the Internet as a WAN connection path back to the rest of an Enterprise network. This chapter takes a wide look at the surprisingly large number of networking functions that must occur on a branch router in such cases. It also gives examples of configurations for IPsec and GRE tunnels, DHCP server, NAT, and DSL.

Part VIII: “Final Preparation”: This short part includes one chapter as well. This chapter does not include any new technical topics:

- **Chapter 20: “Final Preparation”:** This chapter suggests some study strategies for your final preparation before the ROUTE exam.

In addition to the core chapters of the book, the book has several appendixes as well. Some appendixes exist in the printed book, whereas others exist in softcopy form on the CD included with the book.

Printed Appendixes

Appendixes printed in the book include

- **Appendix A, “Answers to the “Do I Know This Already?” Quizzes”:** Includes the answers to all the questions from Chapters 2 through 19.
- **Appendix B, “Conversion Tables”:** Lists a decimal-to-binary conversion table, decimal values 0 through 255, along with the binary equivalents. It also lists a hex-to-decimal conversion table as well.
- **Appendix C, “CCNP ROUTE Exam Updates: Version 1.0”:** Covers a variety of short topics that either clarify or expand upon topics covered earlier in the book. This appendix is updated from time to time, and posted at <http://www.ciscopress.com/title/9781587202537>, with the most recent version available at the time of printing included here as Appendix C. (The first page of the appendix includes instructions on how to check to see if a later version of Appendix C is available online.)

CD Appendixes

The appendixes included on the CD-ROM are

- **Appendix D, “Memory Tables”:** This appendix holds the key tables and lists from each chapter with some of the content removed. You can print this appendix, and as a memory exercise, complete the tables and lists. The goal is to help you memorize facts that can be useful on the exams.
- **Appendix E, “Memory Tables Answer Key”:** This appendix contains the answer key for the exercises in Appendix D.
- **Appendix F, “Completed Planning Practice Tables”:** The end of Chapters 2 through 19 list planning tables that you can complete to help learn the content more deeply. If you use these tables, refer to this appendix for the suggested answers.
- **Glossary:** The glossary contains definitions for all the terms listed in the “Define Key Terms” section at the conclusion of Chapters 2 through 19.

For More Information

If you have any comments about the book, you can submit those via the www.ciscopress.com. Just go to the website, select Contact Us, and type in your message.

Cisco might make changes that affect the ROUTE exam from time to time. You should always check www.cisco.com/go/ccnp for the latest details.

This page intentionally left blank



This chapter covers the following subjects:

LSAs and the OSPF Link State Database: This section examines LSA Types 1, 2, and 3, and how they allow OSPF routers to model the topology and choose the best routes for each known subnet.

The Database Exchange Process: This section details how neighboring routers use OSPF messages to exchange their LSAs.

Choosing the Best Internal OSPF Routes: This section examines how OSPF routers calculate the cost for each possible route to each subnet.

OSPF Topology, Routes, and Convergence

OSPF and EIGRP both use three major branches of logic, each of which populates a different table: the neighbor table, the topology table, and the IP routing table. This chapter examines topics related to the OSPF topology table—the contents, and the processes by which routers exchange this information—and how OSPF routers choose the best routes in the topology table to be added to the IP routing table.

In particular, this chapter begins by looking at the building blocks of OSPF topology, namely the OSPF link state advertisement (LSA). Following that, the chapter examines the process by which OSPF routers exchange LSAs with each other. Finally, the last major section of the chapter discusses how OSPF chooses the best route among many when running the Shortest Path First (SPF) algorithm.

Note that this chapter focuses on OSPF Version 2, the long-available version of OSPF that supports IPv4 routes. Chapter 17, “IPv6 Routing Protocols and Redistribution,” discusses OSPF Version 3, which applies to IPv6.

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz allows you to assess if you should read the entire chapter. If you miss no more than one of these nine self-assessment questions, you might want to move ahead to the “Exam Preparation Tasks.” Table 6-1 lists the major headings in this chapter and the “Do I Know This Already?” quiz questions covering the material in those headings so that you can assess your knowledge of these specific areas. The answers to the “Do I Know This Already?” quiz appear in Appendix A.

Table 6-1 “Do I Know This Already?” Foundation Topics Section-to-Question Mapping

Foundations Topics Section	Questions
LSAs and the OSPF Link State Database	1–3
The Database Exchange Process	4, 5
Choosing the Best OSPF Routes	6–9

1. A network design shows area 1 with three internal routers, area 0 with four internal routers, and area 2 with five internal routers. Additionally, one ABR (ABR1) connects areas 0 and 1, plus a different ABR (ABR2) connects areas 0 and 2. How many Type 1 LSAs would be listed in ABR2's LSDB?
 - a. 6
 - b. 7
 - c. 15
 - d. 12
 - e. None of the other answers is correct.

2. A network planning diagram shows a large internetwork with many routers. The configurations show that OSPF has been enabled on all interfaces, IP addresses correctly configured, and OSPF working. For which of the following cases would you expect a router to create and flood a Type 2 LSA?
 - a. When OSPF is enabled on a LAN interface, and the router is the only router connected to the subnet.
 - b. When OSPF is enabled on a point-to-point serial link, and that router has both the higher router ID and higher interface IP address on the link.
 - c. When OSPF is enabled on a Frame Relay point-to-point subinterface, has the lower RID and lower subinterface IP address, and otherwise uses default OSPF configuration on the interface.
 - d. When OSPF is enabled on a working LAN interface on a router, and the router has been elected BDR.
 - e. None of the other answers is correct.

3. A verification plan shows a network diagram with branch office Routers B1 through B100, plus two ABRs, ABR1, and ABR2, all in area 100. The branches connect to the ABRs using Frame Relay point-to-point subinterfaces. The verification plan lists the output of the **show ip ospf database summary 10.100.0.0** command on a router B1, one of the branches. Which of the following is true regarding the output that could be listed for this command?
 - a. The output lists nothing unless 10.100.0.0 has been configured as a summary route using the **area range** command.
 - b. If 10.100.0.0 is a subnet in area 0, the output lists one Type 3 LSA, specifically the LSA with the lower metric when comparing ABR1's and ABR2's LSA for 10.100.0.0.
 - c. If 10.100.0.0 is a subnet in area 0, the output lists two Type 3 LSAs, one each created by ABR1 and ABR2.
 - d. None, because the Type 3 LSAs would exist only in the ABR's LSDBs.

4. Which of the following OSPF messages contains entire complete LSAs used during the database exchange process?
- LSR
 - LSAck
 - LSU
 - DD
 - Hello
5. Routers R1, R2, R3, and R4 connect to the same 10.10.10.0/24 LAN-based subnet. OSPF is fully working in the subnet. Later, R5, whose OSPF priority is higher than the other four routers, joins the subnet. Which of the following are true about the OSPF database exchange process over this subnet at this point? (Choose two.)
- R5 will send its DD, LSR, and LSU packets to the 224.0.0.5 all-DR-routers multicast address.
 - R5 will send its DD, LSR, and LSU packets to the 224.0.0.6 all-DR-routers multicast address.
 - The DR will inform R5 about LSAs by sending its DD, LSR, and LSU packets to the 224.0.0.6 all-SPF-routers multicast address.
 - The DR will inform R5 about LSAs by sending its DD, LSR, and LSU packets to the 224.0.0.5 all-SPF-routers multicast address.
6. R1 is internal to area 1, and R2 is internal to area 2. Subnet 10.1.1.0/24 exists in area 2 as a connected subnet off R2. ABR ABR1 connects area 1 to backbone area 0, and ABR2 connects area 0 to area 2. Which of the following LSAs must R1 use when calculating R1's best route for 10.1.1.0/24?
- R2's Type 1 LSA
 - Subnet 10.1.1.0/24's Type 2 LSA
 - ABR1's Type 1 LSA in area 0
 - Subnet 10.1.1.0/24's Type 3 LSA in Area 0
 - Subnet 10.1.1.0/24's Type 3 LSA in Area 1
7. Which of the following LSA types describes topology information that, when changed, requires a router in the same area to perform an SPF calculation? (Choose two.)
- 1
 - 2
 - 3
 - 4
 - 5
 - 7

8. The following output was taken from Router R3. A scan of R3's configuration shows that no **bandwidth** commands have been configured in this router. Which of the following answers lists configuration settings could be a part of a configuration that results in the following output? (Choose two.) Only two of the three interface's costs have been set directly.

R3#show ip ospf interface brief

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Se0/0/0.2	3	34	10.10.23.3/29	647	P2P	1/1	
Se0/0/0.1	3	34	10.10.13.3/29	1000	P2P	1/1	
Fa0/0	3	34	10.10.34.3/24	20	BDR	1/1	

- a. An **auto-cost reference-bandwidth 1000** command in router ospf mode
 - b. An **auto-cost reference-bandwidth 2000** command in router ospf mode
 - c. An **ip ospf cost 1000 interface S0/0/0.1** command in router ospf mode
 - d. An **auto-cost reference-bandwidth 64700** command in router ospf mode
9. Which of the following LSA types describe information related to topology or subnets useful for calculating routes for subnets inside the OSPF domain? (Choose three.)
- a. 1
 - b. 2
 - c. 3
 - d. 4
 - e. 5
 - f. 7

Foundation Topics

LSAs and the OSPF Link State Database

Every router that connects to a given OSPF area should learn the exact same topology data. Each router stores the data, composed of individual link state advertisements (LSA), in their own copy of the link state database (LSDB). Then, the router applies the Shortest Path First (SPF) algorithm to the LSDB to determine the best (lowest cost) route for each reachable subnet (prefix/length).

When a router uses SPF to analyze the LSDB, the SPF process has some similarities to how humans put a jigsaw puzzle together—but without a picture of what the puzzle looks like. Humans faced with such a challenge might first look for the obvious puzzle pieces, such as the corner and edge pieces, because they are easily recognized. You might then group puzzle pieces together if they have the same color or look for straight lines that might span multiple puzzle pieces. And of course, you would be looking at the shapes of the puzzle pieces to see which ones fit together.

Similarly, a router's SPF process must examine the individual LSAs and see how they fit together, based on their characteristics. To better appreciate the SPF process, the first section of this chapter examines the three LSA types OSPF uses to describe an Enterprise OSPF topology inside the OSPF domain. By understanding the types of LSAs, you can get a better understanding of what a router might look for to take the LSAs—the pieces of a network topology puzzle if you will—and build the equivalent of a network diagram.

For reference, Table 6-2 lists the various OSPF LSA types. Note that Chapter 9 explains three other LSA types, all used when redistributing routes into the OSPF domain.

Table 6-2 *OSPF LSA Types*

LSA Type	Common Name	Description
1	Router	Each router creates its own Type 1 LSA to represent itself for each area to which it connects. The LSDB for one area contains one Type 1 LSA per router per area, listing the RID and all interface IP addresses on that router that are in that area. Represents stub networks as well.
2	Network	One per transit network. Created by the DR on the subnet, and represents the subnet and the router interfaces connected to the subnet.
3	Net Summary	Created by ABRs to represent subnets listed in one area's type 1 and 2 LSAs when being advertised into another area. Defines the links (subnets) in the origin area, and cost, but no topology data.
4	ASBR Summary	Like a type 3 LSA, except it advertises a host route used to reach an ASBR.



Table 6-2 OSPF LSA Types

LSA Type	Common Name	Description
5	AS External	Created by ASBRs for external routes injected into OSPF.
6	Group Membership	Defined for MOSPF; not supported by Cisco IOS.
7	NSSA External	Created by ASBRs inside an NSSA area, instead of a type 5 LSA.
8	External Attributes	Not implemented in Cisco routers.
9–11	Opaque	Used as generic LSAs to allow for easy future extension of OSPF; for example, type 10 has been adapted for MPLS traffic engineering.

LSA Type 1: Router LSA

An LSA type 1, called a *router LSA*, identifies an OSPF router based on its OSPF router ID (RID). Each router creates a Type 1 LSA for itself and floods the LSA throughout the same area. To flood the LSA, the originating router sends the Type 1 LSA to its neighbors inside the same area, who in turn send it to their other neighbors inside the same area, until all routers in the area have a copy of the LSA.

Besides the RID of the router, this LSA also lists information about the attached links. In particular, the Type 1 LSA lists:

- For each interface on which no DR has been elected, it lists the router's interface subnet number/mask and interface OSPF cost. (OSPF refers to these subnets as *stub networks*.)
- For each interface on which a DR has been elected, it lists the IP address of the DR and a notation that the link attaches to a *transit network* (meaning a type 2 LSA exists for that network).
- For each interface with no DR, but for which a neighbor is reachable, it lists the neighbor's RID.

As with all OSPF LSAs, OSPF identifies a Type 1 LSA using a 32-bit *link state identifier* (LSID). When creating its own Type 1 LSA, each router uses its own OSPF RID value as the LSID.

Internal routers each create a single Type 1 LSA for themselves, but ABRs create multiple Type 1 LSAs for themselves: one per area. The Type 1 LSA in one area will list only interfaces in that area and only neighbors in that area. However, the router still has only a single RID, so all its Type 1 LSAs for a single router list the same RID. The ABR then floods each of its Type 1 LSAs into the appropriate area.

To provide a better backdrop for the upcoming LSA discussions, Figure 6-1 shows a sample internetwork, which will be used in most of the examples in this chapter.




Figure 6-1 Sample OSPF Multi-Area Design

All routers that participate in an area, be they internal routers or ABRs, create and flood a Type 1 LSA inside the area. For example, in Figure 6-1, area 5 has one internal router (R5, RID 5.5.5.5), and two ABRs: R1 with RID 1.1.1.1 and R2 with RID 2.2.2.2. Each of these three routers create and flood their own Type 1 LSA inside area 5 so that all three routers know the same three Type 1 LSAs.

Next, to further understand the details inside a Type 1 LSA, first consider the OSPF configuration of R5 as an example. R5 has three IP-enabled interfaces: Fa0/0, S0/0.1, and S0/0.2. R5 uses point-to-point subinterfaces, so R5 should form neighbor relationships with both R1 and R2 with no extra configuration beyond enabling OSPF, in area 5, on all three interfaces. Example 6-1 shows this baseline configuration on R5.

Example 6-1 R5 Configuration—IP Addresses and OSPF

```
interface FastEthernet0/0
 ip address 10.10.5.5 255.255.255.224
 ip ospf 5 area 5
!
interface s0/0.1 point-to-point
 ip address 10.10.15.5 255.255.255.248
 frame-relay interface-dlci 101
 ip ospf 5 area 5
!
interface s0/0.2 point-to-point
```

```

ip address 10.10.25.5 255.255.255.248
frame-relay interface-dlci 102
ip ospf 5 area 5
!
router ospf 5
  router-id 5.5.5.5
!
R5#show ip ospf interface brief

```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
se0/0.2	5	5	10.10.25.5/29	64	P2P	1/1	
se0/0.1	5	5	10.10.15.5/29	64	P2P	1/1	
fa0/0	5	5	10.10.5.5/27	1	DR	0/0	

```

R5#show ip ospf neighbor

```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	0	FULL/ -	00:00:30	10.10.25.2	Serial0/0.2
1.1.1.1	0	FULL/ -	00:00:38	10.10.15.1	Serial0/0.1

R5's OSPF configuration enables OSPF, for process ID 5, placing three interfaces in area 5. As a result, R5's type 1 LSA will list at least these three interfaces as links, plus it will refer to the two working neighbors. Example 6-2 displays the contents of R5's area 5 LSDB, including the detailed information in R5's Type 1 LSA, including the following:

- The LSID of R5's Type 1 LSA (5.5.5.5)
- Three links that connect to a stub network, each listing the subnet/mask
- Two links that state a connection to another router, one listing R1 (RID 1.1.1.1) and one listing R2 (RID 2.2.2.2)

Example 6-2 R5 Configuration—IP Addresses and OSPF

```

R5#show ip ospf database

```

OSPF Router with ID (5.5.5.5) (Process ID 5)

Router Link States (Area 5)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	835	0x80000002	0x006BDA	2
2.2.2.2	2.2.2.2	788	0x80000002	0x0082A6	2
5.5.5.5	5.5.5.5	787	0x80000004	0x0063C3	5

Summary Net Link States (Area 5)

```

Link ID          ADV Router      Age           Seq#           Checksum
10.10.12.0       1.1.1.1         835          0x80000001   0x00F522
10.10.12.0       2.2.2.2         787          0x80000001   0x00D73C
! lines omitted for brevity

```

```
R5#show ip ospf database router 5.5.5.5
```

```
OSPF Router with ID (5.5.5.5) (Process ID 5)
```

```
Router Link States (Area 5)
```

```
LS age: 796
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: Router Links
```

```
Link State ID: 5.5.5.5
```

```
Advertising Router: 5.5.5.5
```

```
LS Seq Number: 80000004
```

```
Checksum: 0x63C3
```

```
Length: 84
```

```
Number of Links: 5
```

```
Link connected to: another Router (point-to-point)
```

```
(Link ID) Neighboring Router ID: 2.2.2.2
```

```
(Link Data) Router Interface address: 10.10.25.5
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 64
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 10.10.25.0
```

```
(Link Data) Network Mask: 255.255.255.248
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 64
```

```
Link connected to: another Router (point-to-point)
```

```
(Link ID) Neighboring Router ID: 1.1.1.1
```

```
(Link Data) Router Interface address: 10.10.15.5
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 64
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 10.10.15.0
```

```
(Link Data) Network Mask: 255.255.255.248
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 64
```

```

Link connected to: a Stub Network
(Link ID) Network/subnet number: 10.10.5.0
(Link Data) Network Mask: 255.255.255.224
Number of TOS metrics: 0
TOS 0 Metrics: 1
    
```

The first command, **show ip ospf database**, displays a summary of the LSAs known to R5. The output mainly consists of a single line per LSA, listed by LSA ID. The three highlighted lines of this command, in Example 6-2, highlight the RID of the three router (Type 1) LSAs, namely 1.1.1.1 (R1), 2.2.2.2 (R2), and 5.5.5.5 (R5).

The output of the **show ip ospf database router 5.5.5.5** command displays the detailed information in R5's router LSA. Looking at the highlighted portions, you see three stub networks—three interfaces on which no DR has been elected—and the associated subnet numbers. The LSA also lists the neighbor IDs of two neighbors (1.1.1.1 and 2.2.2.2) and the interfaces on which these neighbors can be reached.

Armed with the same kind of information in R1's and R2's Type 1 LSAs, a router has enough information to determine which routers connect, over which stub links, and then use the interface IP address configuration to figure out the interfaces that connect to the other routers. Figure 6-2 shows a diagram of area 5 that could be built just based on the detailed information held in the router LSAs for R1, R2, and R5.

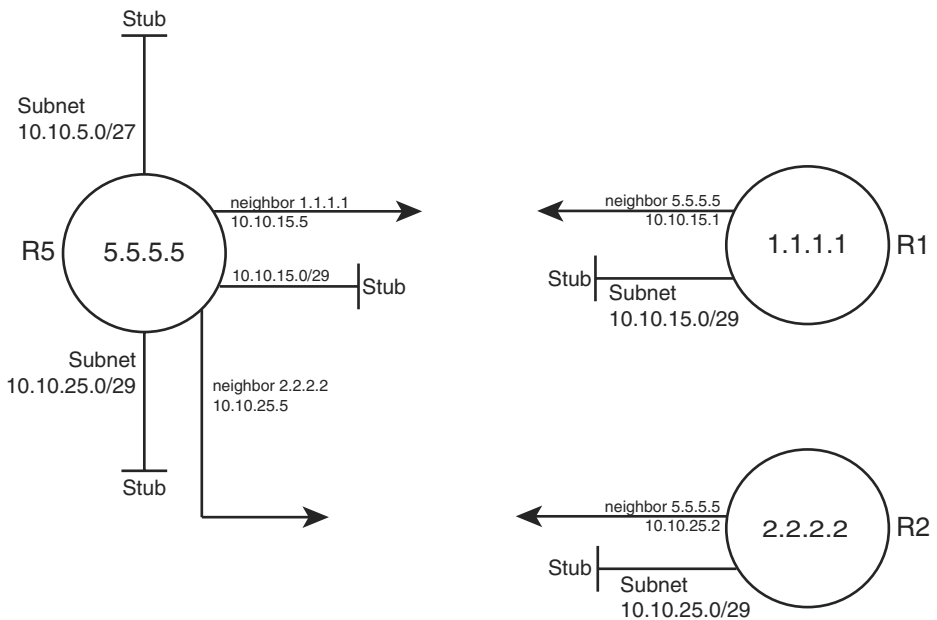


Figure 6-2 Three Type 1 LSAs in Area 5

Note that Figure 6-2 displays only information that could be learned from the Type 1 router LSAs inside area 5. Each Type 1 router LSA lists information about a router but only the details related to a specific area. As a result, Figure 6-2 shows R1's interface in area 5 but none of the interfaces in area 34 nor in area 0. To complete the explanation surrounding Figure 6-2, Example 6-3 lists R1's Type 1 router LSA for area 5.

Example 6-3 R1's Type 1 LSA in Area 5

```
R5#show ip ospf database router 1.1.1.1

          OSPF Router with ID (5.5.5.5) (Process ID 5)

          Router Link States (Area 5)

Routing Bit Set on this LSA
LS age: 1306
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 1.1.1.1
Advertising Router: 1.1.1.1
LS Seq Number: 80000002
Checksum: 0x6BDA
Length: 48
Area Border Router
Number of Links: 2

Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 5.5.5.5
(Link Data) Router Interface address: 10.10.15.1
Number of TOS metrics: 0
TOS 0 Metrics: 64

Link connected to: a Stub Network
(Link ID) Network/subnet number: 10.10.15.0
(Link Data) Network Mask: 255.255.255.248
Number of TOS metrics: 0
TOS 0 Metrics: 64
```

Note: Because the OSPF uses the RID for many purposes inside different LSAs—for instance, as the LSID of a type 1 LSA—Cisco recommends setting the RID to a stable, predictable value. To do this, use the OSPF `router-id value` OSPF subcommand, or define a loopback interface with an IP address, as discussed in Chapter 5's section “Using a Unique OSPF Router ID.”

LSA Type 2: Network LSA

SPF requires that the LSDB model the topology with nodes (routers) and connections between nodes (links). In particular, each link must be between a pair of nodes. When a multiaccess data link exists—for instance, a LAN—OSPF must somehow model that LAN so that the topology represents nodes and links between only a pair of nodes. To do so, OSPF uses the concept of a Type 2 Network LSA.

OSPF routers actually choose whether to use a Type 2 LSA for a multiaccess network based on whether a designated router (DR) has or has not been elected on an interface. So, before discussing the details of the Type 2 network LSA, a few more facts about the concept of a DR need to be discussed.

Background on Designated Routers

As discussed in Chapter 5's section "OSPF Network Types," the OSPF network type assigned to a router interface tells that router whether to attempt to elect a DR on that interface. Then, when a router has heard a Hello from at least one other router, the routers elect a DR and BDR.

OSPF uses a DR in a particular subnet for two main purposes:

- To create and flood a Type 2 network LSA for that subnet
- To aid in the detailed process of database exchange over that subnet

Routers elect a DR, and a backup DR (BDR), based on information in the OSPF Hello. The Hello message lists each router's RID and a priority value. When no DR exists at the time, routers use the following election rules when neither a DR nor BDR yet exists:

- Choose the router with the highest priority (default 1, max 255, set with `ip ospf priority value` interface subcommand).
- If tied on priority, choose the router with highest RID.
- Choose a BDR, based on next-best priority, or if a tie, next-best (highest) RID.

Although the preceding describes the election when no DR currently exists, the rules differ a bit when a DR and BDR already exist. After a DR and BDR are elected, no election is held until either the DR or BDR fails. If the DR fails, the BDR becomes the DR—regardless of whether a higher priority router has joined the subnet—and a new election is held to choose a new BDR. If the BDR fails, a new election is held for BDR, and the DR remains unchanged.

On LANs, the choice of DR matters little from a design perspective, but does matter from an operational perspective. Throughout this chapter, note the cases in which output of `show` commands identify the DR and its role. Now, back to the topic of Type 2 LSAs.

Note: On Frame Relay WAN links, the choice of DR may impact whether OSPF functions at all. This topic is covered in Chapter 8, "OSPF Virtual Links and Frame Relay Operations."



Type 2 Network LSA Concepts

OSPF uses the concept of a Type 2 LSA to model a multiaccess network—a network with more than two routers connected to the same subnet—while still conforming to the “a link connects only two nodes” rule for the topology. For example, consider the network in Figure 6-3 (also shown as Figure 5-4 in the previous chapter). As seen in Chapter 5, all four routers form neighbor relationships inside area 0, with the DR and BDR becoming fully adjacent with the other routers.




Figure 6-3 *Small Network, Four Routers, on a LAN*

OSPF cannot represent the idea of four routers connected via a single subnet by using a link connected to all four routers. Instead, OSPF defines the Type 2 network LSA, used as a *pseudonode*. Each router’s Type 1 router LSA lists a connection to this pseudonode, often called a *transit network*, which is then modeled by a Type 2 network LSA. The Type 2 network LSA itself then lists references back to each Type 1 router LSA connected to it—four in this example, as shown in Figure 6-4.

The elected DR in a subnet creates the Type 2 LSA for that subnet. The DR identifies the LSA by assigning an LSID of the DR’s interface IP address in that subnet. The type 2 LSA also lists the DR’s RID as the router advertising the LSA.

Type 2 LSA show Commands

To see these concepts in the form of OSPF **show** commands, next consider area 34 back in Figure 6-1. This design shows that R3 and R4 connect to the same LAN, which means that

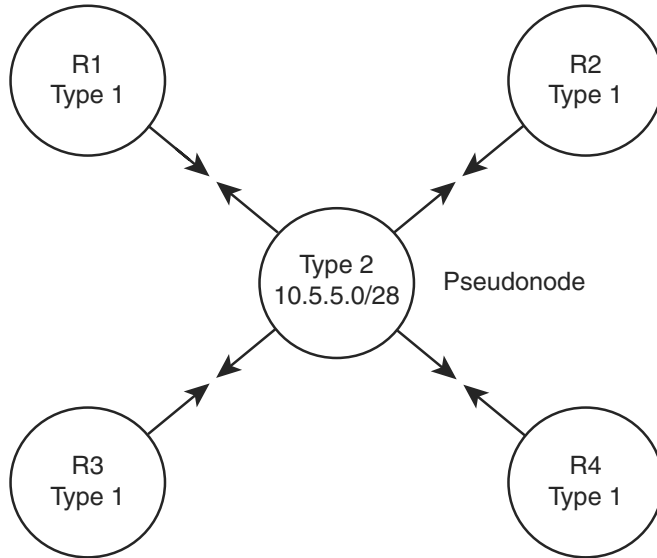


Figure 6-4 OSPF Topology when Using a Type 2 Network LSA

a DR will be elected. (OSPF elects a DR on LANs when at least two routers pass the neighbor requirements and can become neighbors.) If both R3 and R4 default to use priority 1, then R4 wins the election, due to its 4.4.4.4 RID (versus R3's 3.3.3.3). So, R4 creates the Type 2 LSA for that subnet and floods the LSA. Figure 6-5 depicts the area 34 topology, and Example 6-4 shows the related LSDB entries.

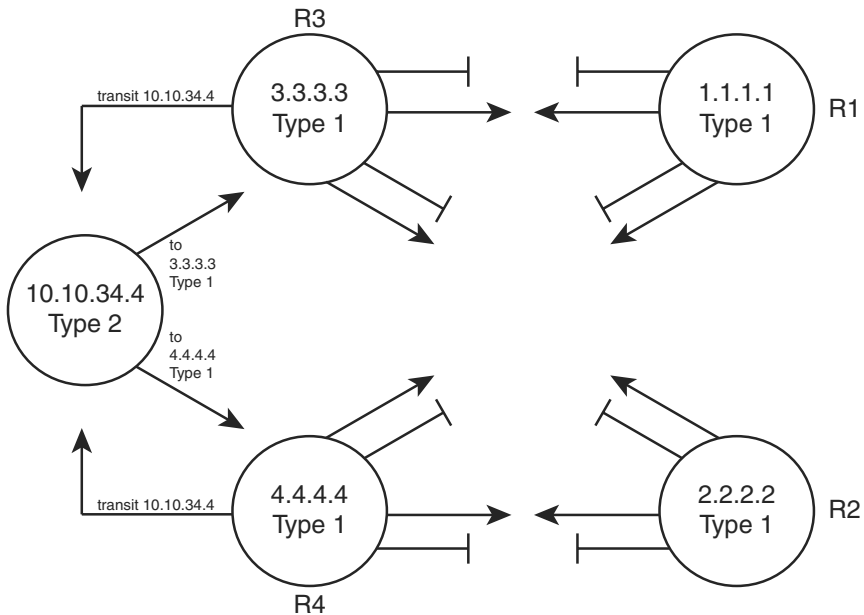


Figure 6-5 Area 34 Topology with Four Type 1 LSAs and One Type 2 LSA

Example 6-4 Area 34 LSAs for R3, Network 10.10.34.0/24

```
R3#show ip ospf database
```

```
OSPF Router with ID (3.3.3.3) (Process ID 3)
```

```
Router Link States (Area 34)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	1061	0x80000002	0x00EA7A	4
2.2.2.2	2.2.2.2	1067	0x80000001	0x0061D2	4
3.3.3.3	3.3.3.3	1066	0x80000003	0x00E2E8	5
4.4.4.4	4.4.4.4	1067	0x80000003	0x007D3F	5

```
Net Link States (Area 34)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.10.34.4	4.4.4.4	1104	0x80000001	0x00AB28

```
Summary Net Link States (Area 34)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.10.5.0	1.1.1.1	1023	0x80000001	0x00BF2
10.10.5.0	2.2.2.2	1022	0x80000001	0x00EC0D

```
! lines omitted for brevity
```

```
R3#show ip ospf database router 4.4.4.4
```

```
OSPF Router with ID (3.3.3.3) (Process ID 3)
```

```
Router Link States (Area 34)
```

```
LS age: 1078
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: Router Links
```

```
Link State ID: 4.4.4.4
```

```
Advertising Router: 4.4.4.4
```

```
LS Seq Number: 80000003
```

```
Checksum: 0x7D3F
```

```
Length: 84
```

Number of Links: 5

Link connected to: another Router (point-to-point)
 (Link ID) Neighboring Router ID: 2.2.2.2
 (Link Data) Router Interface address: 10.10.24.4
 Number of TOS metrics: 0
 TOS 0 Metrics: 64

Link connected to: a Stub Network
 (Link ID) Network/subnet number: 10.10.24.0
 (Link Data) Network Mask: 255.255.255.248
 Number of TOS metrics: 0
 TOS 0 Metrics: 64

Link connected to: another Router (point-to-point)
 (Link ID) Neighboring Router ID: 1.1.1.1
 (Link Data) Router Interface address: 10.10.14.4
 Number of TOS metrics: 0
 TOS 0 Metrics: 64

Link connected to: a Stub Network
 (Link ID) Network/subnet number: 10.10.14.0
 (Link Data) Network Mask: 255.255.255.248
 Number of TOS metrics: 0
 TOS 0 Metrics: 64

Link connected to: a Transit Network
 (Link ID) Designated Router address: 10.10.34.4
 (Link Data) Router Interface address: 10.10.34.4
 Number of TOS metrics: 0
 TOS 0 Metrics: 1

R3#show ip ospf database network 10.10.34.4

OSPF Router with ID (3.3.3.3) (Process ID 3)

Net Link States (Area 34)

Routing Bit Set on this LSA
 LS age: 1161
 Options: (No TOS-capability, DC)
 LS Type: Network Links
 Link State ID: 10.10.34.4 (address of Designated Router)
 Advertising Router: 4.4.4.4

```

LS Seq Number: 80000001
Checksum: 0xAB28
Length: 32
Network Mask: /24
Attached Router: 4.4.4.4
Attached Router: 3.3.3.3

```

The **show ip ospf database** command lists a single line for each LSA. Note that the (highlighted) heading for network LSAs lists one entry, with LSID 10.10.34.4, which is R4's Fa0/0 IP address. The LSID for Type 2 Network LSAs is the interface IP address of the DR that creates the LSA.

The **show ip ospf database router 4.4.4.4** command shows the new style of entry for the reference to a *Transit Network*, which again refers to a connection to a Type 2 LSA. The output lists a LSID of 10.10.34.4, which again is the LSID of the Type 2 LSA.

Finally, the **show ip ospf database network 10.10.34.4** command shows the details of the Type 2 LSA, based on its LSID of 10.10.34.4. Near the bottom, the output lists the attached routers, based on RID. The SPF process can then use the cross-referenced information, as shown in Figure 6-5, to determine which routers connect to this transit network (pseudonode). The SPF process has information in both the Type 1 LSAs that refer to the transit network link to a Type 2 LSA, and the Type 2 LSA has a list of RIDs of Type 1 LSAs that connect to the Type 2 LSA, making the process of modeling the network possible.

OSPF can model all the topology inside a single area using Type 1 and 2 LSAs. When a router uses its SPF process to build a model of the topology, it can then calculate the best (lowest cost) route for each subnet in the area. The next topic completes the LSA picture for internal OSPF routes by looking at Type 3 LSAs, which are used to model interarea routes.

LSA Type 3: Summary LSA

OSPF areas exist in part so that engineers can reduce the consumption of memory and compute resources in routers. Instead of having all routers, regardless of area, know all Type 1 and Type 2 LSAs inside an OSPF domain, ABRs do not forward Type 1 and Type 2 LSAs from one area into another area, and vice versa. This convention results in smaller per-area LSDBs, saving memory and reducing complexity for each run of the SPF algorithm, which saves CPU and improves convergence time.

However, even though ABRs do not flood Type 1 and Type 2 LSAs into other areas, routers still need to learn about subnets in other areas. OSPF advertises these interarea routes using the Type 3 summary LSA. ABRs generate a Type 3 LSA for each subnet in one area, and advertises each Type 3 LSA into the other areas.

For example, if subnet A exists in area 3, then the routers in area 3 learn of that subnet as part of Type 1 and Type 2 LSAs. However, an ABR connected to area 3 will not forward

the Type 1 and Type 2 LSAs into other areas, instead creating a Type 3 LSA for each subnet (including subnet A). The routers inside the other areas can then calculate a route for the subnets (like subnet A) that exist inside another area.

Type 3 summary LSAs do not contain all the detailed topology information, so in comparison to Types 1 and 2, these LSAs summarize the information—hence the name *summary LSA*. Conceptually, a Type 3 LSA appears to be another subnet connected to the ABR that created and advertised the Type 3 LSA. The routers inside that area can calculate their best route to reach the ABR, which gives the router a good loop-free route to reach the subnet listed in a Type 3 LSA.

An example can certainly help in this case. First, consider the comparison shown in the top and bottom of Figure 6-6. The top depicts the topology shown back in Figure 6-1 if that design had used a single area. In that case, every router would have a copy of each Type 1 LSA (shown as a router name in the figure), and each Type 2 (abbreviated as T2 in the figure). The bottom of Figure 6-6 shows the area 5 topology, when holding to the three area design shown in Figure 6-1.

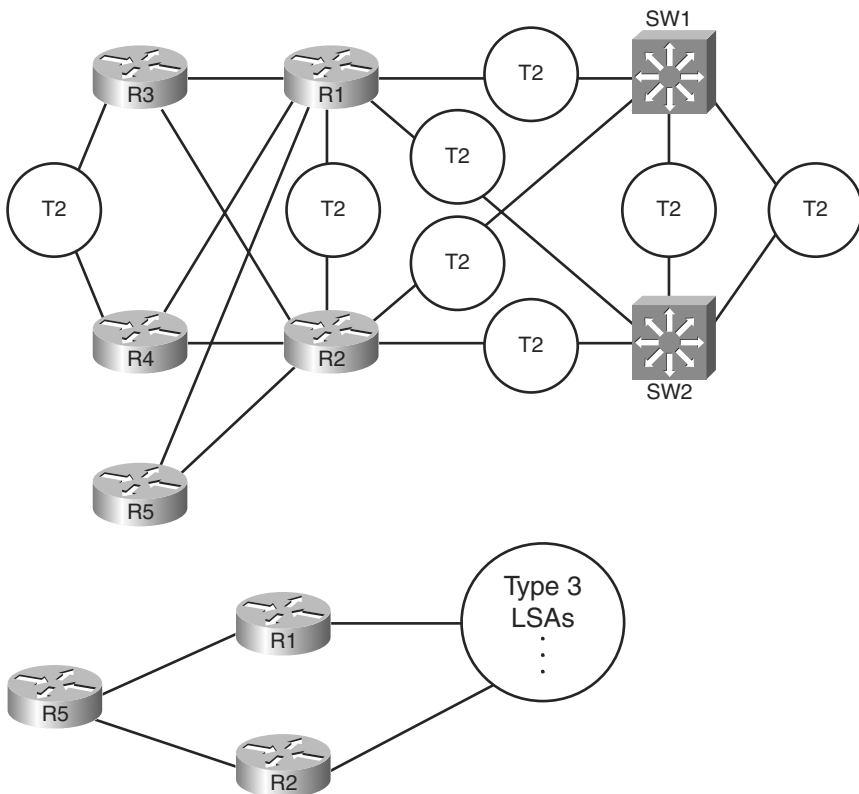


Figure 6-6 Comparing a Single Area LSDB to a Three Area LSDB

The ABR creates and floods each Type 3 LSA into the next area. The ABR assigns an LSID of the subnet number being advertised. It also adds its own RID to the LSA as well, so that routers know which ABR advertised the route. It also includes the subnet mask. The correlation between the advertising router's RID and the LSID (subnet number) allows the OSPF processes to create the part of the topology as shown with Type 3 LSAs at the bottom of Figure 6-6.

Example 6-5 focuses on the Type 3 LSAs in Area 34 of the network shown in Figure 6-1. Ten subnets exist outside area 34. As ABRs, both R1 and R2 create and flood a Type 3 LSA for each of these 10 subnets, resulting in 20 Type 3 LSAs listed in the output of the `show ip ospf database` command inside area 34. Then, the example focuses specifically on the Type 3 LSA for subnet 10.10.99.0/24.

Example 6-5 *Type 3 LSAs in Area 34*

```
R3#show ip ospf database
```

```
OSPF Router with ID (3.3.3.3) (Process ID 3)
```

```
Router Link States (Area 34)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	943	0x80000003	0x00E87B	4
2.2.2.2	2.2.2.2	991	0x80000002	0x005FD3	4
3.3.3.3	3.3.3.3	966	0x80000004	0x00E0E9	5
4.4.4.4	4.4.4.4	977	0x80000004	0x007B40	5

```
Net Link States (Area 34)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.10.34.4	4.4.4.4	977	0x80000002	0x00A929

```
Summary Net Link States (Area 34)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.10.5.0	1.1.1.1	943	0x80000002	0x0009F3
10.10.5.0	2.2.2.2	991	0x80000002	0x00EA0E
10.10.12.0	1.1.1.1	943	0x80000002	0x00F323
10.10.12.0	2.2.2.2	991	0x80000002	0x00D53D
10.10.15.0	1.1.1.1	943	0x80000002	0x0021BA
10.10.15.0	2.2.2.2	993	0x80000003	0x008313
10.10.17.0	1.1.1.1	946	0x80000002	0x00BC55
10.10.17.0	2.2.2.2	993	0x80000002	0x00A864
10.10.18.0	1.1.1.1	946	0x80000002	0x00B15F
10.10.18.0	2.2.2.2	994	0x80000002	0x009D6E

10.10.25.0	1.1.1.1	946	0x80000002	0x00355C
10.10.25.0	2.2.2.2	993	0x80000002	0x009439
10.10.27.0	1.1.1.1	946	0x80000002	0x0058AE
10.10.27.0	2.2.2.2	993	0x80000002	0x0030D3
10.10.28.0	1.1.1.1	947	0x80000002	0x004DB8
10.10.28.0	2.2.2.2	993	0x80000002	0x0025DD
10.10.98.0	1.1.1.1	946	0x80000002	0x004877
10.10.98.0	2.2.2.2	993	0x80000002	0x002A91
10.10.99.0	1.1.1.1	946	0x80000002	0x003D81
10.10.99.0	2.2.2.2	993	0x80000002	0x001F9B

```
R3#show ip ospf database summary 10.10.99.0
```

```
OSPF Router with ID (3.3.3.3) (Process ID 3)
```

```
Summary Net Link States (Area 34)
```

```
Routing Bit Set on this LSA
```

```
LS age: 1062
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 10.10.99.0 (summary Network Number)
```

```
Advertising Router: 1.1.1.1
```

```
LS Seq Number: 80000002
```

```
Checksum: 0x3D81
```

```
Length: 28
```

```
Network Mask: /24
```

```
TOS: 0 Metric: 2
```

```
Routing Bit Set on this LSA
```

```
LS age: 1109
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 10.10.99.0 (summary Network Number)
```

```
Advertising Router: 2.2.2.2
```

```
LS Seq Number: 80000002
```

```
Checksum: 0x1F9B
```

```
Length: 28
```

```
Network Mask: /24
```

```
TOS: 0 Metric: 2
```

Note: The Type 3 Summary LSA is not used for the purpose of route summarization. OSPF does support route summarization, and Type 3 LSAs may indeed advertise such a summary, but the Type 3 LSA does not inherently represent a summary route. The term

summary reflects the idea that the information is sparse compared to the detail inside Type 1 and Type 2 LSAs.

The upcoming section “Calculating the Cost of Inter-area Routes” discusses how a router determines the available routes to reach subnets listed in a Type 3 LSA and how a router chooses which route is best.

Limiting the Number of LSAs

By default, Cisco IOS does not limit the number of LSAs a router can learn. However, it may be useful to protect a router from learning too many LSAs to protect router memory. Also, with a large number of LSAs, the router may be unable to process the LSDB with SPF well enough to converge in a reasonable amount of time.

The maximum number of LSAs learned from other routers can be limited by a router using the `max-lsa number` OSPF subcommand. When configured, if the router learns more than the configured number of LSAs from other routers (ignoring those created by the router itself), the router reacts. The first reaction is to issue log messages. The router ignores the event for a time period, after which the router repeats the warning message. This ignore-and-wait strategy can proceed through several iterations, ending when the router closes all neighborships, discards its LSDB, and then starts adding neighbors again. (The ignore time, and the number of times to ignore the event, can be configured with the `max-lsa` command.)

Summary of Internal LSA Types

OSPF uses Type 1, 2, and 3 LSAs to calculate the best routes for all routes inside the OSPF routing domain. Later, Chapter 9 explains Types 4, 5, and 7, which OSPF uses to calculate routes for external routes—routes redistributed into OSPF.

Table 6-3 summarizes some of the key points regarding OSPF Type 1, 2, and 3 LSAs. In particular for the ROUTE exam, the ability to sift through the output of various `show ip ospf database` commands can be important. Knowing what the OSPF LSID represents can help you interpret the output, and knowing the keywords used with the `show ip ospf database lsa-type lsid` commands can also be very useful. Table 6-3 summarizes these details.

Table 6-3 *Facts about LSA Types 1, 2, and 3*

LSA Type (Number)	LSA Type (Name)	This Type Represents	Display Using <code>show ip ospf database keyword...</code>	LSID Is Equal To	Created By
1	Router	A router	<code>router</code>	RID of router	Each router creates its own



Table 6-3 *Facts about LSA Types 1, 2, and 3*

LSA Type (Number)	LSA Type (Name)	This Type Represents	Display Using show ip ospf database keyword...	LSID Is Equal To	Created By
2	Network	A subnet in which a DR exists	network	DR's IP address in the subnet	The DR in that subnet
3	Summary	Subnet in another area	summary	Subnet number	An ABR

The Database Exchange Process

Every router in an area, when OSPF stabilizes after topology changes occur, should have an identical LSDB for that area. Internal routers (routers inside a single area) have only that area's LSAs, but an ABR's LSDB will contain LSAs for each area to which it connects. The ABR does, however, know which LSAs exist in each area.

OSPF routers flood both the LSAs they create, and the LSAs they learn from their neighbors, until all routers in the area have a copy of each of the most recent LSAs for that area. To manage and control this process, OSPF defines several messages, processes, and neighbor states that indicate the progress when flooding LSAs to each neighbor. This section begins by listing reference information for the OSPF messages and neighbor states. Next, the text describes the flooding process between two neighbors when a DR does not exist, followed by a description of the similar process used when a DR does exist. This section ends with a few items related to how the routers avoid looping the LSA advertisements and how they periodically reflood the information.

OSPF Message and Neighbor State Reference

For reference, Table 6-4 lists the OSPF message types that will be mentioned in the next few pages. Additionally, Table 6-5 lists the various neighbor states as well. Although useful for study, when you are first learning this topic, feel free to skip these tables for now.

Table 6-4 *OSPF Message Types and Functions*

Message Name/number	Description
Hello (1)	Used to discover neighbors, supply information used to confirm two routers should be allowed to become neighbors, to bring a neighbor relationship to a 2-way state, and to monitor a neighbor's responsiveness in case it fails
Database Description (DD or DBD) (2)	Used to exchange brief versions of each LSA, typically on initial topology exchange, so that a router knows a list of that neighbor's known LSAs

Table 6-4 *OSPF Message Types and Functions*

Message Name/number	Description
Link-State Request (LSR) (3)	A packet that lists the LSIDs of LSAs the sender of the LSR would like the receiver of the LSR to supply during database exchange
Link-State Update (LSU) (4)	A packet that contains fully detailed LSAs, typically sent in response to an LSR message
Link-State Acknowledgment (LSAck) (5)	Sent to confirm receipt of an LSU message

Table 6-5 *OSPF Neighbor State Reference*

State	Meaning
Down	No Hellos have been received from this neighbor for more than the dead interval.
Attempt	Used when the neighbor is defined with the neighbor command, after sending a Hello, but before receiving a Hello from that neighbor.
Init	A Hello has been received from the neighbor, but it did not have the local router's RID in it or lists parameters that do not pass the neighbor verification checks. This is a permanent state when Hello parameters do not match.
2Way	A Hello has been received from the neighbor, it has the router's RID in it, and all neighbor verification checks passed.
ExStart	Currently negotiating the DD sequence numbers and master/slave logic used for DD packets.
Exchange	Finished negotiating the DD process particulars, and currently exchanging DD packets.
Loading	All DD packets are exchanged, and the routers are currently sending LSR, LSU, and LSAck packets to exchange full LSAs.
Full	Neighbors are fully adjacent, meaning they believe that their LSDBs for that area are identical. Routing table (re)calculations can begin.



Exchange Without a Designated Router

As discussed in Chapter 5, the OSPF interface network type tells that router whether to attempt to elect a DR on that interface. The most common case for which routers do not elect a DR occur on point-to-point topologies, such as true point-to-point serial links and point-to-point subinterfaces. This section examines the database exchange process on such interfaces, in preparation for the slightly more complex process when using a DR on an OSPF broadcast network type, like a LAN.

Every OSPF neighborhood begins by exchanging Hellos until the neighbors (hopefully) reach the 2-Way state. During these early stages, the routers discover each other by sending multicast Hellos and then check each other's parameters to make sure all items match

(as listed in Chapter 5's Table 5-5). Figure 6-7 shows the details, with the various neighbor states listed on the outside of the figure and the messages listed in the middle.

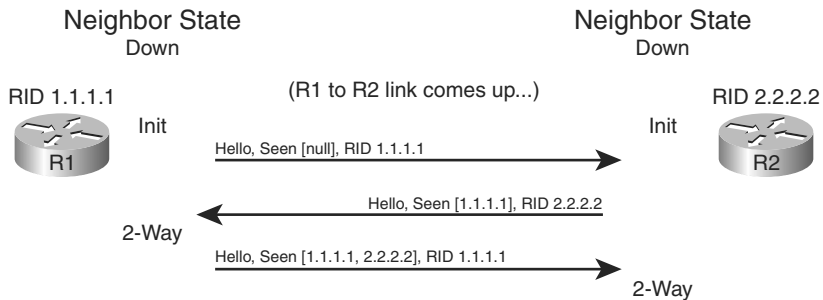


Figure 6-7 Neighbor Initialization—Early Stages

Figure 6-7 shows an example that begins with a failed neighborship, so the neighborship is in a down state. When a router tries to reestablish the neighborship, each router sends a multicast Hello and moves to an INIT state. After a router has both received a Hello and verified that all the required parameters agree, the router lists the other router's RID in the Hello as being seen, as shown in the bottom two Hello messages in the figure. When a router receives a Hello that lists its own RID as having been seen by the other router, the router can transition to 2-Way state.

When a router has reached the 2-Way state with a neighbor, as shown at the end of Figure 6-7, the router then decides whether it should exchange its LSDB entries. When no DR exists, the answer is always “yes.” Each router next follows this general process:

- Step 1.** Discover the LSAs known to the neighbor but unknown to me.
- Step 2.** Discover the LSAs known by both routers, but the neighbor's LSA is more up to date.
- Step 3.** Ask the neighbor for a copy of all the LSAs identified in the first two steps.

Figure 6-8 details the messages and neighbor states used to exchange the LSAs between two neighbors.

Figure 6-8 shows many details. As with Figure 6-7, Figure 6-8 shows neighbor states on the outer edges of the flows (refer to Table 6-5 for reference). Routers display these neighbor states (in the `show ip ospf neighbor` command variants), so a particular state may be useful in determining how far two neighbors have gotten in the database exchange process. The more important neighbor states will be mentioned throughout the chapter.

The inner portions of Figure 6-8 represent the OSPF message flows, with Table 6-2 earlier in the chapter listing the messages for reference. The next several pages examine the process shown in Figure 6-8 in more detail.

Discovering a Description of the Neighbor's LSDB

After a router has decided to move forward from 2-Way state and exchange its LSDB with a neighbor, the routers use the sequence shown in Figure 6-8. The next step in that process

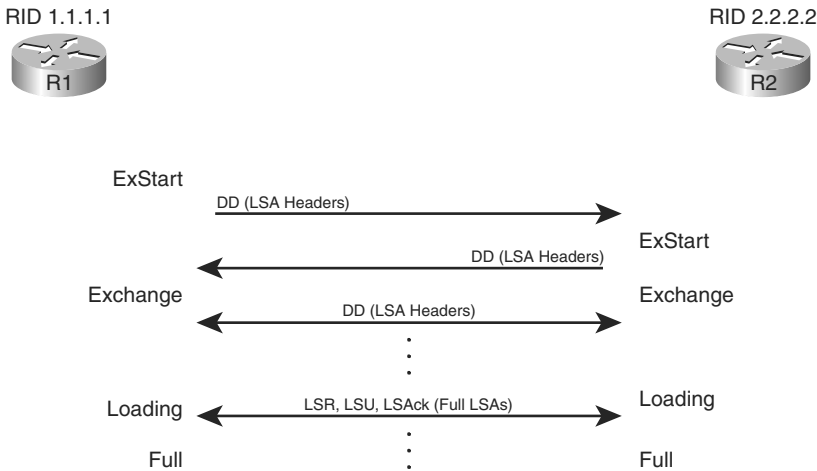


Figure 6-8 Overview of the Database Exchange Process Between Two Neighbors

requires both routers to tell each other the LSIDs of all their known LSAs in that area. The primary goal is for each neighbor to realize which LSAs it does not know, so it can then ask for those full LSAs to be sent. To learn the list of LSAs known by the neighbor, the neighboring routers follow these steps:

- Step 1.** Multicast database description packets (abbreviated as both DD and DBD, depending on the reference) to 224.0.0.5, which is the all SPF routers multicast address.
- Step 2.** When sending the first DD message, transition to the *ExStart* state until one router, the one with the higher RID, becomes master in a master/slave relationship.
- Step 3.** After electing a master, transition the neighbor to the *Exchange* state.
- Step 4.** Continue multicasting DD messages to each other until both routers have the same shared view of the LSIDs known collectively by both routers, in that area.

Note that the DD messages themselves do not list the entire LSAs, but rather LSA headers. These headers include the LSIDs of the LSAs and the LSA sequence number. The LSA sequence number for an LSA begins at value 0x80000001 (hex) when initially created; the router creating the LSA increments the sequence number, and refloods the LSA, whenever the LSA changes. For example, if an interface moves from up to down state, that router changes its Type 1 LSA to list that interface state as down, increments the LSA sequence number, and refloods the LSA.

The master router for each exchange controls the flow of DD messages, with the slave responding to the master's DD messages. The master keeps sending DD messages until it lists all its known LSIDs in that area. The slave responds by placing LSA headers in its DD messages. Some of those LSA headers simply repeat what the slave heard from the master, for the purpose of acknowledging to the master that the slave learned that LSA header

from the master. Additionally, the slave also includes the LSA headers for any LSAs that the master did not list.

This exchange of DD messages ends with each router knowing a list of LSAs that it does not have in its LSDB, but that the other router does have those LSAs. Additionally, each router also ends this process with a list of LSAs that the local router already knows, but for which the other router has a more recent copy (based on the sequence number).

Exchanging the LSAs

When the two neighbors realize that they have a shared view of the list of LSIDs, they transition to the Loading state and start exchanging the full LSAs—but only those that they do not yet know about or those that have changed.

For example, when the two routers in Figure 6-8 first become neighbors, neither router will have a copy of the Type 1 LSA for the other router. So, R1 will request that R2 send its LSA with LSID 2.2.2.2; R2 will send its Type 1 LSA; and R1 will acknowledge receipt. The mechanics work like this:

- Step 1.** Transition the neighbor state to Loading.
- Step 2.** For any missing LSAs, send a *Link State Request* (LSR) message, listing the LSID of the requested LSA.
- Step 3.** Respond to any LSR messages with an *Link State Update* (LSU), listing one or more LSAs in each message.
- Step 4.** Acknowledge receipt by either sending a *Link State Acknowledgment* (LSAck) message (called explicit acknowledgment), or by sending the same LSA that was received back to the other router in an LSU message (implicit acknowledgment).
- Step 5.** When all LSAs have been sent, received, and acknowledged, transition the neighborship to the FULL state (fully adjacent).

Note: Because this section examines the case without a DR, all these messages flow as multicasts to 224.0.0.5, the all SPF routers multicast address, unless the neighbors have been defined with an OSPF **neighbor** command.

By the end of this process, both routers should have an identical LSDB for the area to which the link has been assigned. At that point, the two routers can run the SPF algorithm to choose the currently best routes for each subnet.

Exchange with a Designated Router

Database exchange with a DR differs slightly than database exchange when no DR exists. The majority of the process is similar, with the same messages, meanings, and neighbor states. The big difference is the overriding choice of with whom each router chooses to perform database exchange.

Non-DR routers do not exchange their databases directly with all neighbors on a subnet. Instead, they exchange their database with the DR. Then, the DR exchanges any new/changed LSAs with the rest of the OSPF routers in the subnet.

The concept actually follows along with the idea of a Type 2 LSA as seen earlier in Figure 6-4. Figure 6-9 represents four Type 1 LSAs, for four real routers on the same LAN, plus a single Type 2 LSA that represents the multiaccess subnet. The DR created the Type 2 LSA as part of its role in life.

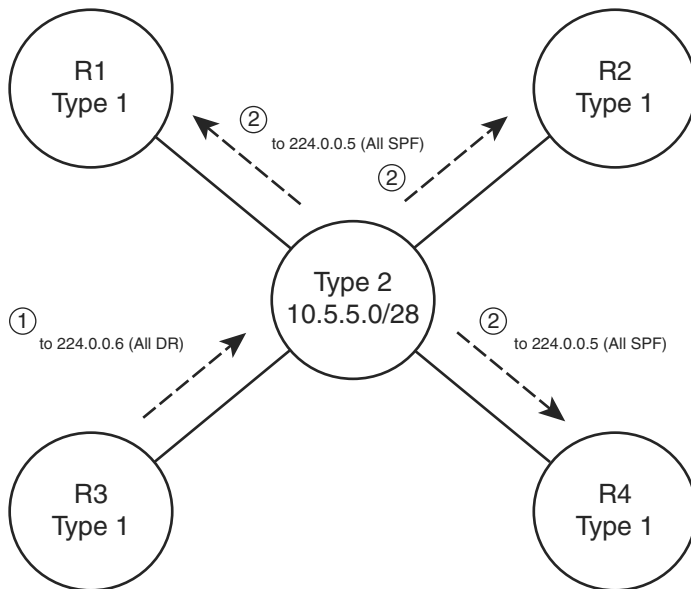


Figure 6-9 *Conceptual View—Exchanging the Database with a Pseudonode*

Figure 6-9 shows two conceptual steps for database exchange. The non-DR router (R3) first exchanges its database with the pseudonode, and then the type 2 pseudonode exchanges its database with the other routers. However, the pseudonode is a concept, not a router; to make the process depicted in Figure 6-9 work, the DR takes on the role of the Type 2 pseudonode. The messages differ slightly as well, as follows:

- The non-DR performs database exchange with the same messages, as shown in Figure 6-9, but sends these messages to the 224.0.0.6 All-DR-routers multicast address.
- The DR performs database exchange with the same messages but sends the messages to the 224.0.0.5 all-SPF-routers multicast address.



Consider these two conventions one at a time. First, the messages sent to 224.0.0.6 are processed by the DR and the BDR only. The DR actively participates, replying to the messages, with the BDR acting as a silent bystander. In effect, this allows the non-DR router to exchange their database directly with the DR and BDR, but with none of the other routers in the subnet.

Next, consider the multicast messages from the DR to the 224.0.0.5 all-SPF-router multicast address. All OSPF routers process these messages, so the rest of the routers—the DROthers to use the IOS term—also learn the newly exchanged LSAs. This process completes the second step shown in the conceptual Figure 6-9, where the DR, acting like the pseudonode, floods the LSAs to the other OSPF routers in the subnet.

The process occurs in the background and can be generally ignored. However, for operating an OSPF network, an important distinction must be made. With a DR in existence, a DROther router performs the database exchange process (as seen in Figure 6-9) with the DR/BDR only and not any other DROther routers in the subnet. For example, in Figure 6-9, R1 acts as DR, R2 acts as BDR, and R3/R4 act as DROther routers. Because the underlying process does not make R3 and R4 perform database exchange with each other, the routers do not reach the FULL neighbor state, remaining in 2-Way state.

Example 6-6 shows the resulting output for the LAN shown in Figure 6-9, with four routers. The output, taken from DROther R3, shows a 2-Way state with R4, the other DROther. It also shows on interface Fa0/0 that its own priority is 2. This output also shows a neighbor count (all neighbors) of 3 and an adjacent neighbor count (all fully adjacent neighbors) of 2, again because the neighborhood between DROthers R3 and R4 is not a full adjacency.

Example 6-6 *Demonstrating OSPF FULL and 2-Way Adjacencies*

```
R3#show ip ospf interface fa0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 172.16.1.3/24, Area 0
  Process ID 75, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DROTHER, Priority 2
  Designated Router (ID) 1.1.1.1, Interface address 172.16.1.1
  Backup Designated router (ID) 2.2.2.2, Interface address 172.16.1.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:02
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 1/1, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 4
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 3, Adjacent neighbor count is 2
    Adjacent with neighbor 1.1.1.1 (Designated Router)
    Adjacent with neighbor 2.2.2.2 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

```
R3#show ip ospf neighbor fa0/0
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	4	FULL/DR	00:00:37	172.16.1.1	FastEthernet0/0
2.2.2.2	3	FULL/BDR	00:00:37	172.16.1.2	FastEthernet0/0
44.44.44.44	1	2WAY/DROTHER	00:00:36	172.16.1.4	FastEthernet0/0

Flooding Throughout the Area

So far in this section, the database exchange process has focused on exchanging the database between neighbors. However, LSAs need to be flooded throughout an area. To do so, when a router learns new LSAs from one neighbor, that router then knows that its other neighbors in that same area may not know of that LSA. Similarly, when an LSA changes, for example, when an interface changes state, a router may learn the same old LSA but with a new sequence number, and again need to flood the changed LSA to other neighbors in that area.

Figure 6-10 shows a basic example of the process. In this case, R2, R3, and R4 have established neighbor relationships, with four LSAs in their LSDB in this area. R1 is again the new router added to the internetwork.




Figure 6-10 Flooding Throughout an Area

First, consider what happens as the new R1-R2 neighborhood comes up and goes through database exchange. When R1 loads, and the link comes up, R1 and R2 reach a full state and have a shared view of the Area 1 LSDB. R2 has learned all R1's new LSAs (should only be R1's Type 1 router LSA), and R1 has learned all the area 1 LSAs known to R2, including the Type 1 LSAs for R3 and R4.

Next, think about the LSDBs of R3 and R4 at this point. The database exchange between R1-R2 did not inform R3 nor R4 about any of the new LSAs known by R1. So, R2, when it learns of R1's Type 1 LSA, sends DD packets to the DR on the R2/R3/R4 LAN. LSR/LSU packets follow, resulting in R3 and R4 learning about the new LSA for R1. If more routers existed in area 1, the flooding process would continue throughout the entire area, until all routers know of the best (highest sequence number) copy of each LSA.

The flooding process prevents the looping of LSAs as a side-effect of the database exchange process. Neighbors use DD messages to learn the LSA headers known by the neighbor, and then only request the LSAs known by the neighbor but not known by the local router. By requesting only unknown LSAs or new versions of old LSAs, routers prevent the LSAs advertisements from looping.

Periodic Flooding

Although OSPF does not send routing updates on a periodic interval, as do distance vector protocols, OSPF does re-flood each LSA every 30 minutes based on each LSA's age variable. The router that creates the LSA sets this age to 0 (seconds). Each router then increments the age of its copy of each LSA over time. If 30 minutes pass with no changes to an LSA—meaning no other reason existed in that 30 minutes to cause a re-flooding of the LSA—the owning router increments the sequence number, resets the timer to 0, and re-floods the LSA.

Because the owning router increments the sequence number and resets the LSAge every 1800 seconds (30 minutes), the output of various **show ip ospf database** commands should also show an age of less than 1800 seconds. For example, referring back to Example 6-5, the Type 1 LSA for R1 (RID 1.1.1.1) shows an age of 943 seconds and a sequence number of 0x80000003. Over time the sequence number should increment once per every 30 minutes, with the LSAge cycle upward toward 1800 and then back to 0 when the LSA is re-flooded.

Note also that when a router realizes it needs to flush an LSA from the LSDB for an area, it actually sets the age of the LSA to the MaxAge setting (3600) and re-floods the LSA. All the other routers receive the LSA, see the age is already at the maximum, causing those routers to also remove the LSA from their LSDBs.

Choosing the Best OSPF Routes

All this effort to define LSA types, create areas, and fully flood the LSAs has one goal in mind: to allow all routers in that area to calculate the best, loop-free routes for all known subnets. Although the database exchange process may seem laborious, the process by which SPF calculates the best routes requires a little less thought, at least to the level re-

quired for the CCNP ROUTE exam. In fact, the choice of the best route for a given subnet, and calculated by a particular router, can be summarized as follows:

- Analyze the LSDB to find all possible routes to reach the subnet.
- For each possible route, add the OSPF interface cost for all outgoing interfaces in that route.
- Pick the route with the lowest total cost.



For humans, if you build a network diagram, note the OSPF cost for each interface (as shown with `show ip ospf interface`), you can easily add up the costs for each router's possible routes to each subnet and tell which route OSPF will choose. The routers must use a more complex SPF algorithm to derive a mathematical model of the topology based on the LSAs. This section examines both the simpler human view of metric calculation and folds in some of the basics of what SPF must do on a router to calculate the best routes. It also goes through the options for tuning the metric calculation to influence the choice of routes.

OSPF Metric Calculation for Internal OSPF Routes

The process of calculating the cost from a router to each subnet may be intuitive to most people. However, spending a few minutes considering the details is worthwhile, in part to link the concepts with the LSAs, and to be better prepared for questions on the ROUTE exam. This section breaks the discussion into three sections: intra-area routes, interarea routes, a short discussion about cases when both intra-area and interarea routes exist for the same subnet, and an explanation of SPF calculations.

Calculating the Cost of Intra-Area Routes

When a router analyzes the LSDB to calculate the best route to each subnet, it does the following:

- Step 1.** Finds all subnets inside the area, based on the stub interfaces listed in the Type 1 LSAs and based on any Type 2 network LSAs
- Step 2.** Runs SPF to find all possible paths through the area's topology, from itself to each subnet
- Step 3.** Calculates the OSPF interface costs for all outgoing interfaces in each route, picking the lowest total cost route for each subnet as the best route



For example, Figure 6-11 shows the routers and links inside area 34, as a subset of the inter-network also shown in Figure 6-1. Figure 6-11 shows the interface numbers and OSPF costs.

Following the basic three-step process, at Step 1, R1 can determine that subnet 10.10.34.0/24 exists in area 34 because of the Type 2 LSA created by the DR in that subnet. For Step 2, R1 can then run SPF and determine four possible routes, two of which are clearly more reasonable to humans: R1-R3 and R1-R4. (The two other possible routes, R1-R3-R2-R4 and R1-R4-R2-R3, are possible and would be considered by OSPF but would clearly be higher cost.) For Step 3, R1 does the simple math of adding the costs of the outgoing interfaces in each route, as follows:

- R1-R3: Add R1's S0/0/0.3 cost (647) and R3's Fa0/0 cost (10), total 657
- R1-R4: Add R1's S0/0/0.4 cost (647) and R4's Fa0/0 cost (10), total 657

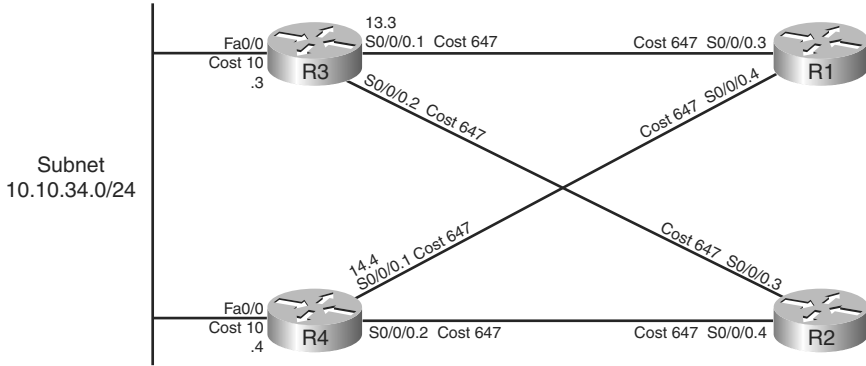


Figure 6-11 Area 34 Portion of Figure 6-1

The metrics tie, so with a default setting of **maximum-paths 4**, R1 adds both routes to its routing table. In particular, the routes list the metric of 657, and the next-hop IP address on the other end of the respectively links: 10.10.13.3 (R3's S0/0/0.1) and 10.10.14.4 (R4's S0/0/0.1).

Note that OSPF supports equal-cost load balancing, but it does not support unequal-cost load balancing. The **maximum-paths** OSPF subcommand can be set as low as 1, with the maximum being dependent on router platform and IOS version. Modern IOS versions typically support 16 or 32 concurrent routes to one destination (maximum).

Calculating the Cost of Interarea Routes

From a human perspective, the cost for interarea routes can be calculated just like for intra-area routes if we have the full network diagram, subnet numbers, and OSPF interface costs. To do so, just find all possible routes from a router to the destination subnet, add up the costs of the outgoing interfaces, and choose the router with the lowest total cost.

However, OSPF routers cannot do the equivalent for interarea routes, because routers internal to one area do not have topological data—LSA Types 1 and 2—for other areas. Instead, ABRs create and flood Type 3 summary LSAs into an area, listing the subnet number and mask, but not listing details about routers and links in the other areas. For example, Figure 6-12 shows both Areas 34 and 0 from Figure 6-1. Then consider how OSPF determines the lowest-cost route from router R3 for subnet 10.10.99.0/24, the Center subnet on the right.

R3 has a large number of possible routes to reach subnet 10.10.99.0/24. For example, just to get from R3 to R1, there are several possibilities: R3-R1, R3-R4-R1, and R3-R2-R1. From R1 the rest of the way to subnet 10.10.99.0/24, many more possibilities exist. The SPF algorithm has to calculate all possible routes inside an area to the ABR, so with more redundancy, SPF's run time goes up. And SPF has to consider all the options, whereas we humans can rule out some routes quickly because they appear to be somewhat ridiculous.

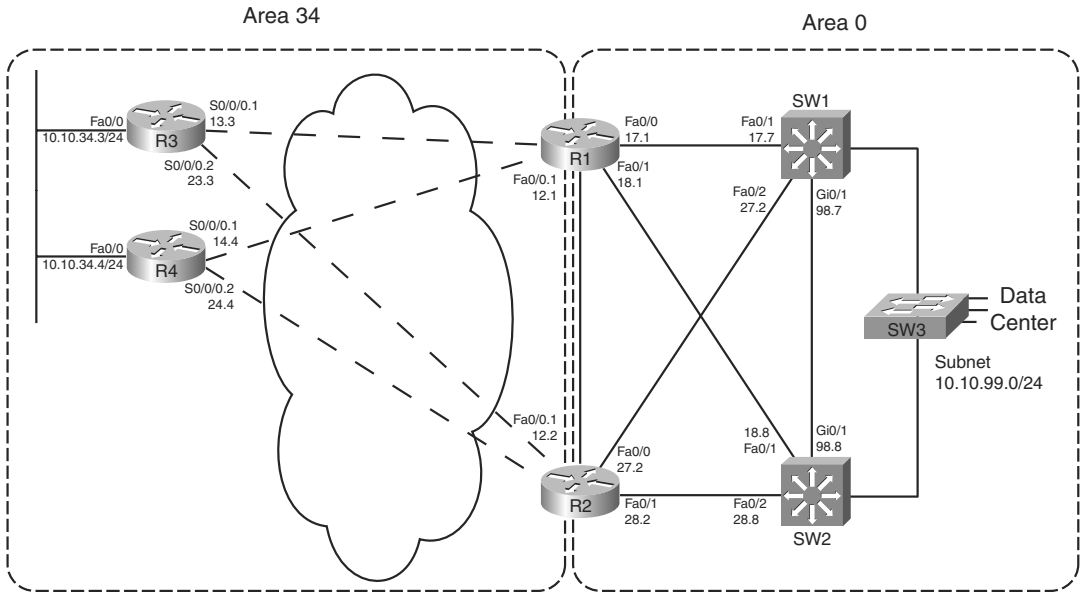


Figure 6-12 Area 34 and Area 0 Portion of Figure 6-1

Because of the area design, with R1 and R2 acting as ABRs, R3 does not process all the topology shown in Figure 6-12. Instead, R3 relies on the Type 3 Summary LSAs created by the ABRs, which have the following information:

- The subnet number/mask represented by the LSA
- The cost of the ABR's lowest-cost route to reach the subnet
- The RID of the ABR

Example 6-7 begins to examine the information R3 will use to calculate its best route for subnet 10.10.99.0/24, on the right side of Figure 6-12. To see these details, Example 6-7 lists several commands taken from R1. It lists R1's best route (actually two that tie) for subnet 10.10.99.0/24, with cost 11. It also lists the Type 3 LSA R1 generated by R1 for 10.10.99.0/24, again listing cost 11, and listing the Type 3 LSA created by ABR R2 and flooded into area 34.

Example 6-7 Route and Type 3 LSA on R1 for 10.10.99.0/24

```
R1#show ip route ospf
 10.0.0.0/8 is variably subnetted, 15 subnets, 3 masks
0    10.10.5.0/27 [110/648] via 10.10.15.5, 00:04:19, Serial10/0/0.5
0    10.10.23.0/29 [110/711] via 10.10.13.3, 00:04:19, Serial10/0/0.3
```

```

0      10.10.24.0/29 [110/711] via 10.10.14.4, 00:04:19, Serial0/0/0.4
0      10.10.25.0/29 [110/711] via 10.10.15.5, 00:04:19, Serial0/0/0.5
0      10.10.27.0/24 [110/11] via 10.10.17.7, 00:04:19, FastEthernet0/0
      [110/11] via 10.10.12.2, 00:04:19, FastEthernet0/0.1
0      10.10.28.0/24 [110/11] via 10.10.18.8, 00:04:19, FastEthernet0/1
      [110/11] via 10.10.12.2, 00:04:19, FastEthernet0/0.1
0      10.10.34.0/24 [110/648] via 10.10.14.4, 00:04:19, Serial0/0/0.4
      [110/648] via 10.10.13.3, 00:04:19, Serial0/0/0.3
0      10.10.98.0/24 [110/11] via 10.10.18.8, 00:04:19, FastEthernet0/1
      [110/11] via 10.10.17.7, 00:04:19, FastEthernet0/0
0      10.10.99.0/24 [110/11] via 10.10.18.8, 00:04:19, FastEthernet0/1
      [110/11] via 10.10.17.7, 00:04:19, FastEthernet0/0

```

```
R1#show ip ospf database summary 10.10.99.0
```

```
OSPF Router with ID (1.1.1.1) (Process ID 1)
```

```
! omitting output for area 5...
```

```
Summary Net Link States (Area 34)
```

```
LS age: 216
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 10.10.99.0 (summary Network Number)
```

```
Advertising Router: 1.1.1.1
```

```
LS Seq Number: 80000003
```

```
Checksum: 0x951F
```

```
Length: 28
```

```
Network Mask: /24
```

```
TOS: 0 Metric: 11
```

```
LS age: 87
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 10.10.99.0 (summary Network Number)
```

```
Advertising Router: 2.2.2.2
```

```
LS Seq Number: 80000002
```

```
Checksum: 0x7938
```

```
Length: 28
```

```
Network Mask: /24
```

```
TOS: 0 Metric: 11
```

Note: The examples use default bandwidth settings, but with all routers configured with the `auto-cost reference-bandwidth 1000` command. This command is explained in the upcoming section “Reference Bandwidth.”

For routers in one area to calculate the cost of an interarea route, the process is simple when you realize that the Type 3 LSA lists the ABR’s best cost to reach that interarea subnet. To calculate the cost:

- Step 1.** Calculate the intra-area cost from that router to the ABR listed in the type 3 LSA.
- Step 2.** Add the cost value listed in the Type 3 LSA. (This cost represents the cost from the ABR to the destination subnet.)



A router applies these two steps for each possible route to reach the ABR. Following the example of router R3 and subnet 10.10.99.0/24, Figure 6-13 shows the components of the calculation.

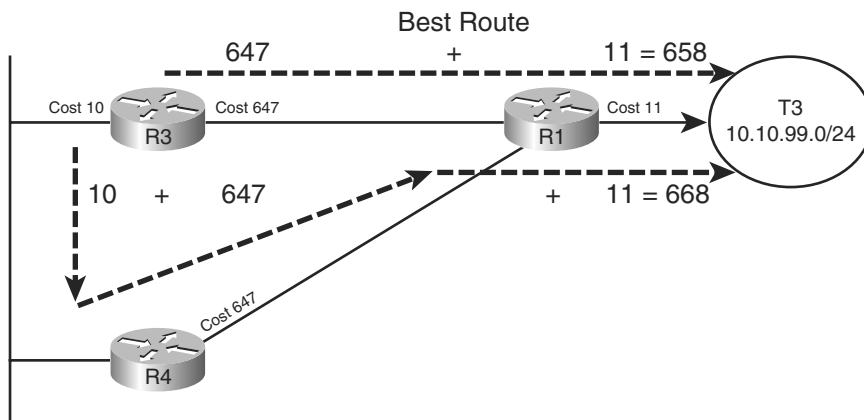


Figure 6-13 R3’s Calculation of Cost for 10.10.99.0/24

Figure 6-13 shows the calculation of both routes, with intra-area cost to reach R1 either 647 or 657 in this case. For both routes, the cost listed in the Type 3 LSA sourced by R1, cost 11, is added.

When more than one ABR exists, as is the case as shown in Figure 6-12, each ABR should have created a Type 3 LSA for the subnet. In fact, the output in Example 6-7 showed the Type 3 LSA for 10.10.99.0/24 created by both R1 and another created by R2. For instance, in the internetwork used throughout this chapter, ABRs R1 and R2 would create a Type 3 LSA for 10.10.99.0/24. So, in this particular example, R3 would also have to calculate the best route to reach 10.10.99.0/24 through ABR R2. Then, R3 would choose the best route among all routes for 10.10.99.0/24.

Each router repeats this process for all known routes to reach the ABR, considering the Type 3 LSAs from each ABR. In this case, R3 ties on metrics for one route through R1 and one through R2, so R3 adds both routes to its routing table, as shown in Example 6-8.

Example 6-8 *Route and Type 3 LSA on R3 for 10.10.99.0/24*

```
R3#show ip route 10.10.99.0 255.255.255.0
Routing entry for 10.10.99.0/24
  Known via "ospf 3", distance 110, metric 658, type inter area
  Last update from 10.10.13.1 on Serial0/0/0.1, 00:08:06 ago
  Routing Descriptor Blocks:
  * 10.10.23.2, from 2.2.2.2, 00:08:06 ago, via Serial0/0/0.2
    Route metric is 658, traffic share count is 1
  10.10.13.1, from 1.1.1.1, 00:08:06 ago, via Serial0/0/0.1
    Route metric is 658, traffic share count is 1

R3#show ip route ospf
  10.0.0.0/8 is variably subnetted, 15 subnets, 3 masks
O IA   10.10.5.0/27 [110/1304] via 10.10.23.2, 00:07:57, Serial0/0/0.2
        [110/1304] via 10.10.13.1, 00:07:57, Serial0/0/0.1
O IA   10.10.12.0/24 [110/657] via 10.10.23.2, 00:08:17, Serial0/0/0.2
        [110/657] via 10.10.13.1, 00:08:17, Serial0/0/0.1
! lines omitted for brevity
O IA   10.10.99.0/24 [110/658] via 10.10.23.2, 00:08:17, Serial0/0/0.2
        [110/658] via 10.10.13.1, 00:08:17, Serial0/0/0.1
```

Besides the information that matches the expected outgoing interfaces per the figures, the output also flags these routes as interarea routes. The first command lists “type inter area” explicitly, and the **show ip route ospf** command lists the same information with the code “O IA,” meaning OSPF, interarea. Simply put, interarea routes are routes for which the subnet is known from a Type 3 summary LSA.

Special Rules Concerning Intra-area and Interarea Routes on ABRs

OSPF has a couple of rules concerning intra-area and interarea routes that take precedence over the simple comparison of the cost calculated for the various routes. The issue exists when more than one ABR connects to the same two areas. Many designs use two routers between the backbone and each nonbackbone area for redundancy, so this design occurs in many OSPF networks.

The issue relates to the fact that with two or more ABRs, the ABRs themselves, when calculating their own routing tables, can calculate both an intra-area route and interarea route for subnets in the backbone area. For example, consider the perspective of Router R1 from the last several examples, as depicted in Figure 6-14.

Conceptually, R1 could calculate both the intra-area route and interarea route to 10.10.99.0/24. However, the OSPF cost settings could be set so that the lower cost route




Figure 6-14 *R1's Choice: Intra-Area or Interarea Route to 10.10.99.0/24*

for R1 actually goes through area 34, to ABR R2, and then on through Area 0 to 10.10.99.0/24. However, two OSPF rules prevent such a choice by R1:

- Step 1.** When choosing the best route, an intra-area route is always better than a competing interarea route, regardless of metric.
- Step 2.** If an ABR learns a Type 3 LSA inside a nonbackbone area, the ABR ignores that LSA when calculating its own routes.

Because of the first rule, R1 would never choose the interarea route if the intra-area route were available. The second rule goes further, stating that R1 could never choose the interarea route at all—R1 simply ignores that LSA for the purposes of choosing its own best IP routes.

Metric and SPF Calculations

Before moving on to discuss how to influence route choices by changing the OSPF interface costs, first take a moment to consider the CPU-intensive SPF work done by a router. SPF does the work to piece together topology information to find all possible routes to a destination. As a result, SPF must execute when the intra-area topology changes, because changes in topology impact the choice of best route. However, changes to Type 3 LSAs do not drive a recalculation of the SPF algorithm, because the Type 3 LSAs do not actually describe the topology.

To take the analysis a little deeper, remember that an internal router, when finding the best interarea route for a subnet, uses the intra-area topology to calculate the cost to reach the ABR. When each route is identified, the internal router adds the intra-area cost to the ABR, plus the corresponding Type 3 LSA's cost. A change to the Type 3 LSA—it fails, comes back up, or the metric changes—does impact the choice of best route, so the changed Type 3 LSA must be flooded. However, no matter the change, the change does not affect the topology between a router and the ABR—and SPF focuses on processing that topology data. So, only changes to Type 1 and 2 LSAs require an SPF calculation.

You can see the number of SPF runs, and the elapsed time since the last SPF run, using several variations on the **show ip ospf** command. Each time a Type 3 LSA changes and is flooded, SPF does not run, and the counter does not increment. However, each time a Type 1 or 2 LSA changes, SPF runs, and the counter increments. Example 6-9 highlights the counter that shows the number of SPF runs on that router, in that area, and the time since the last run. Note that ABRs list a group of messages per area, showing the number of runs per area.

Example 6-9 *Example with New Route Choices but No SPF Run*

```
R3#show ip ospf | begin Area 34
Area 34
  Number of interfaces in this area is 3
  Area has no authentication
  SPF algorithm last executed 00:41:02.812 ago
  SPF algorithm executed 15 times
  Area ranges are
  Number of LSA 25. Checksum Sum 0x0BAC6B
  Number of opaque link LSA 0. Checksum Sum 0x000000
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
```

Metric Tuning

Engineers have a couple of commands available that allow them to tune the values of the OSPF interface cost, thereby influencing the choice of best OSPF route. This section discusses the three methods: changing the reference bandwidth, setting the interface bandwidth, and setting the OSPF cost directly.

Changing the Reference Bandwidth

OSPF calculates the default OSPF cost for an interface based on the following formula:

$$\frac{\text{Reference-bandwidth}}{\text{interface-bandwidth}}$$

The reference-bandwidth, which you can set using the **auto-cost reference-bandwidth bandwidth** router subcommand, sets the numerator of the formula for that one router, with a unit of Mbps. This setting may be different on different routers, but Cisco recommends using the same setting on all routers in an OSPF routing domain.

For example, serial interfaces default to a bandwidth setting of 1544, meaning 1544 Kbps. The reference bandwidth defaults to 100, meaning 100 Mbps. After converting the reference bandwidth units to Kbps to match the bandwidth, the cost, calculated per the defaults, for serial links would be

$$\frac{100,000}{1544} = 64$$

Note: OSPF always rounds down when the calculation results in a fraction.

The primary motivation for changing the reference bandwidth is to accommodate good defaults for higher-speed links. With a default of 100 Mbps, the cost of FastEthernet interfaces calculates to cost 1. However, the minimum OSPF cost is 1, so Gigabit Ethernet and 10 Gigabit interfaces also then default to OSPF cost 1. By setting the OSPF reference bandwidth so that there is some difference in cost between the higher speed links, OSPF can then choose routes that use those higher speed interfaces.

Note: Although Cisco recommends that all routers use the same reference bandwidth, the setting is local to each router.

Note that in the examples earlier in this chapter, the bandwidth settings used default settings, but the **auto-cost reference-bandwidth 1000** command was used on each router to allow different costs for FastEthernet and Gigabit interfaces.

Setting Bandwidth

You can indirectly set the OSPF cost by configuring the **bandwidth speed** interface subcommand. In such cases, the formula shown in the previous section is used, just with the configured bandwidth value.

While on the topic of the interface bandwidth subcommand, a couple of seemingly trivial facts may matter to your choice of how to tune the OSPF cost. First, on serial links, the bandwidth defaults to 1544. On subinterfaces of those serial interfaces, the same bandwidth default is used.

On Ethernet interfaces, if not configured with the **bandwidth** command, the interface bandwidth matches the actual speed. For example, on an interface that supports autonegotiation for 10/100, the bandwidth is either 100,000 (kbps, or 100 Mbps) or 10,000 (Kbps, or 10 Mbps) depending on whether the link currently runs at 100 or 10 Mbps, respectively.

Configuring Cost Directly

The most controllable method to configure OSPF costs, but the most laborious, is to configure the interface cost directly. To do so, use the **ip ospf cost value** interface subcommand, substituting your chosen value as the last parameter.

Verifying OSPF Cost Settings

Several commands can be used to display the OSPF cost settings of various interfaces. Example 6-10 shows several, along with the configuration of all three methods for changing the OSPF cost. In this example, the following has been configured:

- The reference bandwidth is set to 1000.
- Interface S0/0/0.1 has its bandwidth set to 1000 Kbps.
- Interface Fa0/0 has its cost set directly to 17.

Example 6-10 *R3 with OSPF Cost Values Set*

```

router ospf 3
  auto-cost reference-bandwidth 1000
interface S0/0/0.1
  bandwidth 1000
interface fa0/0
  ip ospf cost 17

R3#show ip ospf interface brief

Interface      PID   Area      IP Address/Mask    Cost   State Nbrs F/C
Se0/0/0.2     3     34        10.10.23.3/29     647   P2P   1/1
Se0/0/0.1     3     34        10.10.13.3/29    1000  P2P   1/1
Fa0/0         3     34        10.10.34.3/24    17    BDR   1/1

R3#show ip ospf interface fa0/0
FastEthernet0/0 is up, line protocol is up
  Internet Address 10.10.34.3/24, Area 34
  Process ID 3, Router ID 3.3.3.3, Network Type BROADCAST, Cost: 17
  Enabled by interface config, including secondary ip addresses
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 4.4.4.4, Interface address 10.10.34.4
  Backup Designated router (ID) 3.3.3.3, Interface address 10.10.34.3
! lines omitted for brevity

```

Exam Preparation Tasks

Planning Practice

The CCNP ROUTE exam expects test takers to review design documents, create implementation plans, and create verification plans. This section provides some exercises that may help you to take a step back from the minute details of the topics in this chapter so that you can think about the same technical topics from the planning perspective.

For each planning practice table, simply complete the table. Note that any numbers in parentheses represent the number of options listed for each item in the solutions in Appendix F, “Completed Planning Practice Tables.”

Design Review Table

Table 6-6 lists several design goals related to this chapter. If these design goals were listed in a design document, and you had to take that document and develop an implementation plan, what implementation options come to mind? For any configuration items, a general description can be used, without concern about the specific parameters.

Table 6-6 *Design Review*

Design Goal	Possible Implementation Choices Covered in This Chapter
The design sets specific limits to the number of Type 1 and 2 LSAs in each area. Describe how to predict the number of each type of LSA.	
How could you tune OSPF metrics to favor 10 Gbps links over 1 Gbps and 1 Gig over 100 Mbps (2)?	
The design shows one physical path from ABR1 to core subnet 1 inside area 0, and one longer area 1 path to the same subnet. What can be done to ensure both paths can be used?	

Implementation Plan Peer Review Table

Table 6-7 shows a list of questions that others might ask, or that you might think about, during a peer review of another network engineer’s implementation plan. Complete the table by answering the questions.

Table 6-7 *Notable Questions from This Chapter to Consider During an Implementation Plan Peer Review*

Question	Answers
What conditions must be true for a router to create/flood a Type 2 LSA? (2)	
The plan shows Frame Relay with all point-to-point subinterfaces. By default, will a DR/BDR be elected?	
The plan shows a reference bandwidth change planned for all routers with high speed links, but not all other routers. What is the impact? (2)	
The plan shows many different WAN links speeds but with the interface bandwidths not matching the actual speed. All OSPF cost changes are made explicitly with the <code>ip ospf cost</code> interface subcommand. Do the incorrect bandwidths cause any OSPF problems?	

Create an Implementation Plan Table

To practice skills useful when creating your own OSPF implementation plan, list in Table 6-8 configuration commands related to the configuration of the following features. You may want to record your answers outside the book and set a goal to complete this table (and others like it) from memory during your final reviews before taking the exam.

Table 6-8 *Implementation Plan Configuration Memory Drill*

Feature	Configuration Commands/Notes
Tune metrics by changing the formula for calculating OSPF cost based on interface bandwidth.	
Tune metrics by changing interface bandwidth.	
Change metrics by setting cost directly.	
Set the number of equal-cost OSPF routes allowed in a router's routing table.	
Influence the choice of DR on a LAN. (2)	

Choose Commands for a Verification Plan Table

To practice skills useful when creating your own OSPF verification plan, list in Table 6-9 all commands that supply the requested information. You may want to record your answers outside the book and set a goal to complete this table (and others like it) from memory during your final reviews before taking the exam.

Table 6-9 *Verification Plan Memory Drill*

Information Needed	Command(s)
Display a summary of the OSPF database.	
Display all Type 1 router LSAs known to a router.	
Display the details of a particular Type 1 router LSA.	
Display all Type 2 network LSAs known to a router.	
Display the details of a particular Type 2 router LSA.	
Display all Type 3 summary LSAs known to a router.	
Display the details of a particular Type 3 router LSA.	
Display a list of OSPF-enabled interfaces on a router.	
Determine on which interfaces a router has formed at least one OSPF neighborship.	
Determine the number of fully adjacent neighbors on an interface.	
Determine which transit networks connect to a Type 1 LSA.	
Determine the router that created and flooded a Type 3 LSA.	
Determine the router that created and flooded a Type 2 LSA.	
Determine the router that created and flooded a Type 1 LSA.	
Display the IP address of the current DR and BDR on a LAN.	
Display the OSPF interface cost (metric).	
Display all OSPF-learned routes.	
Display statistics about the number of SPF algorithm runs.	

Note: Some of the entries in this table may not have been specifically mentioned in this chapter but are listed in this table for review and reference.

Review All the Key Topics



Review the most important topics from inside the chapter, noted with the Key Topics icon in the outer margin of the page. Table 6-10 lists a reference of these key topics and the page numbers on which each is found.

Table 6-10 *Key Topics for Chapter 6*

Key Topic Element	Description	Page Number
Table 6-2	OSPF LSA types	179
List	Two main functions of a DR	186
Table 6-3	Facts about LSA Types 1, 2, and 3	195
Table 6-4	OSPF Message Types	196
Table 6-5	OSPF neighbor states	197
List	Key differences between database exchange with and without a DR	201
List	Three steps a router considers when choosing the best OSPF IP routes	205
List	Three steps to calculate OSPF costs for intra-area routes	205
List	Two steps for calculating OSPF costs for interarea routes	209

Complete the Tables and Lists from Memory

Print a copy of Appendix D, “Memory Tables,” (found on the CD), or at least the section for this chapter, and complete the tables and lists from memory. Appendix E, “Memory Tables Answer Key,” also on the CD, includes completed tables and lists to check your work.

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary.

Link State Identifier (LSID), Designated Router (DR), Backup Designated Router (BDR), Internal Router, Area Border Router (ABR), All SPF Routers multicast, All DR multicast, Link State Advertisement, Database Description (DD) packet, Link State Request (LSR) packet, Link State Acknowledgement (LSA) packet, Link State Update (LSU) packet, Router LSA, Network LSA, Summary LSA, Type 1 LSA, Type 2 LSA, Type 3 LSA, Reference bandwidth, SPF calculation

This page intentionally left blank

This page intentionally left blank

Index

A

- ABR (area border routers), 143**
 - interarea routes, 210–211
 - intra-area routes, 210–211
 - OSPF
 - route filtering, 226–230*
 - route summarization, 232–235*
 - virtual links, 260–262*
- ACK (acknowledgement) messages**
 - EIGRP, 61, 65
 - LSAck messages, 200
- ACL (access control lists)**
 - EIGRP route filtering, 102–105
 - route maps, referencing ACL from, 112
- active timers, EIGRP convergence process, 87**
- AD (Administrative Distance), route redistribution with multiple redistribution points**
 - preventing domain loops via per-route AD settings, 350–354
 - preventing domain loops with AD, 346–349
- AD (Advertised Distance), building EIGRP IP routing tables, 69–72**
- administrative Weight (BGP Path Control), influencing, 500–501**
- aggregation, Internet routing, 392–393**
- areas, 143**
 - interarea routes
 - ABR, 210–211*
 - OSPF route selection, 206–210*
 - intra-area routes
 - ABR, 210–211*
 - OSPF route selection, 205–206*
 - OSPF database exchange process, flooding throughout an area, 203–204
- AS External LSA (link state advertisements), 180**
- AS_Path**
 - AS_Path Prepend command
 - BGP Path Control, 517–519*
 - increasing length via, 517–519*
 - BGP PA, 494
 - Weight (BGP Path Control), influencing, 508–511
- AS_SEQ path attribute, BGP ASN, 397–399**
- ASBR (Autonomous System Border Routers)**
 - ASBR Summary LSA (link state advertisements), 179
 - OSPF route summarization, 235–236
- ASN (autonomous system numbers), BGP ASN**
 - AS_SEQ path attribute, 397–399
 - public/private ASN, 400–402

authentication

eBGP neighborships, 429–430

EIGRP

configuration checklist, 40

*configuration example,
41–42*

*controlling EIGRP neighbor-
ships, 39–43*

*key chain time-based logic,
40–41*

verification example, 42–43

OSPF, 159–161

OSPF virtual links

*configuring authentication,
265–267*

*configuring without authentica-
tion, 262–264*

autoconfiguration (stateless)

global unicast addressing

*calculating Interface ID via
EUI-64, 547–548*

*finding DNS IP addresses via
stateless DHCP, 548–549*

*NDP router advertisements,
546–547*

IPv6 addressing, 561–562

automatic 6to4 tunneling, 627–634

auto-summary, network command and BGP route injections, 445

B

backbone routers, 143

bandwidth

EIGRP IP routing tables, configuring
for, 72–73

OSPF route selection

*changing reference bandwidth,
212–213*

setting bandwidth, 213

WAN bandwidth control, EIGRP
topology tables, 67–69

BDR (backup designated routers),
143

BEE (Boson Exam Environment),
CCNP Route exam preparation,
673, 678–679

best path algorithm, BGP Path
Control, 494–497

BFD (Bi-directional Forwarding
Detection feature), EIGRP, 33

BGP (Border Gateway Protocol),
387

administrative control of neighbor
status, 434–435

ASN

*AS_SEQ path attribute,
397–399*

public/private ASN, 400–402

BGP-4 (RFC 4760), 569

- eBGP, 399–400, 419
 - authentication*, 429–430
 - BGP table verification*, 436–443
 - multihop concepts in neighborships*, 428
 - neighbor configurations*, 423–424
 - neighborship requirements*, 425–426
 - Path Control*, 498
 - public IP address advertisements*, 443–448
 - redundancy configuration*, 429–430
 - redundancy issues between neighbors*, 426–428
 - verification commands for eBGP-learned routes*, 441–442
- EIGRP versus, 396–397
- iBGP, 399–400, 419, 455
 - avoiding routing loops*, 471–476
 - BGP Path Control*, 498
 - BGP table entries*, 464–468
 - configuring*, 460–463
 - between Internet-connected routers*, 459–471
 - next-hop reachability issues*, 468–471
 - routing loops, avoiding*, 471–476
 - verifying*, 463–464
- injecting routes
 - network command*, 443–445
 - route redistribution*, 446–448
- message type table, 436
- neighborships
 - administrative control of neighbor status*, 434–435
 - clearing*, 481–483
 - neighbor state reference table*, 430–431
 - verifying neighbors*, 430–434
- OSPF versus, 396–397
- outbound routing to the Internet
 - BGP versus*, 402–404
 - full BGP updates*, 410–411
 - partial BGP updates*, 410–411
 - path selection via BGP in dual-homed Internet designs*, 407–410
- PA
 - AS_Path*, 494
 - next-hop IP addresses*, 494
 - PA table*, 495
- Path Control, 491
 - best path algorithm*, 494–497
 - best path steps*, 498
 - eBGP*, 498
 - iBGP*, 498
 - increasing AS_Path length via AS_Path Prepend command*, 517–519
 - influencing inbound routes via MED*, 519–522
 - influencing outbound routes*, 500–519
 - influencing Weight*, 500–513
 - IP routes based on best paths*, 513–516
 - maximum-paths command*, 516
 - MED*, 519–522
 - memorization tips for best paths*, 499–500
 - Origin PA*, 498
 - PA*, 494–500
 - RIB failures*, 515–516

public IP address assignments
network command, 443–445
route redistribution, 446–448

route filtering, 476–477
clearing neighborships,
 481–483
displaying results, 483–485
filtering based on prefix/length,
 478–481

synchronization, iBGP routing loops,
 475–476

table verification
BGP update messages, 436–437
examining BGP table components,
 438–440
 NLRI, 437
viewing BGP table subsets,
 440–443

verifying, 463–464

verifying neighbors, 430–434

binary-to-decimal conversion tables,
 702–703

branch routing, 647

broadband Internet access, 650–652
configuring DHCP servers,
 664
configuring DSL, 661–663
configuring NAT, 663–664
DSL concepts, 659–661

DHCP servers/clients, 652–653

dynamic routing over GRE tunnels,
 658–659

floating static routes, 658

IPSec tunnels, 654–655

medium/large branches, 657–658

office security, 653–654

small branches, 656–657

VPN configuration, 667–669

**broadband Internet access, branch
 routing**, 650–652, 659–661

DHCP servers, 664

DSL, 661–663

NAT, 663–664

C

CCNP Route exams

implementation planning

focus for CCNP plans, 15

structured methodologies of, 15

styles of plans, 13–14

typical elements of, 14

planning-related exam topics

choosing commands for verification plan tables, 13

design review tables, 12

exam topics not requiring CLI,
 4–5

implementation plan peer review tables, 12

implementation plan tables, 13

preparing for, 5, 10–11

preparing for

activating practice exams, 674

BEE, 673, 678–679

chapter-ending review tools,
 676

Cisco Learning Network, 675

downloading practice exams,
 674

exam engine, 673, 678–679

memory tables, 675

subnetting practice, 677–678

suggested plans for final review/study, 676–679

- relating exam topics to network engineer jobs
 - design planning*, 7
 - fictitious company/network staffing scenarios*, 6–7
 - implementation planning*, 7–10
 - summary of network engineer's role*, 10
 - verification planning*, 9
- updates, 705–706
- verification plans
 - styles of plans*, 13–14
 - typical elements of*, 16
- CD installation (CCNP Route exam preparation), 673–674
- CEF (Cisco Express Forwarding), 366
- chapter-ending review tools (CCNP Route exam preparation), 676
- Cisco Learning Network, 675
- convergence
 - EIGRP convergence process, 32–33
 - active timers*, 87
 - fast convergence to feasible successors*, 78–80
 - feasibility conditions*, 79
 - going active*, 83–88
 - optimizing*, 78–91
 - successor routes*, 78–80
 - unequal metric route load sharing*, 88–91
 - verification of feasible successors*, 80–83
 - OSPF convergence process, 153–156
- conversion tables
 - binary-to-decimal conversion tables, 702–703
 - hex-to-decimal conversion tables, 701

D

- DAD (Duplicate Address Detection), 555
- database exchanges, OSPF database exchange process
 - exchanges
 - DR exchanges*, 200–203
 - exchanges without a DR*, 197–198
 - LSA exchanges*, 200
 - flooding throughout an area, 203–204
 - neighbor LSDB descriptions, describing, 198–200
 - neighbor state reference table, 197
 - OSPF message types, 196–197
 - periodic flooding, 204
- Dead timers, OSPF neighborships, 153–156
- decimals
 - binary-to-decimal conversion tables, 702–703
 - hex-to-decimal conversion tables, 701
- default keywords, PBR logic ordering, 370–371
- default routing
 - EIGRP default routing
 - advertising static default routes*, 127–128
 - configuring default networks*, 128–131
 - Internet routers*, 126–127
 - OSPF default routing, 221, 236–239
- default-information originate command, OSPF default routing, 237–239
- delays, configuring for EIGRP IP routing tables, 72–73
- design planning (CCNP Route exams), 7
- design review tables (CCNP Route exams), 12

DHCP (Dynamic Host Configuration Protocol)

branch routing, 652–653, 664

stateful DHCP, global unicast addressing, 545

stateless autoconfiguration, 548–549

distributed lists, OSPF route filtering, 230–231

distribute-list command, route redistribution filtering, 343

DMVPN (Dynamic Multipoint Virtual Private Networks), 664

DNA (Do Not Age) bits (LSA), 262

DNS IP addresses, finding via stateless DHCP, 548–549

DR (designated routers), 143

Network LSA, 186–191

OSPF

database exchanges, 200–203

database exchanges without, 197–198

OSPF over multipoint Frame Relay, 271–272

DSL (Digital Subscriber Lines), branch router configuration for broadband access, 659–663

dual stacks (IPv4/IPv6), 611–612

dual-homed Internet design, outbound routing to the Internet

full BGP updates, 410–411

partial BGP updates, 410–411

path selection via BGP, 407–410

preferred path routing, 405–407

dual-multihomed Internet design, outbound routing to the Internet, 412–413

dynamic multipoint tunneling, 626

automatic 6to4 tunnels, 627–634

ISATAP tunneling, 634–639

dynamic routing, branch routing, 658–659

E

eBGP (external BGP), 399–400, 419

BGP Path Control, 498

BGP table verification, 436

neighborships

authentication, 429–430

multihop concepts, 428

neighbor configurations, 423–424

redundancy configuration, 429–430

redundancy issues between neighbors, 426–428

requirements for forming, 425–426

public IP address assignments

network command, 443–445

route redistribution, 446–448

verification commands for eBGP-learned routes, 441–442

EIGRP (Enhanced Interior Gateway Routing Protocol), 19, 569

ACK messages, 61, 65

authentication

configuration checklist, 40

configuration example, 41–42

controlling EIGRP neighborships, 39–43

key chain time-based logic, 40–41

verification example, 42–43

BFD feature, 33

BGP versus, 396–397

CCNA review

calculating best routes for routing tables, 30

- configuration review*, 23–25
- exchanging topology information*, 29–30
- internals review*, 29
- verification review*, 25–29
- convergence process, 32–33
 - active timers*, 87
 - fast convergence to feasible successors*, 78–80
 - feasibility conditions*, 79
 - going active*, 83–88
 - optimizing*, 78–91
 - successor routes*, 78–80
 - unequal metric route load sharing*, 88–91
 - verification of feasible successors*, 80–83
- default routing
 - advertising static default routes*, 127–128
 - configuring default networks*, 128–131
 - Internet routers*, 126–127
- EIGRP for IPv4 comparisons, 581–582
- EIGRP for IPv6
 - configuring*, 582–584
 - EIGRP for IPv4 comparisons*, 581–582
 - verifying*, 584–587
- feature summary table, 31
- IP routing tables, building
 - calculating FD/RD metrics*, 69–72
 - configuring bandwidth*, 72–75
 - configuring delays*, 72–73
 - configuring k-values*, 75–76
 - metric tuning*, 72–78
 - Offset Lists*, 76–78
- neighborships
 - configuration settings that prevent relationships*, 46–48
 - configuring Hello/Hold timers*, 33–34
 - configuring metric components via k-values*, 47–48
 - controlling via EIGRP authentication*, 39–43
 - controlling via static configurations*, 43–45
 - Frame Relay*, 49
 - manipulating Hello/Hold timers*, 32–33
 - MetroE*, 51
 - MPLS VPN*, 50
 - neighbor requirements*, 46, 152–153
 - neighborships over WAN*, 48–51
 - passive interface feature*, 36–39
 - verifying Hello/Hold timers*, 34–36
- RID, 48
- route filtering, 101
 - ACL references*, 102–105
 - IP prefix list references*, 105–110
 - route maps*, 110–114
- route redistribution
 - baseline configuration examples*, 298–299
 - configuring with default metric components*, 300–302
 - default AD defeats loop from EIGRP to OSPF to EIGRP*, 346–347
 - default AD defeats loop from OSPF to EIGRP to OSPF*, 346–347
 - redistribute command reference*, 297–298
 - verifying redistribution*, 302–305

- route summarization
 - auto-summary*, 124–126
 - benefits/trade-offs*, 120
 - calculating summary routes*, 116
 - choosing where to summarize routes*, 116–117
 - configuring*, 120–124
 - influencing summary route selection*, 117–118
 - suboptimal forwarding*, 118–120
 - summary route design*, 114–115
- topology tables, building
 - contents of update messages*, 61–64
 - seeding topology tables*, 60
 - Split Horizon*, 64
 - Split Horizon defaults on Frame Relay multipoint subinterfaces*, 65–67
 - update process*, 64–65
 - WAN bandwidth control*, 67–69
 - WAN issues for topology exchanges*, 65–69
- updates
 - update messages*, 30, 61–64
 - update process*, 64–65
- Ethernet, MetroE (Metropolitan Ethernet)
- EIGRP neighborships, 51
- OSPF neighborships, 167–169
- EUI-64, Interface ID calculation for global unicast addresses, 547–548
- exam engine (CCNP Route exam preparation), 673, 678–679
- exams (CCNP Route)
 - implementation planning
 - focus for CCNP plans*, 15
 - structured methodologies of*, 15
 - styles of plans*, 13–14
 - typical elements of*, 14
 - planning-related exam topics
 - choosing commands for verification plan tables*, 13
 - design review tables*, 12
 - exam topics not requiring CLI*, 4–5
 - implementation plan peer review tables*, 12
 - implementation plan tables*, 13
 - preparing for*, 5, 10–11
 - preparing for
 - activating practice exams*, 674
 - BEE*, 673, 678–679
 - chapter-ending review tools*, 675
 - Cisco Learning Network*, 675
 - downloading practice exams*, 674
 - exam engine*, 673, 678–679
 - memory tables*, 675
 - subnetting practice*, 677–678
 - suggested plans for final review/study*, 675–677
 - relating exam topics to network engineer jobs
 - design planning*, 7
 - fictitious company/network staffing scenarios*, 6–7
 - implementation planning*, 7–10
 - summary of network engineer's role*, 10
 - verification planning*, 9
 - updates, 705–706
 - verification plans
 - styles of plans*, 13–14
 - typical elements of*, 16
- explicit acknowledgements. See LSAck (Link State Acknowledgement) messages
- External Attributes LSA (link state advertisements), 180

F

FD (Feasible Distance), building EIGRP IP routing tables, 69–72

feasibility conditions, EIGRP convergence process, 79

feasible successors, verifying (EIGRP convergence process), 78–80

final review/study, suggested plans for, 676–679

floating static routes, branch routing, 658

flooding process, OSPF database exchange process, 203–204

Frame Relay**EIGRP**

neighborships, 49

topology tables, building, 65–67

IP subnetting design over Frame Relay, 267–268

OSPF neighborships, Frame Relay point-to-point subinterfaces, 166

OSPF over multipoint Frame Relay
configuring operations, 274–282
configuring using multipoint subinterfaces, 269–270
configuring using physical interfaces, 268–269

DR, 271–272

IP subnetting design over Frame Relay, 267–268

mapping issues with partial mesh topologies, 272–273

NBMA, 275–279

neighbor discovery, 270–271

network type point-to-multipoint, 279–281

static neighbor definition, 270–271

verifying operations, 274–282

G

GET VPN (Group Encrypted Transport Virtual Private Networks), 665

global unicast addressing, 533

assigning, 544

stateful DHCP, 545

stateful IPv6 address configuration, 549

stateless autoconfiguration, 545–549

automatic 6to4 tunnels, 632–634

global route aggregation, 534–536

prefix assignment example, 539–541

subnetting, 541–543

going active, EIGRP convergence process, 83

SIA routes, 87–88

stub router impact on Query Scope, 84–86

summary route impact on Query Scope, 86–87

GRE (Generic Route Encapsulation) tunneling, 619–620, 624–625

dynamic routing over GRE tunnels, branch routing, 658–659

VPN, configuring in, 666–667

Group Membership LSA (link state advertisements), 180

H

Hello messages, OSPF, 152

Hello timers**EIGRP**

configuring in, 33–34

manipulating in, 32–33

verifying in, 34–36

- OSPF neighborships, optimizing convergence via Hello timers, 153–156
 - hex-to-decimal conversion tables, 701
 - Hold timers, EIGRP**
 - configuring in, 33–34
 - manipulating in, 32–33
 - verifying in, 34–36
-
- iBGP (internal BGP), 399–400, 419, 455**
 - BGP Path Control, 498
 - BGP table entries, 464–468
 - configuring, 460–463
 - between Internet-connected routers, 459–471
 - next-hop reachability issues, 468
 - changing next-hop addresses via next-hop-self command, 469–471*
 - recursive route table lookups, 469*
 - routing loops, avoiding, 471
 - BGP synchronization, 475–476*
 - iBGP mesh topologies, 472–475*
 - IGP redistribution, 475–476*
 - verifying, 463–464
 - IGP (Interior Gateway Protocol) redistribution**
 - advanced IGP redistribution, 329
 - multiple redistribution points, 344–357*
 - route maps, 332–343*
 - basic IGP redistribution, 289
 - EIGRP route redistribution, 297–305*
 - need for route redistribution, 292–294*
 - OSPF route redistribution, 305–323*
 - redistribution concepts, 294–297*
 - redistribution processes, 294–297*
 - iBGP routing loops, avoiding, 475–476
 - IPv6 addressing, 595
 - redistributing with route maps, 598–599*
 - redistributing without route maps, 596–598*
 - implementation planning (CCNP Route exams), 7–9**
 - documenting, 10
 - focus for CCNP plans, 15
 - implementation plan tables, 13
 - peer review tables, 12
 - structured methodologies of, 15
 - styles of plans, 13–14
 - typical elements of, 14
 - interarea routes**
 - ABR (area border routers), 210–211
 - OSPF route selection, 206–210
 - internal routers, 143**
 - Internet routing, 387, 390**
 - aggregation, 392–393
 - EIGRP default routing, 126–127
 - NAT, 393
 - outbound routing to the Internet
 - BGP versus default routing, 402–404*
 - dual-homed Internet design, 405–411*
 - dual-multihomed Internet design, 412–413*
 - single-homed Internet design, 404–405*
 - single-multihomed Internet design, 411–412*
 - PAT, 393–394

- private IP address assignments, 394–395
- public IP address assignments, 391–392
- intra-area routes**
 - ABR, 210–211
 - OSPF route selection, 205–206
- IP addresses, Internet routing**
 - private IP address assignments, 394–395
 - public IP address assignments, 391–392
- IP MTU (maximum transmission unit)**
 - mismatches, OSPF neighborships, 157–159
- IP prefix lists**
 - concepts of, 105–107
 - EIGRP route filtering, 105–110
 - matching, samples of, 107–108
 - route maps, referencing prefix lists from, 112
- IP routing tables, building EIGRP IP routing tables**
 - calculating FD/RD metrics, 69–72
 - configuring
 - bandwidth*, 72–75
 - delays*, 72–73
 - k-values*, 75–76
 - metric tuning, 72–78
 - Offset Lists, 76–78
- IP SLA (Service Level Agreements), 363**
 - concepts of, 373–374
 - configuring, 374–376
 - PBR, 372, 381
 - tracking IP SLA operations to influence routing
 - configuring PBR to track IP SLA*, 381
 - static routes*, 378–381
 - verifying, 376–377
- IP subnetting design over Frame Relay, 267–268**
- IPSec tunnels, branch routing, 654–655**
- IPsec VPN (IP security virtual private networks), configuring, 665–666**
- IPv4 addresses**
 - EIGRP for IPv4, 581–582
 - IPv6 coexistence with, 607
 - dual stacks*, 611–612
 - NAT Protocol Translation*, 617–619
 - tunneling*, 612–617
- IPv6 addressing, 529, 532**
 - connected neighbors, 560
 - connected routes, 560
 - DAD, 555
 - EIGRP for IPv6
 - configuring*, 582–584
 - EIGRP for IPv4 comparisons*, 581–582
 - verifying*, 584–587
 - global unicast addressing, 533
 - assigning*, 544–549
 - global route aggregation*, 534–536
 - prefix assignment example*, 539–541
 - stateful DHCP*, 545
 - stateful IPv6 address configuration*, 549
 - stateless autoconfiguration*, 545–549
 - subnetting*, 541–543
- IGP redistribution, 595**
 - redistributing with route maps*, 598–599
 - redistributing without route maps*, 596–598

inverse neighbor discovery, 555–556

IPv4 coexistence with, 607

- dual stacks*, 611–612
- NAT Protocol Translation*, 617–619
- tunneling*, 612–617

Layer 2 mapping, neighbor address protocol, 554–555

multicast groups joined by IPv6 router interfaces, 559–560

multicast IPv6 addressing, 553–554

neighbor tables, 561

OSPFv3

- configuring*, 590–592
- OSPFv2 comparisons*, 588–589
- verifying*, 592–595

overview of, 550

prefixes

- conventions for writing*, 537–539
- terminology*, 543–544

representing IPv6 addresses, conventions for, 536–537

RIPng, 573

- configuring*, 575–578
- RIP-2 comparisons*, 574
- verifying*, 578–580

router configurations, 556–559

routing protocol updates, 573

stateless autoconfiguration, 561–562

static IPv6 addresses, router configurations, 557–559

static IPv6 routes, 599–601

tunneling

- general concepts*, 612–614
- GRE tunneling*, 619–620, 624–625
- MCT*, 619–624
- point-to-multipoint tunnels*, 615–616, 626–640

- point-to-point tunnels*, 614–615
- static point-to-point tunnels*, 619–626
- tunneling comparison table*, 617

unicast IPv6 addressing, 550, 553

- link local unicast addresses*, 551–552
- unique local addresses*, 551

ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) tunneling, 634–639

J - K - L

key chain time-based logic, EIGRP authentication, 40–41

k-values, configuring EIGRP metric components, 47–48, 75–76

LAN (local area networks), OSPF neighborships

- neighbor discovery, enabling*, 150–152
- LAN adjacencies*, 149–161

large branches (branch routing), 657–658

Layer 2 address mapping, neighbor address protocol, 554–555

link local unicast addresses, 551–552

load balancing, EIGRP convergence process, 88–91

Local_Pref, influencing Weight (BGP Path Control), 507–508

- Local_Pref Internetwork example, 508–511
- route maps, 511–513

LSA (link state advertisements), 143

- ASBR Summary LSA, 179

- DNA bits, 262
- External Attributes LSA, 180
- AS External LSA, 180
- Group Membership LSA, 180
- Net Summary LSA, 179
- Network LSA, 179, 196
 - concepts of*, 187
 - DR, 186
 - show commands*, 187–191
- NSSA External LSA, 180
- Opaque LSA, 180
- OSPF
 - database exchange process*, 200
 - route filtering*, 226–230
 - route redistribution*, 311–318
- OSPF LSDB, 179
 - limiting number of LSA*, 195
 - Network LSA*, 186–191
 - Router LSA*, 180–186
 - Summary LSA*, 191–195
- Router LSA, 179–186, 195
- Summary LSA, 191–196
- LSAck (Link State Acknowledgement) messages, 200
- LSDB (link state databases), OSPF
 - LSDB, 140–142
 - LSA types*, 179–180
 - neighbor LSDB descriptions*,
discovering, 198–200
 - Network LSA*, 186–191
 - number of LSA*, *limiting*, 195
 - Router LSA*, 180–186
 - Summary LSA*, 191–195
- LSID (link state identifiers), 180
- LSR (Link State Request) messages, 200
- LSU (Link State Updates), 142, 200

M

- maximum-paths command, BGP Path Control, 516
- MCT (manually configured tunnels), 619
 - configuring, 620–623
 - verifying, 623–624
- MED (Multi-Exit Discriminators)
 - concepts of, 519–520
 - configuring, 521–522
 - features of, 521
- medium/large branches (branch routing), 657–658
- memory tables, CCNP Route exam preparation, 675
- MetroE (Metropolitan Ethernet)
 - EIGRP neighborships, 51
 - OSPF neighborships, 167–169
- MPLS VPN (Multiprotocol Label Switching Virtual Private Networks)
 - EIGRP neighborships, 50
 - OSPF neighborships, 166–167
- MTU (maximum transmission units) mismatches, OSPF neighborships, 157–159
- multicast IPv6 addressing, 553–554
- multipoint Frame Relay, OSPF over
 - configuring
 - operations*, 274–282
 - via multipoint subinterfaces*,
269–270
 - via physical interfaces*, 268–269
 - DR, 271–272
 - IP subnetting design over Frame Relay,
267–268
 - mapping issues with partial mesh
topologies, 272–273
 - NBMA, 275–279

neighbor discovery, 270–271
 network type point-to-multipoint,
 279–281
 network type point-to-multipoint
 nonbroadcast, 281–282
 static neighbor definition, 270–271
 verifying operations, 274–282

N

N (BGP best path memorization tip),
 500

NAT (network address translation)

branch router configuration for broad-
 band access, 663–664

Internet routing, 393

NAT Protocol Translation, 617–619

NBMA (Network Type Nonbroadcast),
 OSPF over multipoint Frame Relay,
 275–279

NDP router advertisements, stateless
 autoconfiguration of global unicast
 addresses, 546–547

neighbor address protocol, Layer 2
 address mapping, 554–555

neighbor tables, IPv6 addressing, 561

neighbor weight command, influencing
 Weight (BGP Path Control),
 506–507

neighborships

EIGRP neighborships

configuration settings that
prevent relationships, 46–48

configuring Hello/Hold timers,
 33–34

configuring metric components
via k-values, 47–48

controlling via EIGRP authentica-
tion, 39–43

controlling via static configura-
tions, 43–45

Frame Relay, 49

manipulating Hello/Hold timers,
 32

MetroE, 51

MPLS VPN, 50

neighbor requirements, 46,
 152–153

neighborships over WAN,
 48–51

passive interface feature, 36–39
verifying Hello/Hold timers,

34–36

OSPF neighborships

authentication, 159–161

enabling neighbor discovery on
LAN, 150–152

Frame Relay point-to-point
subinterfaces, 166

LAN adjacencies, 149–161

MetroE, 167–169

MPLS VPN, 166–167

MTU mismatches, 157–159
neighbor requirements, 46,
 152–153

network types, 162–163

optimizing convergence via
Hello/Dead timers, 153–156

point-to-point links, 164–166

unique RID, 156–157

WAN adjacencies, 162–169

Net Summary LSA (link state advertise-
ments), 179

network command, BGP route injec-
tions, 443–445

network engineer jobs, relating CCNP
Route exam topics to

design planning, 7

- fictitious company/network staffing scenarios, 6–7
 - implementation planning, 7–10
 - summary of network engineer's role, 10
 - verification planning, 9
 - Network LSA (link state advertisements), 179, 196**
 - concepts of, 187
 - OSPF LSDB, DR, 186
 - show commands, 187–191
 - network type point-to-multipoint, OSPF over multipoint Frame Relay, 279–281**
 - network type point-to-multipoint non-broadcast, OSPF over multipoint Frame Relay, 281–282**
 - next-hop IP addresses, BGP PA, 494**
 - next-hop-self command, iBGP next-hop reachability issues, 469–471**
 - NLRI (Network Layer Reachability Information), BGP table verification, 437**
 - NSSA (not-so stubby areas), 240–241, 248–250**
 - NSSA External LSA (link state advertisements), 180
 - OSPF route redistribution, external routes in NSSA areas, 320–323
 - totally NSSA, 240–241
- O**
-
- office security, branch routing, 653–654**
 - Offset Lists, building EIGRP IP routing tables, 76–78**
 - OMNI (BGP best path memorization tip), 500**
 - Opaque LSA (link state advertisements), 180**
 - Origin PA (Path Attributes), BGP PA, 498**
 - OSPF (Open Shortest Path First), 137, 569**
 - authentication, 159–161
 - BGP versus, 396–397
 - commonly used terms table, 142–143
 - configuration review, 144–146
 - convergence, optimizing via Hello/Dead timers, 153–156
 - database exchange process
 - discovering neighbor LSDB descriptions, 198–200*
 - exchanges without a DR, 197–198*
 - exchanging LSA, 200*
 - exchanging with DR, 200–203*
 - flooding throughout an area, 203–204*
 - neighbor state reference table, 197*
 - OSPF message types, 196–197*
 - periodic flooding, 204*
 - default routing, 221, 236–239
 - feature summary table, 149
 - Hello messages, 152
 - LSDB, 140–142
 - limiting number of LSA, 195*
 - LSA types, 179–180*
 - Network LSA, 186–191*
 - Router LSA, 180–186*
 - Summary LSA, 191–195*
 - multipoint Frame Relay
 - configuring operations, 274–282*
 - configuring using multipoint subinterfaces, 269–270*
 - configuring using physical interfaces, 268–269*
 - DR, 271–272*
 - IP subnetting design over Frame Relay, 267–268*

- mapping issues with partial mesh topologies, 272–273*
- NBMA, 275–279
- neighbor discovery, 270–271*
- network type point-to-multipoint, 279–281*
- network type point-to-multipoint nonbroadcast, 281–282*
- static neighbor definition, 270–271*
- verifying operations, 274–282*
- neighborships
 - authentication, 159–161*
 - enabling neighbor discovery on LAN, 150–152*
 - Frame Relay point-to-point subinterfaces, 166*
 - LAN adjacencies, 149–161*
 - MetroE, 167–169*
 - MPLS VPN, 166–167*
 - MTU mismatches, 157–159*
 - neighbor requirements, 46, 152–153*
 - network types, 162–163*
 - optimizing convergence via Hello/Dead timers, 153–156*
 - point-to-point links, 164–166*
 - unique RID, 156–157*
 - WAN adjacencies, 162–169*
- network types, 274
- OSPFv3
 - configuring, 590–592*
 - OSPFv2 comparisons, 588–589*
 - verifying, 592–595*
- route filtering, 221, 225
 - filtering OSPF routes added to routing tables, 230–231*
 - filtering with distributed lists, 230–231*
 - Type 3 LSA filtering, 226–230*
- route redistribution
 - configuring with minimal parameters, 306–310*
 - default AD defeats loop from EIGRP to OSPF to EIGRP, 347–348*
 - default AD defeats loop from OSPF to EIGRP to OSPF, 347–348*
 - E1/E2 route comparisons, 319–320*
 - external routes in NSSA areas, 320–323*
 - external type 2 route LSA, 311–318*
 - external type 2 route metrics, 311–318*
 - redistribute command reference, 305–306*
 - redistributing into OSPF as E1 routes, 318–319*
 - setting OSPF metrics, 310–311*
- route selection, 204
 - calculating interarea route costs, 206–210*
 - calculating intra-area route costs, 205–206*
 - changing reference bandwidth, 212–213*
 - configuring cost directly, 213*
 - metric calculation for internal routes, 205–211*
 - metric calculations, 211–212*
 - metric tuning, 212–214*
 - setting bandwidth, 213*
 - special rules for interarea routes, 210–211*
 - special rules for intra-area routes, 210–211*
 - SPF calculations, 211–212*
 - verifying cost settings, 213–214*

- route summarization, 221, 231
 - manual summarization at ABR, 232–235*
 - manual summarization at ASBR, 235–236*
- stub routers, 221, 239
- stubby areas, 239
 - configuring, 241–243*
 - NSSA, 240–241, 248–250*
 - summary table, 250*
 - totally NSSA, 240–241*
 - totally stubby areas, 240–241, 246–248*
 - types of, 240–241*
 - verifying, 243–246*
- verification review, 146–148
- virtual links
 - concepts of, 260–262*
 - configuring authentication, 265–267*
 - configuring without authentication, 262–264*
 - verifying, 264–265*

P

- partial mesh topologies, OSPF over multipoint Frame Relay mapping issues, 272–273
- passive interface feature, EIGRP neighborships, 36–39
- PAT (port address translation), Internet routing, 393–394
- PBR (Policy-Based Routing), 363, 366
 - configuring, 368–370
 - default keyword and logic ordering, 370–371
 - IP precedence, setting, 371–372
 - IP SLA, 372, 381
 - locally created packets, applying PBR to, 371
 - matching packets, 367–368
 - setting routes, 367–368
- periodic flooding, OSPF database exchange process, 204
- planning-related exam topics (CCNP Route exams)
 - choosing commands for verification plan tables, 13
 - design review tables, 12
 - exam topics not requiring CLI, 4–5
 - implementation plan peer review tables, 12
 - implementation plan tables, 13
 - preparing for, 5, 10–11
- point-to-multipoint tunneling, 615–616, 626
 - automatic 6to4 tunnels, 627–634
 - ISATAP tunneling, 634–639
- point-to-point links, OSPF neighborships, 164–166
- point-to-point tunneling, 614–615
 - GRE tunneling, 619–620, 624–625
 - MCT, 619–624
- practice exams (CCNP Route exams), downloading, 674
- prefix lists
 - concepts of, 105–107
 - EIGRP route filtering, 105–110
 - matching, samples of, 107–108
 - route maps, referencing prefix lists from, 112
- preparing for CCNP Route Exams
 - activating practice exams, 674
 - BEE, 673, 678–679
 - chapter-ending review tools, 675
 - Cisco Learning Network, 675

downloading practice exams, 674
 exam engine, 673, 678–679
 memory tables, 675
 subnetting practice, 677–678
 suggested plans for final review/study,
 675–677

PVC (permanent virtual circuits), EIGRP
 neighborships on Frame Relay, 49

Q - R

QoS (Quality of Service), PBR, 371–372

Query Scope, EIGRP convergence
 process

stub router impact on Query Scope,
 84–86

summary route impact on Query Scope,
 86–87

RD (Reported Distance), building EIGRP
 IP routing tables, 69–72

recursive route table lookups, iBGP next-
 hop reachability issues, 469

redistribute command reference

EIGRP route redistribution, 297–298

OSPF route redistribution, 305–306

reference bandwidth, OSPF route selec-
 tion, 212–213

review tools (chapter-ending), 676

reviews (final), suggested plans for,
 676–679

RIB failures, BGP Path Control,
 515–516

RID (Router ID)

EIGRP RID, 48

OSPF neighborships, unique RID,
 156–157

RIP-2 (Routing Information Protocol-2),
 RIPng comparisons, 574

RIPng (Routing Information Protocol
 next generation), 569, 573

configuring, 575–578

RIP-2 comparisons, 574

verifying, 578–580

route filtering

BGP route filtering, 476–477

clearing neighborships, 481–483

displaying results, 483–485

*filtering based on prefix/length,
 478–481*

EIGRP route filtering, 101

ACL references, 102–105

IP prefix list references, 105–110

route maps, 110–114

OSPF route filtering, 221, 225

*filtering OSPF routes added to
 routing tables, 230–231*

*filtering with distributed lists,
 230–231*

Type 3 LSA filtering, 226–230

route maps

ACL references, 112

concepts of, 111–112

EIGRP route filtering, 110–114

IGP redistribution in IPv6 addressing
*redistributing with route maps,
 598–599*

*redistributing without route
 maps, 596–598*

prefix list references, 112

route redistribution, 332–333

*configuring metric settings,
 339–341*

*filtering redistributed routes,
 334–339*

- redistribution filtering via distribute-list command*, 343
 - verifying metric settings*, 341–342
- Weight (BGP Path Control), influencing, 504–506, 511–513
- route redistribution**
 - advanced IGP redistribution, 329
 - multiple redistribution points*, 344–357
 - route maps*, 332–343
 - basic IGP redistribution, 289
 - EIGRP route redistribution*, 297–305
 - need for route redistribution*, 292–294
 - OSPF route redistribution*, 305–323
 - redistribution concepts*, 294–297
 - redistribution processes*, 294–297
 - BGP route injections, 446–448
 - EIGRP route redistribution
 - baseline configuration examples*, 298–299
 - configuring with default metric components*, 300–302
 - redistribute command reference*, 297–298
 - verifying redistribution*, 302–305
 - multiple redistribution points, 344
 - domain loop problems with multiple routing domains*, 349–357
 - preventing domain loops via per-route AD settings*, 350–354
 - preventing domain loops via route-tag filtering using distribute lists*, 355–357
 - preventing domain loops via subnet filtering while redistributing*, 354–355
 - preventing domain loops with AD*, 346–349
 - preventing routing domain loops with higher metrics*, 345
- OSPF route redistribution
 - configuring with minimal parameters*, 306–310
 - E1/E2 route comparisons*, 319–320
 - external routes in NSSA areas*, 320–323
 - external type 2 route LSA*, 311–318
 - external type 2 route metrics*, 311–318
 - redistribute command reference*, 305–306
 - redistributing into OSPF as E1 routes*, 318–319
 - setting OSPF metrics*, 310–311
- route maps, 332–334
 - configuring metric settings*, 339–341
 - filtering redistributed routes*, 334–339
 - redistribution filtering via distribute-list command*, 343
 - verifying metric settings*, 341–342
- route summarization**
 - EIGRP route summarization
 - auto-summary*, 124–126
 - benefits/trade-offs*, 120
 - calculating summary routes*, 116
 - choosing where to summarize routes*, 116–117
 - configuring*, 120–124
 - influencing summary route selection*, 117–118

- suboptimal forwarding*, 118–120
 - summary route design*, 114–115
 - OSPF route summarization, 221, 231
 - manual summarization at ABR*, 232–235
 - manual summarization at ASBR*, 235–236
 - Router LSA (link state advertisements), 179–186, 195
 - routing loops
 - iBGP routing loops, avoiding, 471
 - BGP synchronization*, 475–476
 - iBGP mesh topologies*, 472–475
 - IGP redistribution*, 475–476
 - route redistribution with multiple redistribution points
 - domain loop problems with multiple routing domains*, 349–357
 - preventing domain loops via per-route AD settings*, 350–354
 - preventing domain loops via route-tag filtering using distribute lists*, 355–357
 - preventing domain loops via subnet filtering while redistributing*, 354–355
 - preventing domain loops with AD*, 346–349
 - preventing routing domain loops with higher metrics*, 345
 - routing tables, OSPF route filtering, 230–231
 - RTP (Reliable Transport Protocol), EIGRP updates, 30, 65
- S**
-
- security, branch routing, 653–654
 - show commands, Network LSA (link state advertisements), 187–191
 - SIA (Stuck-In-Active) routes, EIGRP convergence process, 87–88
 - Simulation mode (BEE), 678
 - single-homed Internet design, outbound routing to the Internet, 404–405
 - single-multihomed Internet design, outbound routing to the Internet, 411–412
 - small branches (branch routing), 656–657
 - SPF (shortest path first), OSPF, 142, 211–212
 - Split Horizon, building EIGRP topology tables, 64–67
 - stateful DHCP (Dynamic Host Configuration Protocol), global unicast addressing, 545
 - stateless autoconfiguration
 - global unicast addressing
 - calculating Interface ID via EUI-64*, 547–548
 - finding DNS IP addresses via stateless DHCP*, 548–549
 - NDP router advertisements*, 546–547
 - IPv6 addressing, 561–562
 - stateless DHCP (Dynamic Host Configuration Protocol), finding DNS IP addresses via stateless DHCP, 548–549
 - static configurations, EIGRP neighborships, 43–45
 - static default routes, EIGRP default routing, 127–128
 - static IPv6 addresses, router configurations, 557–559
 - static point-to-point tunnels, 619
 - GRE tunneling, 619–620, 624–625
 - MCT, 619
 - configuring*, 620–623
 - verifying*, 623–624

static routes

- floating static routes, branch routing, 658
- IP SLA, tracking operations to influence routing, 378–381
- static default routes, EIGRP default routing, 127–128
- static IPv6 routes, 599–601

stub routers

- EIGRP convergence process, stub router impact on Query Scope, 84–86
- OSPF, 221, 239

stubby areas

- NSSA, 240–241, 248–250
- OSPF stubby areas, 239
 - configuring*, 241–243
 - summary table*, 250
 - totally stubby areas*, 246–248
 - types of*, 240–241
 - verifying*, 243–246
- totally NSSA, 240–241
- totally stubby areas, 240–241

Study mode (BEE), 678**subnetting**

- global unicast addressing, 541–543
- practicing, 677–678

successor routes, EIGRP convergence process, 78–80**Summary LSA (link state advertisements), 191–196****summary routes**

- EIGRP convergence process, summary route impact on Query Scope, 86–87
- EIGRP route summarization
 - auto-summary*, 124–126
 - benefits/trade-offs*, 120
 - calculating summary routes*, 116

choosing where to summarize routes, 116–117

configuring, 120–124

influencing summary route selection, 117–118

suboptimal forwarding, 118–120

summary route design, 114–115

OSPF route summarization, 221, 231

manual summarization at ABR, 232–235

manual summarization at ASBR, 235–236

synchronization (BGP), iBGP routing loops, 475–476**T****tables**

binary-to-decimal conversion tables, 702–703

conversion tables, 701–703

design review tables (CCNP Route exams), 12

EIGRP IP routing tables, building, 69–78

bandwidth, 72–75

calculating FD/RD metrics, 69–72

configuring, 72–76

delays, 72–73

k-values, 75–76

metric tuning, 72–78

Offset Lists, 76–78

feature summary table (OSPF), 149

hex-to-decimal conversion tables, 701

implementation plan tables, 13

memory tables, CCNP Route exam preparation, 675

neighbor tables, IPv6 addressing, 561

peer review tables, 12

summary tables (stubby areas), 250

topology tables, building

- contents of update messages, 61–64*
- seeding topology tables, 60*
- Split Horizon, 64*
- Split Horizon defaults on Frame Relay multipoint subinterfaces, 65–67*
- update process, 64–65*
- WAN bandwidth control, 67–69*
- WAN issues for topology exchanges, 65–69*

tunneling comparison table, 617

verification plan tables, choosing commands for, 13

tests. See CCNP Route exams

topology tables

- EIGRP topology tables, building

 - contents of update messages, 61–64*
 - seeding topology tables, 60*
 - Split Horizon defaults on Frame Relay multipoint subinterfaces, 65–67*
 - update process, 64–65*
 - WAN bandwidth control, 67–69*
 - WAN issues for topology exchanges, 65–69*

totally NSSA (not-so stubby areas), 240–241

totally stubby areas, 240–241, 246–248

tunneling

- general concepts, 612–614
- GRE tunnels, configuring in VPN, 666–667
- IPSec tunnels, branch routing, 654–655
- point-to-multipoint tunnels, 615–616, 626–640

- point-to-point tunnels, 614–615
 - GRE tunneling, 619–620, 624–625, 658–659*
 - MCT, 619–624*
 - static point-to-point tunnels, 619–626*
- tunneling comparison table, 617
- virtual tunnel interfaces, 664

U

unicast IPv6 addressing, 550, 553

- link local unicast addresses, 551–552
- unique local addresses, 551

unique local IPv6 addresses, 551

updates

- BGP updates
 - outbound routing to the Internet, 410–411*
 - update messages, table verification, 436–437*
- CCNP Route exams, 705–706
- EIGRP
 - update messages, 30, 61–64*
 - update process, 64–65*
- IPv6 addressing, routing protocol updates, 573
- LSU, 200
 - outbound routing to the Internet, full BGP updates, 410–411

V

verifying

- EIGRP for IPv6, 584–587
- feasible successors (EIGRP convergence process), 78–80

Hello/Hold timers, 34–36

iBGP, 463–464

verifying neighbors, 430–434

IP SLA, 376–377

MCT, 623–624

OSPF

OSPF over multipoint Frame Relay operations, 274–282

OSPFv3, 592–595

Path Control, 430–434, 463–464

RIPng, 578–580

route redistribution

EIGRP, 302–305

metric settings, 341–342

route selection, verifying cost settings, 213–214

stubby areas, 243–246

verification planning (CCNP Route exams), 9

styles of plans, 13–14

typical elements of, 16

verification plan tables, choosing commands for, 13

virtual links, 264–265

virtual links (OSPF)

concepts of, 260–262

configuring

authentication, 265–267

without authentication, 262–264

verifying, 264–265

virtual tunnel interfaces, 664

VPN (virtual private networks)

branch routing, 667–669

DMVPN, 664

GET VPN, 665

GRE tunnels, configuring, 666–667

IPsec VPN, configuring, 665–666

virtual tunnel interfaces, 664

W - X - Y - Z

WAN (wide area networks)

EIGRP

IP routing tables, building, 73–75

neighborships, 48–51

topology tables, building, 65–69

OSPF neighborships, WAN adjacencies, 162–169

Weight (BGP Path Control), influencing

administrative Weight, 500–501, 504–506

AS_Path length Internetwork example, 508–511

local preferences, setting, 507–508

Local_Pref Internetwork example, 508–511

route maps, 511–513

sample Internetworks, 501–504

setting via

neighbor weight command, 506–507

route maps, 504–506

WLLA (BGP best path memorization tip), 500